

UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

ON ANALYSIS AND OPTIMIZATION OF LOW EARTH ORBIT SATELLITE
CONSTELLATIONS

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By CADEN JAMES MATTHEWS

Norman, Oklahoma

2026

ON ANALYSIS AND OPTIMIZATION OF LOW EARTH ORBIT SATELLITE
CONSTELLATIONS

A THESIS APPROVED FOR THE
SCHOOL OF AEROSPACE AND MECHANICAL ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. Diogo Merguizo Sanchez, Chair

Dr. Iman Ghamarian

Dr. Prakash Vedula

©Copyright by CADEN JAMES MATTHEWS 2026
All Rights Reserved.

Abstract

This thesis presents a comprehensive framework for analyzing, modeling, and optimizing large-scale Low Earth orbit (LEO) satellite constellations, utilizing Starlink as a case study. By reframing constellation behavior as a dynamic programming problem, the research effectively decomposes high-dimensional data, derived from historical Two-Line Elements (TLE), into Reduced Orbital Parameter Vectors (ROPVs). These vectors are embedded into low-dimensional submanifolds using Uniform Manifold Approximation and Projection (UMAP) and clustered using Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) to identify distinct “satellite families.” Within each family, critical physical parameters are inferred and validated amidst atmospheric uncertainties, informed by dynamic thermospheric models and solar–geomagnetic indices, enhanced by deep learning-based solar radio flux forecasting. Furthermore, a high-fidelity orbital simulation is integrated with physics-informed reduced order modeling techniques to parameterize the low-thrust performance of krypton Hall-effect thrusters, enabling satellite constellation multi-objective optimization for phasing, station-keeping, and propellant allocation.

Contents

Abstract	iv
1 Introduction, Motivation, and Scope	1
1.1 Low Earth Orbit Satellite Constellations	1
1.1.1 LEO Constellations as Coupled Dynamical Systems	1
1.1.2 LEO Constellation Taxonomy	6
1.1.3 Canonical Geometries and Coverage Theory	10
1.1.4 State Representations for Constellation Analysis	14
1.1.5 LEO Perturbation Environment	18
1.1.6 Secular Evolution, Differential Drift, and Shell Maintenance	23
1.1.7 Relative Motion, Phasing, and Station-Keeping	25
1.1.8 Conjunctions, Debris Feedback, and Disposal Constraints	28
1.1.9 Modeling Hierarchy, Public Data Products, and the Scientific Gap	30
1.2 Starlink Satellite Constellation	34
1.2.1 Gen 1 Starlink Case Study	34
1.2.2 Orbital Architecture	36
1.2.3 Catalog-Space Geometry	37
1.2.4 Aggregate Historical Morphology	41
1.2.5 Lifecycle Corridors	45
1.2.6 Secular Dynamics and Drift	52
1.2.7 Maturity, Turnover, and Inferred Operational State	55
2 High-Fidelity Orbital Simulator Development	59
2.1 Propagation Architecture, State Definition, and Numerical Strategy	59
2.1.1 Simulation Scope and Propagated State	59
2.1.2 Frames, Epochs, and Time Variables	60
2.1.3 Adaptive Integration and Event Handling	61
2.2 Initial-Condition Construction and Multi-Satellite Case Assembly	62

2.2.1	Manual and Catalog-Driven Initialization	62
2.2.2	Epoch-Conditioned Major-Body Initialization	64
2.2.3	Single-Satellite and Ensemble Case Assembly	64
2.3	Earth Gravity and Conservative Perturbation Modeling	65
2.3.1	Central Gravity and Truncated Spherical Harmonics	66
2.3.2	Third-Body Gravitational Forcing	67
2.3.3	Conservative-Force Assembly in the Propagated Equations	68
2.4	Solar Radiation Pressure, Eclipse Geometry, and Effective Optical Response	68
2.4.1	Solar Radiation Pressure Force Formulation	69
2.4.2	Eclipse Geometry and Partial Shadowing	69
2.4.3	Effective Optical Response and Projected-Area Surrogates	70
2.5	Thermospheric Drag Coupling and Effective Aerodynamic State	71
2.5.1	Relative Velocity and Drag Acceleration	71
2.5.2	Static Reference Atmosphere	72
2.5.3	Dynamic Atmosphere Interface	73
2.5.4	Effective Aerodynamic State	74
2.6	Phase-Conditioned Low-Thrust Scheduling and Variable-Mass Dynamics	74
2.6.1	Maneuver Phase Taxonomy and Schedule Construction	75
2.6.2	Along-Track Thrust Direction and Mass-Depletion Law	76
2.6.3	Segment Boundaries and Coast Intervals	77
2.6.4	Thrust-Integrated Diagnostics	78
2.7	Parallel Batch Propagation and Output Products	79
2.7.1	Ensemble Execution and Worker Orchestration	79
2.7.2	Stored State and Orbital Element Products	80
3	Data Decomposition and Constellation Clustering	81
3.1	Constellation Decomposition as a Family-Discovery Problem	81
3.1.1	Large Constellations Require Family-Wise Decomposition	81
3.1.2	Satellite Families	82
3.1.3	Validation Logic for Labeled, Semi-Labeled, and Unlabeled Cases	84
3.2	Public Data and Orbit-Geometry Representation	85
3.2.1	TLE-Derived Public Data and Epoch Synchronization	86
3.2.2	Limitations of Raw Classical Elements in Dense Constellations	87
3.2.3	Reduced Orbital Parameter Vector	88
3.2.4	Local Metric Geometry	90

3.3	Neighborhood-Preserving Embedding and Density-Based Family Extraction	91
3.3.1	Neighborhood-Preserving Latent Charting	91
3.3.2	Trustworthiness	94
3.3.3	Hierarchical Density Extraction	96
3.3.4	Density-Based Clustering Validation	97
3.4	Positive-Control Recovery of Known Architecture	99
3.4.1	GPS as an External Benchmark	100
3.4.2	OneWeb as an Architecture-Conditioned Benchmark	104
3.5	Starlink as a Shell-Conditioned Discovery Problem	108
3.5.1	Shell-Conditioned Analysis	109
3.5.2	The 53.0° Shell	110
3.5.3	The 53.2° Shell	111
3.5.4	The 70.0° Shell	111
3.5.5	The 97.6° Shell	112
3.5.6	Cross-Shell Comparison	112
3.6	Global Satellite Synthesis	116
3.6.1	Local-to-Global Remapping	116
3.6.2	Global UMAP Structure	117
3.6.3	Centroid-Level Feature Diagnostics	119
3.6.4	Cluster Orbital Element Diagnostics	121
3.7	Interpretation, Limits, and Methodological Boundaries	122
3.7.1	Latent Space	123
3.7.2	TLE and SGP4 Data Products	123
3.7.3	Hyperparameter Sensitivity, Stochasticity, and Validation	124
3.7.4	Methodological Limitations	125
4	Forecast-Conditioned Thermospheric Forcing	126
4.1	Forecast-Conditioned Thermospheric Forcing for Future Satellite Propagation . . .	126
4.1.1	Future Atmospheric Forcing	126
4.1.2	Scope of the Present Formulation	128
4.2	Space-Weather Driver Data	128
4.2.1	The Solar Radio Flux Driver Family	129
4.2.2	Geomagnetic Activity Drivers	130
4.2.3	Historical Driver Archive	131
4.3	Harmonic Phase Encoding	133

4.3.1	Annual Harmonic Basis	133
4.3.2	Solar-Cycle Harmonic Basis	134
4.3.3	Construction of the Forecast State	137
4.4	Hybrid ResNeXt-iTransformer Forecasting Architecture	138
4.4.1	Forecast State, Sequence Length, and Targets	138
4.4.2	MultiScale ResNeXt Backbone	140
4.4.3	Inverted Transformer Encoder	142
4.4.4	Output Head and Recursive Rollout	144
4.5	Training Protocol and Forecast Validation	145
4.5.1	Temporal Split and Cross-Validation	145
4.5.2	Optimization Configuration	146
4.5.3	Fold-Level Forecast	148
4.6	Residual Structure	150
4.6.1	Residual Distribution	151
4.6.2	Horizon Residual Behavior	152
4.6.3	Spectral Content	153
4.7	Multi-Year Forecasts of Solar and Geomagnetic Drivers	155
4.7.1	Forecast of Observed Solar Radio Flux	155
4.7.2	Forecast of Adjusted Solar Radio Flux	156
4.7.3	Forecast of URSI-D Solar Radio Flux	157
4.7.4	Conditional Stochastic Forecast of a_p	158
4.8	NRLMSIS 2.0 Coupling	160
4.8.1	Mapping Forecast Drivers	160
4.8.2	Density Products for Future Satellite Propagation	163
5	Physical Parameter Estimation via Bayesian Inference	165
5.1	TLE Decay Observables	165
5.1.1	Late-Life Satellite Decay as an Inverse Problem	165
5.1.2	Semi-Major Axis from TLE Mean Motion	166
5.1.3	Anchored Residuals and the Inference-Prediction Split	167
5.1.4	Retained Satellite Decay Arcs	167
5.2	Forecast-Conditioned Forward Model for Satellite Drag Simulation	169
5.2.1	Equations of Motion and Retained Perturbations	169
5.2.2	NRLMSIS Density Realization	170
5.2.3	Offline Density Grids and Runtime Interpolation	171

5.2.4	Terminal Event and Propagation Horizon	172
5.3	Effective Ballistic-Coefficient Parameterization	172
5.3.1	Effective Drag Representation	173
5.3.2	Random-Tumbling Effective-Area Surrogate	173
5.3.3	Identifiability Under TLE-Derived Decay	174
5.3.4	Effective Parameterization	175
5.4	Bayesian Inference and MCMC Implementation	175
5.4.1	Likelihood on Anchored Semi-Major-Axis Residuals	175
5.4.2	Prior on the Inferred Aerodynamic State	176
5.4.3	Random-Walk Metropolis and Warmup Adaptation	177
5.4.4	Diagnostic Criteria and Archive Limitations	178
5.5	Per-Satellite and Constellation-Level Ballistic Coefficients	179
5.5.1	Per-Satellite Posterior Summaries	180
5.5.2	Conditional Posterior Concentration	181
5.5.3	Cross-Satellite Heterogeneity	182
5.5.4	Constellation-Level Posterior Structure	182
5.6	Posterior-Predictive Satellite Simulation and Re-Entry Uncertainty	184
5.6.1	Propagation of Posterior Draws	184
5.6.2	Semi-Major Axis Trajectory Fan	185
5.6.3	Re-Entry Timing Uncertainty	186
6	Physics-Constrained System Identification of Low-Thrust Dynamics	187
6.1	Mean-Element Observation Model	187
6.1.1	Catalog Mean-Element Observation Model	187
6.1.2	Phase-Conditioned Effective Parameterization	189
6.1.3	Inference Scope and Effective Interpretation	191
6.2	Data Construction, Segmentation, and Observable Structure	192
6.2.1	TLE Ingestion, Phase Intervals, and Trajectory-Arc Construction	192
6.2.2	Phase Imbalance and Segment-Duration Structure	195
6.2.3	Observable Channels for Stage A and Stage B	197
6.3	Reduced Dynamical and Propulsion Model	198
6.3.1	Mean-Element Trajectory and Phase-Conditioned Net Acceleration	198
6.3.2	Krypton Beam Relations and Effective Propulsion Variables	202
6.3.3	Surrogate Chemistry and Feasibility Regularization	203
6.4	Trajectory Estimation and Learned Phase Structure	205

6.4.1	Composite Trajectory Objective and Weighting Structure	205
6.4.2	Optimization Behavior and Convergence	207
6.4.3	Learned Global and Phase-Level Parameter Evolution	208
6.5	Trajectories, Residual Structure, and Failure Modes	211
6.5.1	Aggregate Fit Quality and Phase-Resolved Error Hierarchy	211
6.5.2	Observed-Predicted Structure and Residual Geometry	212
6.5.3	Robustness to Timing and Model Perturbation	213
6.6	Posterior Inference and Identifiability	215
6.6.1	Reduced Posterior Parameterization and SNPE Setup	215
6.6.2	Posterior Marginals and Compensating Parameter Structure	217
6.6.3	Thruster Performance	219
6.6.4	Practical Identifiability and Low-Rank Structure	220
7	Recommendations, Discussion, and Implications	223
7.1	Integrated Findings Relative to Initial Hypotheses	224
7.2	Importance and Applications for Satellite Constellations	226
7.3	Limitations and Sources of Uncertainty	228
7.4	Strategic Recommendations	230
8	Conclusion	233
8.1	Concluding Synthesis	233
8.2	Future Research Directions	235
A	Code	257

List of Figures

1.1	Low Earth Orbit Satellite Coupled Dynamical System	2
1.2	Classical Constellation to Large LEO Fleet Comparison	4
1.3	Lower Operational LEO and Higher Operational LEO Comparison	9
1.4	Walker Delta and Walker Star Comparison	10
1.5	Canonical Satellite Constellation Families	12
1.6	Single-Satellite Visibility Geometry	14
1.7	Local Orbital Frames	18
1.8	Perturbation Hierarchy for LEO Propagation	19
1.9	Differential J_2 -Driven Nodal Regression and Shell-Plane Shear	25
1.10	Propagation Hierarchy for LEO Propagation	31
1.11	Combined Three-Dimensional Geometry of the Principal Gen 1 Starlink Divisions .	38
1.12	Combined RAAN versus Mean Anomaly Geometry of the Principal Gen 1 Starlink Divisions	39
1.13	Combined Orbital Vector Representations of the Principal Gen 1 Starlink Divisions	41
1.14	Aggregate Inclination of the Gen 1 Starlink Satellite Constellation	42
1.15	Aggregate Semi-Major Axis of the Gen 1 Starlink Satellite Constellation	43
1.16	Aggregate Eccentricity of the Gen 1 Starlink Satellite Constellation	45
1.17	Inclination versus Semi-Major Axis for the 53° Complex	47
1.18	Inclination versus Time of the 53° Complex	48
1.19	Inclination versus Semi-Major Axis for the 70° Division	49
1.20	Inclination versus Time of the 70° Division	50
1.21	Inclination versus Semi-Major Axis for the 97.6° Division	51
1.22	Inclination versus Time of the 97.6° Division	52
1.23	Grouped First-Order J_2 Nodal Precession Rates of the Principal Gen 1 Starlink Divisions	53
1.24	Grouped First-Order J_2 Apsidal Precession Rates of the Principal Gen 1 Starlink Divisions	55

1.25	Shell Compactness/Occupancy Spread Chronology	56
1.26	Sustainability-State Counts Chronology	57
1.27	Disposal-Onset Chronology	57
3.1	Conceptual Workflow for Constellation Decomposition as Satellite Family Discovery	85
3.2	Reduced Orbital Parameter Vector Structure for the Global Starlink Population . . .	89
3.3	UMAP-HDBSCAN Grid Search Space	93
3.4	Trustworthiness Evaluation Workflow	95
3.5	Density-Based Clustering Validation Workflow	98
3.6	GPS Condensed Tree	101
3.7	GPS UMAP Clusters	102
3.8	GPS Orbital Elements Pairplot	103
3.9	OneWeb Condensed Tree	105
3.10	OneWeb UMAP Clusters	106
3.11	OneWeb Orbital Elements Pairplot	107
3.12	Starlink Condensed Tree	113
3.13	Starlink UMAP Clusters	114
3.14	Starlink Orbital Elements Pairplot	115
3.15	Global Starlink UMAP Clusters	118
3.16	Centroid-Level Correlation for the Starlink Satellite Constellation	120
3.17	Global Starlink Orbital Elements Pairplot	122
4.1	Historical 10.7 cm Solar Radio Flux	129
4.2	Historical Daily Sunspot Number	132
4.3	Seasonal Fourier Harmonic Channels	134
4.4	Solar Cycle Fourier Harmonic Channels	136
4.5	Model Architecture	140
4.6	ResNeXt-101 Backbone	141
4.7	iTransformer Encoder	143
4.8	Training Loss	147
4.9	Cross-Validation MAE by Fold	149
4.10	Cross-Validation RMSE by Fold	149
4.11	Histogram of Prediction Residuals	151
4.12	Mean Forecast Residual over the Prediction Horizon	152
4.13	Lomb–Scargle Diagnostics of 10.7 cm Solar Radio Flux	154

4.14	Lomb–Scargle Diagnostics of 10.7 cm Solar Radio Flux Residuals	154
4.15	Observed 10.7 cm Solar Radio Flux Forecast	156
4.16	Adjusted 10.7 cm Solar Radio Flux Forecast	157
4.17	URSI-D 10.7 cm Solar Radio Flux Forecast	158
4.18	Conditional Stochastic a_p Forecast	159
5.1	Retained Late-Life Semi-major Axis Histories	168
5.2	Starlink Ballistic Coefficients Forest	181
5.3	Pairwise Wasserstein-1 Distances between Per-Satellite Posterior Distributions . . .	183
5.4	Posterior-Predictive Semi-Major Axis Trajectories	185
6.1	Mission Phases for Gen 1 Starlink TLE History	190
6.2	Segment Counts by Maneuver Phase	195
6.3	Distribution of Segment Duration and Observed Semi-Major Axis Increment	196
6.4	Stage A Log-Loss Decomposition	207
6.5	Stage A Trajectory Training	207
6.6	Epoch-by-Epoch Evolution of the Phase-Conditioned Parameters	209
6.7	Final Effective Phase Parameters	210
6.8	Distribution of Trajectory Path RMSE by Phase	212
6.9	Observed versus Predicted Endpoint Structure	213
6.10	Sensitivity of Validation RMSE to Uniform Timing Shifts	214
6.11	Posterior Correlation Structure for Reduced Stage B Parameters	218
6.12	Singular Values of the Local Stage B Sensitivity Operator	221

List of Tables

5.1	Starlink Decay Arcs for Ballistic Coefficient Inference	168
5.2	Posterior Summaries and Monte Carlo Diagnostics	179
5.3	Per-Satellite Posterior Summaries	180
6.1	Effective Phase-Resolved Thruster Performance	219

Chapter 1

Introduction, Motivation, and Scope

1.1 Low Earth Orbit Satellite Constellations

1.1.1 LEO Constellations as Coupled Dynamical Systems

Low Earth orbit (LEO) constellations may be represented as distributed dynamical systems, whose collective evolution arises from shared perturbative forcing and coordinated operational behavior [1, 2, 3]. The emergence of large LEO fleets underscores the importance of this system-level representation, as population growth, replenishment, conjunction exposure, and disposal reliability become coupled architectural properties rather than isolated spacecraft attributes [4, 5]. Although geometric coverage and telemetry are fundamental to constellation design and operation, those components provide only a partial description of the underlying dynamics. Fleet integrity depends on maintaining shell coherence under secular drift and atmospheric drag, preserving permissible along-track separations, and ensuring timely disposal before debris-persistence thresholds are exceeded in adjacent altitude bands [6, 7, 8, 9, 10]. The present analysis, therefore, focuses on the coupled evolution of multiple satellites within a shared orbital environment.

A constellation may be represented as an engineered multi-satellite architecture with state variables distributed across N members. Let the state of satellite s at epoch t be denoted by $\mathbf{x}_s(t)$. The aggregate constellation state is then defined as

$$\mathbf{X}(t) = \begin{bmatrix} \mathbf{x}_1(t)^\top & \mathbf{x}_2(t)^\top & \cdots & \mathbf{x}_N(t)^\top \end{bmatrix}^\top \quad (1.1)$$

At a schematic level, the evolution of the satellite fleet can be written as

$$\dot{\mathbf{X}} = \mathbf{F}(\mathbf{X}, t; \Theta_{\text{env}}, \Theta_{\text{veh}}) + \mathbf{G}(\mathbf{X}, t)\mathbf{U} + \mathbf{W}(t) \quad (1.2)$$

where Θ_{env} denotes the environmental parameter set governing shared perturbative forcing, including geopotential harmonics, thermospheric density, solar radiation pressure (SRP), and third-body accelerations. For satellite s , let the vehicle-specific parameter set be written as

$$\Theta_{\text{veh},s} = \{m_s, A_{D,s}, C_{D,s}, A_{\text{SRP},s}, C_{R,s}, T_{\text{max},s}, I_{\text{sp},s}, m_{p,s}\}, \quad s = 1, \dots, N \quad (1.3)$$

where Θ_{veh} denotes the vehicle-specific parameter set governing satellite-to-satellite differences in drag response, SRP response, and propulsion capability, including satellite mass (m_s), effective drag area ($A_{D,s}$), drag coefficient ($C_{D,s}$), effective SRP area ($A_{\text{SRP},s}$), reflectivity coefficient ($C_{R,s}$), available thrust magnitude ($T_{\text{max},s}$), specific impulse ($I_{\text{sp},s}$), and remaining propellant mass ($m_{p,s}$). The fleet-level vehicle parameter set is then represented as

$$\Theta_{\text{veh}} = \{\Theta_{\text{veh},1}, \Theta_{\text{veh},2}, \dots, \Theta_{\text{veh},N}\} \quad (1.4)$$

The aggregate control input $\mathbf{U}$ and disturbance term $\mathbf{W}(t)$ account for maneuver execution and model uncertainty, respectively [1, 2, 3]. This formulation highlights the principal coupling channels within large satellite constellations: shared environmental forcing, vehicle heterogeneity, shell geometry, maneuver strategy, and disposal protocol.




Figure 1.1: Low Earth Orbit Satellite Coupled Dynamical System

As depicted in Figure 1.1, the analysis is conducted on three hierarchical scales: single-satellite evolution under non-Keplerian perturbations, emergent orbital shell and plane geometry

resulting from collective trajectories, and fleet-level control under uncertainty [1, 2, 3, 11]. Uniform perturbations manifest differently at each scale. Earth’s oblateness induces secular precession and, in the absence of control, shearing of plane structures. Atmospheric drag decreases orbital energy and influences shell cohesion through inter-satellite ballistic variability [2, 12]. The satellite constellation thus serves as a structured many-body system in which shared forcing and differential responses jointly determine geometry and operational complexity.

The translational dynamics associated with this structure can be described by the perturbed equation of motion (EOM) as

$$\ddot{\mathbf{r}} = -\frac{\mu_{\oplus}}{r^3}\mathbf{r} + \mathbf{a}_{J_2} + \mathbf{a}_{\text{drag}} + \mathbf{a}_{\text{SRP}} + \mathbf{a}_{3b} + \cdots \quad (1.5)$$

At this stage, the equation serves only as a formal scaffold, with its constituent terms detailed in subsequent sections [1, 2]. Shell coherence requires bounded differential responses to perturbations, while along-track structure depends on matched secular drift rates. Effective disposal outcomes are determined by the interplay among orbital altitude, atmospheric drag response, and control authority [12, 8, 10]. Equation (1.5) marks the transition from a purely geometric description to a dynamical characterization of satellite constellations.

From Classical Constellation Geometry to LEO Fleets

Classical constellation design originated as a coverage-geometry problem, in which the principal objective was to achieve a desired coverage pattern with the fewest possible space vehicles. Walker’s canonical framework represented constellations as sets of equal-period circular orbits coordinated through inter-plane phasing and orbital-plane symmetry to produce continuous global coverage [13]. The central design variables were the total number of satellites (N), the number of orbital planes (P), the inclination (i), and the inter-plane phasing parameter (f). Core analytical outputs included coverage bounds, revisit intervals, and symmetry classifications. Under this interpretation, a constellation was treated primarily as a geometric structure whose operational function followed from its spatial configuration [13, 14].

Large LEO fleets substantially alter both the scale and the operational logic of constellation design. These systems are characterized by multiple orbital shells, frequent deployments, continuous replenishment, real-time conjunction mitigation, and stringent disposal protocols [4, 5, 7, 8, 10]. The primary design emphasis, therefore, shifts from purely coverage-driven satellite placement to the management of orbital volume, suppression of shell drift, mitigation of conjunction risk, and prompt re-entry of failed satellites to limit the growth of long-lived debris.

This transition is both qualitative and quantitative, as Figure 1.2 summarizes the progression from classical coverage constellation geometry to continuously managed LEO fleet dynamics. Classical constellations were commonly treated as quasi-static geometric patterns subjected to limited perturbative correction. Contemporary LEO fleets operate as persistently controlled dynamical systems [1, 2, 3]. Individual satellites pass through overlapping deployment, phasing, shell transfer, station-keeping, collision avoidance, disposal, and replacement cycles, so the constellation architecture must be understood as a temporally sustained orbital process. Large LEO fleets are therefore more naturally described through mean-element evolution, relative-motion structure, and shell-level dynamical models [11, 15, 3].



Figure 1.2: Classical Constellation to Large LEO Fleet Comparison

Limitations of the Keplerian Description

The Keplerian paradigm forms the foundation of classical orbital mechanics; however, its descriptive capacity is limited for large LEO satellite constellations [1, 2, 3]. In LEO, secular precession due to Earth’s oblateness and persistent semi-major axis decay from atmospheric drag produce continuous and significant evolution even in orbits that would otherwise be regarded as stable [2, 12]. Within large satellite constellations, differential evolution of orbital elements across satellites, amplified by small deployment offsets and control anomalies, leads to pronounced drift in the right ascension of the ascending node (RAAN), divergence in mean motion (n), and phase decoherence. These effects can collectively degrade the geometry of a satellite constellation over time [1, 2, 11].

Relative geometry serves as a principal descriptor of satellite constellation dynamics. Metrics such as slot occupancy, phase separation, collision margin, and maneuver coordination are inherently collective and cannot be captured through isolated single-satellite analyses. Formation-flying theory provides the appropriate mathematical framework, enabling long-horizon characterization of perturbed evolution through mean and relative orbital elements [11, 15, 3]. The same perspec-

tive extends naturally from compact formations to large satellite constellations, where the primary controlled quantity is the admissible relative geometry under shared perturbative forcing.

The operational cadence of large LEO satellite constellations introduces additional analytical demands. Conjunction assessment becomes a persistent apparatus, while collision screening, maneuver allocation, deployment, and disposal must be managed as continuous system processes. System-level observables such as conjunction burden, maneuver frequency, disposal robustness, and risk rate therefore emerge as ensemble quantities that require explicit system-theoretic analysis [6, 7, 8, 10].

Data provenance imposes equally important constraints, as the reliability and completeness of orbital data fundamentally shape the scope of any subsequent analysis. Public orbital archives lack continuous truth-state records, covariance matrices, maneuver logs, and fully specified force-model histories. Two-Line Elements (TLEs) are mean-element catalog products interpreted through simplified perturbation models such as Simplified General Perturbations 4 (SGP4) [16, 17, 18, 19]. Current Consultative Committee for Space Data Systems (CCSDS) conventions formalize these interpretive constraints by prescribing explicit metadata for reference frames, time systems, and mean-element message content [20]. These limitations directly shape the admissible analytical approach, because public catalog products constitute the primary empirical basis for the present work. Rigorous inference is possible when the observational model remains consistent with the semantics and modeling assumptions of the underlying catalog data [16, 17, 20, 18].

Scope

This analysis adopts a focused disciplinary perspective while addressing a broad range of dynamical phenomena, including orbital mechanics, perturbation theory, relative dynamics, station-keeping, conjunction assessment, disposal mechanics, and uncertainty quantification within large LEO satellite constellations [1, 2, 3, 7, 8]. Core analytical themes include shell structure, shell coherence, inter-satellite geometry, secular evolution, and the extraction of interpretable physical insight from public orbital data. Communications payload performance, routing, and radio frequency (RF) design are excluded, except where orbital constraints arise through target altitude, inclination, local-time configuration, plane count, or through effects such as attitude-dependent drag and solar radiation pressure.

The structure of this thesis reflects this analytical scope, as modeling large LEO satellite constellations at the individual-satellite level is computationally prohibitive. To address this challenge, the clustering framework maps TLE-derived orbital histories into reduced orbital representations that support identification of latent satellite families through manifold learning and density-based

clustering. The ballistic-coefficient inference framework complements this approach by treating TLE-based decay histories as observables and probabilistically estimating effective drag with a high-fidelity force model. Finally, the inverse problem of thruster parameters completes this framework by extracting effective propulsion signatures and maneuver-phase behavior from public data. Together, these methods define a unified workflow in which public orbital data reveal latent structure, support family-wise physical inference, and enable uncertainty-aware, high-fidelity propagation at scales relevant to contemporary LEO satellite constellations [16, 17, 20, 18].

The following subsections develop perturbation theory, frame conventions, shell-evolution laws, conjunction formalisms, and disposal mechanics in greater detail. Within this perspective, the constellation is treated as an engineered orbital architecture subject to shared external forcing, differential vehicle response, and collective operational constraints. The analysis proceeds across three hierarchical scales: single-satellite propagation, shell-and-plane geometry, and fleet-level control under uncertainty. The principal challenges at these scales include differential drift, conjunction throughput, and disposal reliability, which are addressed through structure discovery in public data, family-wise physical inference, and uncertainty-aware dynamical modeling.

1.1.2 LEO Constellation Taxonomy

The altitude regime provides the primary taxonomic axis because atmospheric drag, secular precession, orbital lifetime, debris persistence, and post-mission disposal depend strongly on orbital height within the LEO protected region [1, 2, 12, 21]. Orbital architecture provides a secondary classification axis, characterized by shells, planes, slots, and phasing arrangements [13, 14, 22]. The operational class provides a tertiary distinction, separating constellations into polar, inclined, sun-synchronous, repeating-ground-track, and near-circular regimes. These dimensions are coupled: changes in altitude introduced to improve coverage, reduce latency, or modify revisit cadence directly alter nodal regression, mean motion, aerodynamic lifetime, and end-of-life removal behavior [1, 2, 23].

Altitude Regime and Internal Structure

LEO is commonly treated as the region below approximately 2000 km, although the perturbation environment within this region is strongly stratified [21, 10, 1]. The transition from rapid, drag-assisted object removal to prolonged orbital persistence occurs well below the upper LEO boundary, motivating a more resolved internal classification. Very Low Earth Orbit (VLEO) is defined here as the region below approximately 450 km, where aerodynamic interaction is strong, drag

compensation becomes a primary design constraint, and orbital lifetimes are short compared with higher LEO bands [23]. The 450 to 650 km band supports many broadband satellite constellations and remains highly sensitive to atmospheric drag, with direct consequences for station-keeping, collision avoidance, and end-of-life disposal [4, 5, 12]. Between 650 and 1200 km, reduced aerodynamic forcing leads to longer debris-retention durations and greater dependence on reliable post-mission disposal. The upper LEO region, from 1200 to 2000 km, is characterized by extended debris persistence and weak passive self-clearing on operational timescales [21, 7, 8, 9].

Atmospheric drag is the dominant non-conservative perturbation throughout much of the LEO region. Secular nodal precession produced by Earth’s oblateness depends strongly on both semi-major axis and inclination. For near-circular LEO orbits, the leading-order secular rate of right ascension of the ascending node (RAAN) drift induced by the J_2 perturbation may be formulated as

$$\dot{\Omega}_{J_2} \approx -\frac{3}{2} J_2 n \left(\frac{R_{\oplus}}{a} \right)^2 \cos i \quad (1.6)$$

where $n = \sqrt{\mu_{\oplus}/a^3}$ [1, 2]. Atmospheric drag modifies the semi-major axis, which in turn modulates mean motion, orbital lifetime, and phasing behavior. The J_2 perturbation drives rapid variation in RAAN as a function of semi-major axis and inclination, while drag produces altitude decay, mean-motion variation, and phase dispersion over operationally relevant timescales [1, 2, 12]. At altitudes below roughly 600 km, atmospheric drag can materially assist passive removal of failed satellites and debris. Above roughly 700 to 800 km, passive removal becomes increasingly sensitive to ballistic coefficient, solar-cycle state, and atmospheric-density uncertainty. Near and above 1000 km, orbital residence times may extend far beyond ordinary mission and regulatory planning horizons [12, 21, 7, 8].

Shells, Planes, Slots, and Phasing

A shell is defined as a group of orbits with nearly identical altitude and inclination, forming a dynamically coherent layer of objects with similar geometric and operational behavior. In operational satellite constellations, shells are commonly implemented as Walker-type circular patterns further divided into multiple planes with prescribed phasing [13, 14]. The shell serves as the primary unit for deployment, occupancy regulation, and end-of-life disposal, with constraints often enforced at the aggregate shell level as well as at the individual satellite level.

A plane consists of satellites sharing a common inclination and nearly common RAAN, with members distributed primarily by argument of latitude (u) to produce the intended along-track spacing within that orbital band. When Walker notation is invoked explicitly, the shell architecture

is described by the inclination, the total number of satellites T , the number of planes (P), and the inter-plane phasing parameter (F) [13, 14].

A slot is a designated occupancy location within the constellation architecture, defined by separation requirements and by the need to satisfy minimum-distance and relative-motion constraints under perturbations and maneuvers [11, 15, 3].

Phasing operates at two hierarchical levels in constellation geometry. Intra-plane phasing governs along-track separations within a plane, whereas inter-plane phasing establishes offsets between adjacent planes and therefore influences seam geometry, deployment sequencing, and station-keeping protocols [13, 14]. Multi-shell architectures combine several shell-plane-phasing regimes across different altitudes and inclinations, producing a coupled system in altitude, geometry, and control. Repeating-ground-track and flower-constellation design generalize this phasing perspective by explicitly coordinating orbital recurrence, Earth rotation, and satellite distribution across the constellation architecture [22].

LEO Operational Classes

LEO operational classes are defined by the interaction between orbital geometry and the perturbative environment [1, 2]. Polar and near-polar systems provide global access to high latitudes and concentrate nodal crossings and interaction regions there, significantly altering local traffic patterns. Inclined constellations, with inclinations below 90° , are optimized for coverage of densely populated mid-latitude regions. Walker-Delta architectures distribute inclined planes across the full range of RAAN, using phasing to regularize access intervals and reduce congestion in polar crossing regions [13, 14].

Sun-synchronous orbits (SSOs) form a valuable near-polar class, exploiting J_2 -driven nodal regression to synchronize orbital-plane precession with the apparent motion of the Sun. This synchronization preserves an approximately constant local solar time, which is essential for illumination consistency in many remote-sensing missions [1, 2]. Repeating-ground-track (RGT) orbits achieve periodicity by coordinating orbital period, nodal precession, and Earth's rotation, so that the ground track recurs after a prescribed cycle [22, 1]. Near-circular satellite constellations improve operational regularity because low eccentricity sharpens shell definition and limits secular differential drift. Each operational class, therefore, exhibits distinct perturbation sensitivities, recurrence properties, and control requirements within orbital structure [13, 22, 1, 2].

Altitude Regimes

The primary sustainability distinction within operational LEO is between lower shells near 500 to 600 km and higher shells near 1000 to 1200 km. In lower operational LEO, atmospheric drag can materially assist post-mission disposal of defunct satellites on policy-relevant timescales. In higher operational LEO, passive clearance is weaker, so end-of-life removal requires propulsion, drag augmentation, or other active disposal measures [12, 23, 21, 7, 8]. Lower shells are therefore associated with stronger environmental attenuation and shorter debris residence times, whereas higher shells can retain failed satellites and fragments for prolonged periods [6, 21, 9].

Lower operational LEO still contains important hazards because reduced debris persistence is accompanied by higher traffic density, frequent shell crossings, and greater operational throughput. Figure 1.3 provides a comparison of the lower and higher operational LEO regimes. Although stronger atmospheric drag shortens orbital lifetimes, lower shells often host dense traffic, frequent shell crossings, and sustained maneuver activity. In higher operational LEO, slower drag-driven decay increases the long-term consequences of satellite failures and fragmentation events. International debris-mitigation frameworks emphasize post-mission disposal, minimization of breakups, and collision avoidance as core requirements, while current U.S. regulations for FCC-licensed satellites ending their mission in, or passing through, the LEO region require uncontrolled atmospheric reentry disposal as soon as practicable and no later than five years after mission end [21, 7, 8, 9, 10].




Figure 1.3: Lower Operational LEO and Higher Operational LEO Comparison

1.1.3 Canonical Geometries and Coverage Theory

Walker Delta and Walker Star

Walker-type satellite constellations constitute the principal class of equal-period circular-orbit architectures [13, 14, 24]. The canonical Walker configuration is denoted by $i : T/P/F$, where i is the inclination, T is the total number of satellites, P is the number of orbital planes, and F is the inter-plane phasing parameter. In a uniform distribution, the number of satellites per plane is

$$S = \frac{T}{P} \quad (1.7)$$

and the argument-of-latitude offset between corresponding satellites in adjacent planes is

$$\Delta u = F \left(\frac{360^\circ}{T} \right) \quad (1.8)$$

Walker Delta architectures distribute inclined orbital planes uniformly over the full 360° range in RAAN, providing regular coverage for near-global and mid-latitude regions, whereas Walker Star architectures, which form the near-polar subclass, distribute the orbital planes over 180° in RAAN and produce distinctive polar seam geometries, as shown in Figure 1.4 [13, 14, 24].




Figure 1.4: Walker Delta and Walker Star Comparison

Ballard Rosettes

Ballard rosette-type satellite constellations, which predate Walker's formulation, establish a broader symmetry framework for circular multi-plane architectures [14]. The Ballard methodology emphasizes global inertial patterning rather than a compact discrete parameter set. Walker's parameterization may be interpreted as a regular subset within this broader symmetry class. Modern synthesis algorithms frequently recover or extend these foundational geometries [14, 13, 25].

Streets of Coverage


Streets of Coverage (SoC) constellations are designed primarily for polar and near-polar missions that seek continuous single-fold or multi-fold coverage above a specified latitude while minimizing the number of satellites. Coverage paths are phased to maintain persistent observation within the target region, without requiring the same global symmetry constraints that characterize Walker families [26].

This approach is characterized by latitude-constrained continuous access rather than full-Earth symmetry. That distinction produces tradeoffs in satellite count, resilience, deployment strategy, station-keeping, and disposal behavior between Walker and SoC architectures [26].


Flower Constellations

Flower-type satellite constellations expand the design space by incorporating repeating-ground-track (RGT) recurrence directly into constellation construction [22]. Recurrence in the Earth-fixed frame is embedded a priori through integer recurrence parameters and shared orbital elements, generating a repeated ground-track family.


Flower constellations prioritize temporal recurrence and revisit structure in the Earth-fixed frame, whereas Walker and SoC families emphasize instantaneous coverage symmetry. The design objective shifts from spatial coverage at a given epoch to controlled, recurring access at fixed terrestrial locations, which makes Flower constellations a distinct architectural class [22]. Figure 1.5 places these canonical families in a comparative schematic, visually highlighting the key architectural differences and similarities among Walker, Ballard, SoC, and Flower constellations.




Walker Constellation



Ballard Rosette Constellation



Streets-of-Coverage Constellation



Flower Constellation

Figure 1.5: Canonical Satellite Constellation Families

Visibility Geometry and Coverage Metrics

Single-satellite coverage is modeled through a spherical-cap representation on Earth's surface [24, 1, 2]. The governing relations are expressed as

$$\sin \alpha = \frac{R_{\oplus}}{R_{\oplus} + h} \cos \varepsilon_{\min} \quad (1.9)$$

and

$$\beta = \frac{\pi}{2} - \varepsilon_{\min} - \alpha \quad (1.10)$$

Here, $R_{\oplus}$ denotes Earth's reference radius, h the orbital altitude, $r = R_{\oplus} + h$ the orbital radius, and $\varepsilon_{\min}$ the minimum ground elevation angle. The nadir angle is denoted by α , and the geocentric

half-angle of the visible cap by β [24, 1, 2]. Eliminating α gives

$$\beta = \arccos\left(\frac{R_{\oplus}}{R_{\oplus} + h} \cos \varepsilon_{\min}\right) - \varepsilon_{\min} \quad (1.11)$$

where β defines the instantaneous extent of surface visibility for a specified altitude and elevation constraint.

The area of the visible cap is formulated as

$$A_{\text{cap}} = 2\pi R_{\oplus}^2 (1 - \cos \beta) \quad (1.12)$$

and the corresponding fraction of Earth's surface covered is

$$f_{\text{cap}} = \frac{A_{\text{cap}}}{4\pi R_{\oplus}^2} = \frac{1 - \cos \beta}{2} \quad (1.13)$$

The expressions provide the geometric basis for coverage metrics for which system-level quantities are derived either from the visibility cap itself or from the associated sensor footprint. Increasing altitude increases β and A_{cap} , thereby reducing the number of satellites required for a given instantaneous coverage objective. Higher altitudes, however, also introduce greater impediments for orbital lifetime management and end-of-life disposal. Conversely, stricter elevation constraints reduce β , requiring a denser constellation to maintain coverage quality [13, 24, 1].

A simple geometric lower bound for instantaneous single global coverage is

$$N_{\min} \gtrsim \frac{4\pi R_{\oplus}^2}{A_{\text{cap}}} = \frac{2}{1 - \cos \beta} \quad (1.14)$$

Equation (1.14) is only a lower bound since the relation does not account for overlap, recurrence, phasing, seam effects, or architectural feasibility, so its role is interpretive rather than prescriptive [13, 14, 24].

As shown in Figure 1.6, coverage metrics for satellites extend beyond instantaneous access. Revisit time quantifies the interval between consecutive observations of a target, outage time measures periods without access, and time in view denotes the temporal fraction during which continuous visibility is maintained. k -fold coverage specifies the minimum number of satellites simultaneously in view. Each metric addresses a distinct mission requirement and therefore requires separate analytical treatment [13, 24, 26]. The spherical cap is the fundamental geometric unit, while overall performance depends on the temporal evolution, overlap, and recurrence of these caps.




Figure 1.6: Single-Satellite Visibility Geometry

Mixed Discrete-Continuous Optimization in Constellation Design

Integrating canonical satellite constellations with coverage metrics clarifies the structural character of the constellation design problem. Orbital altitude (h) and inclination (i) are continuous design variables, whereas total satellite count (T), plane count (P), satellites per plane (S), recurrence integers, Walker phasing (F), and coverage multiplicity (k) are discrete variables. In Walker architectures, phasing is represented through integer values, whereas in Flower constellations, admissible geometries are additionally restricted by recurrence relations [13, 22]. Although mission objectives are often expressed through continuous performance measures such as revisit time, outage time, and availability, the feasible design set is determined by combinatorial architectural choices together with dynamical admissibility.

Satellite constellation synthesis is therefore naturally posed as a mixed-integer nonlinear optimization problem. The primary objective, therefore, is to embed those footprints within an orbital architecture that is geometrically efficient, temporally admissible, and dynamically maintainable under perturbative and operational constraints [13, 26, 22, 25, 27].

1.1.4 State Representations for Constellation Analysis

Satellite constellation analysis depends on the physical meaning assigned to the state, and that meaning is inseparable from the chosen representation [1, 2, 3]. For satellite s , an absolute Cartesian state may be written as $(\mathbf{r}_s, \mathbf{v}_s)$. Classical orbital elements (COEs) may be written as $(a, e, i, \Omega, \omega, \nu)$. Mean catalog states, nonsingular element sets, relative orbital elements (ROEs),

and local orbital frames (LOFs) each encode different dynamical content and support distinct analytical tasks. The selected state representation determines which perturbations are modeled explicitly, governs the temporal behavior of the variables, and establishes the conditions under which public catalog data may be compared consistently with high-fidelity orbital propagation. State specification is, therefore, a foundational modeling decision affecting both analytical tractability and interpretive clarity.

Analytical workflows are naturally organized around five hierarchical state families. Absolute Cartesian states support direct force-model integration and numerical propagation of the equations of motion. COEs enable geometric interpretation of orbit size, shape, and orientation. Mean and nonsingular element sets emphasize secular structure and long-timescale dynamical trends rather than instantaneous configuration, while also reducing singular behavior in limiting cases such as circular or equatorial orbits. ROEs define shell-level and pairwise geometry in a form well suited to spacing, formation structure, and differential motion analysis. LOFs support maneuver design, covariance transformation, and conjunction risk assessment by expressing relative motion and uncertainty on a dynamically meaningful basis. Transformations between these representations are nontrivial because each introduces distinct assumptions regarding coordinates, averaging, periodic content, and dynamical theory [1, 2, 3, 17, 20, 28].

Osculating and Mean Elements

Osculating elements define the instantaneous Keplerian conic tangent to the perturbed trajectory at a specified epoch, and therefore represent the orbit in its momentary geometric configuration under the full local effect of the applied perturbations. Mean elements, by contrast, remove selected periodic terms so that the resulting representation is governed primarily by secular and long-period evolution rather than short-timescale oscillatory content [29, 30, 1, 2, 31]. This structural decomposition may be written as

$$\alpha_{\text{osc}}(t) = \bar{\alpha}(t) + \alpha_{\text{sp}}(t) + \alpha_{\text{lp}}(t), \quad (1.15)$$

where $\bar{\alpha}(t)$ denotes the mean state, $\alpha_{\text{sp}}(t)$ denotes the short-period component, and $\alpha_{\text{lp}}(t)$ denotes the long-period component, each defined within the adopted averaging framework. This decomposition is not unique, as the mean state is meaningful only relative to the underlying perturbation theory or catalog convention [29, 30, 31, 1, 2].

This theoretical dependence becomes more pronounced when public catalog data are introduced. Catalog mean elements are inseparable from the propagation and averaging conventions

under which they are generated, so the mean element cannot be interchanged with osculating states without introducing systematic inconsistency. Analytical rigor, therefore, requires preserving the mean-versus-osculating distinction throughout catalog interpretation, perturbation analysis, and inverse modeling [16, 17, 20].

Classical Orbital Elements

COEs remain the canonical element set because of formal compactness and geometric interpretability. The standard form can be represented as

$$\alpha_{\text{COE}} = [a, e, i, \Omega, \omega, \nu]^\top \quad (1.16)$$

where a and e define orbit size and shape, i , Ω , and ω define orbital orientation, and ν specifies the instantaneous position along the conic [1, 2]. For averaged dynamical analyses, the true anomaly is often replaced by the mean anomaly, represented as

$$\alpha_{\text{COE}} = [a, e, i, \Omega, \omega, M]^\top \quad (1.17)$$

These variables remain central due to mapping directly onto coverage geometry, secular evolution, and traditional constellation-design methodology.

Analytical degeneracies, however, limit the usefulness of COEs as internal coordinates. Near-circular orbits render the line of apsides and argument of perigee (ω) poorly defined, whereas near-equatorial orbits render the ascending node ambiguous and destabilize RAAN. These complications are prevalent in large LEO constellations, where near-circular shells and estimation workflows approach geometric singularities. COEs, therefore, remain essential for interpretation, but are not ideal internal variables for long-horizon estimation or optimization [2, 3, 15, 32].

Nonsingular and Averaged Formulations

Nonsingular element sets address these pathologies by replacing unstable geometric variables with combinations that remain regular across the operational domain [32, 33, 2, 3]. The equinoctial framework is a representative example, as it is constructed to avoid the classical singular behavior at the circular and equatorial orbit limits. The purpose is analytic smoothness and continuity near circular and equatorial orbits, where conventional element definitions become poorly conditioned. Such representations are consequently well suited to robust propagation, estimation, and optimization in large LEO satellite constellations, particularly when numerical stability, differentiability,

and consistent behavior across various orbital regimes are required.

Mean elements suppress short-period oscillations that can obscure the secular dynamics governing shell maintenance, phase drift, and long-horizon relative motion. Both averaged and nonsingular states therefore provide a cleaner separation between rapid phase evolution and slow geometric variation than Cartesian states or osculating COEs. This separation improves the interpretability of state evolution over operational timescales [29, 30, 1, 2, 3, 31].

Relative Orbital Elements

When analysis extends beyond isolated trajectories to pairwise or fleet-level geometry, ROEs provide a more analytically robust representation than Cartesian position and velocity differences formed directly in a rotating frame. A widely used representation for a quasi-nonsingular ROE vector is

$$\delta\alpha = [\delta a, \delta\lambda, \delta e_x, \delta e_y, \delta i_x, \delta i_y]^\top \quad (1.18)$$

where δa is the normalized relative semi-major axis, $\delta\lambda$ is the relative longitudinal phase, $(\delta e_x, \delta e_y)$ are the relative eccentricity-vector components, and $(\delta i_x, \delta i_y)$ are the relative inclination-vector components. Each component carries a direct geometric role, with δa governing secular along-track drift, $\delta\lambda$ governing phasing, $(\delta e_x, \delta e_y)$ governing in-plane relative motion, and $(\delta i_x, \delta i_y)$ governing cross-track motion [15, 11, 3].

The representation is especially well suited to satellite constellation analysis, as shell coherence, slot occupancy, phasing uniformity, and structural drift can be characterized through incrementally varying relative variables. Cartesian states remain necessary for force-model integration and particular covariance workflows, but often obscure the large-scale geometric organization of the satellite constellation. ROEs therefore provide an essential intermediate abstraction between detailed dynamics and architecture-level analysis [15, 11].

Local Orbital Frames

LOFs remain indispensable for maneuver execution, covariance propagation, and conjunction assessment [1, 3, 28]. The frames are defined from the instantaneous position and velocity of the reference trajectory. In the RTN (Radial-Transverse-Normal) frame, the radial axis (R) is aligned with the position vector ($\mathbf{r}$), the transverse axis (T) lies in the orbital plane (p), and the normal axis (N) completes the orthogonal triad. In the TNW (Tangential-Normal-Binormal) frame, the tangential axis (T) is aligned with velocity vector ($\mathbf{v}$), the binormal axis (W) is aligned with the specific angular momentum vector ($\mathbf{h}$), and the normal axis (N) completes the right-handed system.

In the CCSDS LVLH (Local Vertical Local Horizontal) frame, the origin is at the satellite center of mass, the Z axis is opposite the gravimetric position vector, the Y axis is opposite the specific angular momentum vector, and the X axis completes the orthonormal triad [28, 20, 1, 3].

LOFs provide operational clarity because the frames resolve motion and control into dynamically meaningful directions. Maneuvers can be specified in radial, along-track, and cross-track components, as illustrated in Figure 1.7. Covariance propagation becomes geometrically interpretable, and conjunction uncertainty can be expressed in encounter-relevant bases. Strict frame discipline is essential, since nomenclature and axis ordering vary across publications and software, particularly for LVLH. All covariance products, maneuver definitions, and relative-state transformations in this work therefore state the employed LOF explicitly [28, 1, 3].




Figure 1.7: Local Orbital Frames

1.1.5 LEO Perturbation Environment

LEO dynamics arise from the combined action of conservative and nonconservative forces. The relative importance of each perturbation depends on altitude, inclination, satellite geometry, effective area-to-mass properties, thermospheric state, and the propagation horizon [1, 2, 3]. Over extended timescales, Earth's nonspherical gravity field establishes the dominant geometric structure of orbital evolution. Atmospheric drag, by contrast, governs dissipative evolution and constitutes the principal source of predictive uncertainty in long-horizon propagation throughout much of LEO [1, 2, 12, 34, 35, 36]. Solar radiation pressure (SRP), third-body perturbations, and attitude-dependent surface forces are typically of secondary importance, although the forces become significant for precision orbit determination, objects with substantial effective area-to-mass ratios, or over sufficiently long propagation intervals [1, 2, 3]. The LEO perturbation environment may be interpreted

as a stratified hierarchy of effects, as visualized in Figure 1.8, where this hierarchy is organized across four characteristic timescales.




Figure 1.8: Perturbation Hierarchy for LEO Propagation

Over a single orbital period, short-period variations arise from instantaneous perturbing accelerations. Over multi-orbit to daily intervals, atmospheric drag and SRP are modulated by local solar time, eclipse transitions, terrestrial rotation, and geomagnetic variability. At the operational timescale, cumulative drag decay, J_2 -driven precession, and iterative station-keeping determine shell continuity and phasing. Climatological and lifetime timescales include seasonal variability, solar rotation, geomagnetic storms, and the solar cycle, which modulate thermospheric density and alter the structure of long-horizon propagation [34, 35, 36, 1, 2].

Cowell Formulation and Variation of Parameters

Cowell's formulation integrates the Cartesian EOM directly and incorporates all retained force-model contributions without requiring transformation to an alternative coordinate set [1, 2, 3]. The

perturbed EOM, introduced as Equation (1.5), is again written as

$$\ddot{\mathbf{r}} = -\frac{\mu_{\oplus}}{r^3}\mathbf{r} + \mathbf{a}_{\text{harm}} + \mathbf{a}_{\text{drag}} + \mathbf{a}_{\text{SRP}} + \mathbf{a}_{\text{3b}} + \dots$$

where $\mathbf{a}_{\text{harm}}$ collects the nonspherical-gravity contributions, including J_2 , J_3 , and higher-order harmonics. This form is well-suited to high-fidelity propagation, adaptive numerical integration, and state-transition or covariance analysis in Cartesian space [1, 2].

Variation of parameters (VOP) alternatively describes the evolution of orbital elements under perturbation, thereby replacing direct propagation of the Cartesian state with an evolution formulation of a chosen element set. In practice, the Gauss variational equations (GVEs) provide a convenient route for both propagation and interpretation of how external accelerations modify orbital geometry [1, 2, 3]. In satellite constellation mechanics, shell evolution, phase drift, and relative geometry are often more naturally expressed in terms of slowly varying element-based variables than in rapidly oscillatory inertial coordinates. Cowell propagation is therefore well suited to numerical force-model integration, whereas element-based dynamics support interpretation, reduced-order modeling, and control-oriented analysis [1, 2, 3].

Earth Gravity, J_2 , and Higher Harmonics

The nonspherical component of Earth’s gravity field constitutes the primary conservative perturbation affecting low Earth orbiters [1, 2, 37]. Geopotential models, including EGM2008, employ comprehensive spherical harmonic expansions; however, for constellation-scale analyses, retaining only the second zonal harmonic is generally sufficient to represent the principal deviation from the central potential [37, 1, 2]. The coefficient J_2 characterizes Earth’s oblateness and determines the dominant secular rates of nodal and apsidal drift.

For near-circular orbits, the first-order secular nodal regression is given by

$$\dot{\Omega}_{J_2} \approx -\frac{3}{2}J_2 n \left(\frac{R_{\oplus}}{a}\right)^2 \cos i \quad (1.19)$$

where $n = \sqrt{\mu_{\oplus}/a^3}$ [1, 2]. This expression reveals the dominant shell-level effect of Earth’s oblateness. Small variations in semi-major axis or inclination result in differential nodal drift, which can quickly lead to plane decorrelation if active station-keeping is not implemented. At low Earth altitudes, J_2 primarily determines the secular structure of constellation shells [1, 2, 11].

Higher-order zonal, tesseral, and sectorial harmonics introduce additional periodic and secular corrections. These terms are essential for accurately modeling resonance phenomena, repeating-

ground-track requirements, or long-duration orbit predictions [37, 1, 2]. Although a J_2 -only model captures first-order shell evolution, full-spectrum geopotential expansions are necessary to achieve the accuracy required for high-precision or long-term analyses. The truncation level of the harmonic expansion should be carefully matched to the mission’s allowable dynamical error [37, 1].

Atmospheric Drag, Ballistic Coefficient, and Thermospheric Uncertainty

Atmospheric drag is the primary nonconservative perturbation in LEO, governing semi-major axis decay, orbital lifetime, and dissipative orbital evolution [1, 2, 12, 34, 35, 36]. A standard drag acceleration model is expressed as

$$\mathbf{a}_{\text{drag}} = -\frac{1}{2} \rho \frac{C_D A_D}{m} \|\mathbf{v}_{\text{rel}}\| \mathbf{v}_{\text{rel}} = -\frac{1}{2} \rho \frac{1}{BC} \|\mathbf{v}_{\text{rel}}\| \mathbf{v}_{\text{rel}} \quad (1.20)$$

where ρ is the atmospheric mass density, C_D is the drag coefficient, A_D is the effective drag area, m is the satellite mass, and $\mathbf{v}_{\text{rel}}$ is the velocity relative to the co-rotating atmosphere. The ballistic coefficient is denoted as

$$BC = \frac{m}{C_D A_D} \quad (1.21)$$

which constrains only the combined influence of C_D , A_D , and m , so the individual parameters remain underdetermined when orbital decay is the only observation channel [1, 2, 12].

For most operational satellites in LEO, the ambient flow lies in the transitional-to-free-molecular regime, rendering continuum aerodynamic models unsuitable [1, 2, 38]. Effective atmospheric drag is governed by gas-surface interaction, atmospheric composition, energy accommodation, projected area, and time-dependent attitude. The drag coefficient serves as an effective parameter where the numerical value depends on the flow regime, geometric configuration, and model assumptions [38, 39, 36].

Thermospheric mass density is the primary source of uncertainty in drag modeling for satellite constellations. Empirical thermospheric models such as NRLMSISE-00, NRLMSIS 2.0, JB2008, and DTM2020 provide climatological density specification across the upper atmosphere as functions of altitude, location, time, and solar and geomagnetic forcing [34, 35, 40, 41]. NRLMSISE-00 incorporates anomalous-oxygen corrections that improve upper-thermosphere specification, whereas NRLMSIS 2.0 extends empirical density modeling from the surface to the exobase and captures broader spatial, temporal, and activity-dependent variability [34, 35]. JB2008 supplements classical proxy-driven density modeling with additional solar and geomagnetic indices, and DTM2020 provides a semi-empirical temperature-density-composition framework developed specifically for orbit computation [40, 41]. Physics-based frameworks such as TIE-GCM further

describe the coupled thermosphere-ionosphere system from first principles, but practical orbit prediction remains strongly limited by uncertainty in external forcing and density specification [42, 36]. Multiplicative uncertainty in density is the dominant contributor to long-horizon orbit-prediction error for drag-sensitive missions. Errors in density proxies, empirical atmosphere models, projected-area histories, and gas-surface interaction assumptions accelerate semi-major axis and phase divergence and thereby amplify uncertainty growth [34, 35, 39, 36, 1].

Solar Radiation Pressure

SRP constitutes a principal surface force in high-fidelity orbital dynamics for satellite constellations [1, 2]. In the cannonball approximation, the satellite is modeled as a sphere with uniform optical properties, so the SRP acceleration may be written as

$$\mathbf{a}_{\text{SRP}} = \chi_{\text{ecl}} P_{\odot} C_R \frac{A_{\text{SRP}}}{m} \left(\frac{\text{AU}}{r_{\odot}} \right)^2 \hat{\mathbf{s}} \quad (1.22)$$

where $P_{\odot}$ is the solar radiation pressure at one astronomical unit (AU), C_R is the effective reflectivity coefficient, A_{SRP}/m is the Sun-facing area-to-mass ratio, $r_{\odot}$ is the Sun-satellite distance, $\hat{\mathbf{s}}$ is the unit vector from the Sun to the satellite, and χ_{ecl} is the eclipse factor [1, 2]. Both SRP and atmospheric drag scale with area-to-mass ratio and surface properties; however, SRP arises from photon momentum transfer, whereas drag arises from gas-surface interaction.

For compact LEO satellites, SRP inherently has a limited influence on semi-major axis evolution relative to atmospheric drag. Nevertheless, SRP remains important in precise orbit determination, covariance analysis, and systems with high area-to-mass ratios [1, 2]. Large deployable appendages can amplify SRP forces and torques, especially through changing projected area and eclipse transitions. High-fidelity LEO force models retain SRP even when the translational magnitude is smaller than that of drag [1, 2, 3].

Solar and Lunar Third-Body Gravitational Effects

Third-body gravitational perturbations are represented in an Earth-centered inertial frame through the direct-minus-indirect acceleration form [1, 2]

$$\mathbf{a}_{3b,b} = \mu_b \left(\frac{\mathbf{r}_b - \mathbf{r}}{\|\mathbf{r}_b - \mathbf{r}\|^3} - \frac{\mathbf{r}_b}{\|\mathbf{r}_b\|^3} \right), \quad b \in \{\odot, \ominus\} \quad (1.23)$$

where μ_b is the gravitational parameter of the perturbing body and $\mathbf{r}_b$ is the position vector from Earth to that body. The indirect term accounts for the fact that the Earth-centered integration frame

is not inertial. High-fidelity orbital models include solar and lunar gravity as the dominant external point-mass perturbations [1, 2].

At LEO, third-body accelerations are negligible relative to J_2 in secular shell evolution and relative to atmospheric drag in dissipative evolution. As altitude increases, or as accuracy and propagation requirements become more stringent, third-body perturbations become necessary for precision trajectory prediction and orbit determination. The omission of third-body perturbations is therefore appropriate only in coarse models with limited fidelity requirements [1, 2].

Translation-Attitude Coupling

Decoupling translational and rotational motion is generally inadequate for satellites with large exposed surfaces, high area-to-mass ratios, or limited attitude-control authority [1, 3, 38]. Atmospheric drag and SRP both depend strongly on projected geometry and surface orientation, so translational accelerations are coupled to attitude through quantities such as A_D , A_{SRP} , C_D , and C_R . Environmental loading also produces aerodynamic and radiative torques, so orbital and attitude motion must often be treated as a coupled dynamical system [1, 2, 3, 38].

This coupling governs several operationally important scenarios in which differential-drag phasing uses attitude maneuvers to modulate the projected drag area and regulate along-track separation [15, 43]. High-area satellites and drag-augmentation devices exhibit force characteristics that depend directly on deployed geometry. Objects without full attitude control, including inoperative satellites and debris, experience time-varying drag and SRP through tumbling-induced changes in projected area. Inverse estimation of unknown surface or attitude-related parameters from orbital evolution also relies on this intrinsic translation-attitude coupling [1, 2, 3, 38].

1.1.6 Secular Evolution, Differential Drift, and Shell Maintenance

The maintenance of a satellite shell is governed by the evolution of mean orbital elements [1, 2, 11]. Persistent influences, such as the J_2 perturbation, atmospheric drag, and control inputs, continuously modify the shell geometry away from any fixed configuration. Sustaining shell structure requires close alignment of the secular drift rates that determine plane spacing, mean motion, and along-track phasing. Shell coherence is fundamentally a rate-matching condition governed by the temporal evolution of mean-element rates [1, 2, 11, 15].

The primary conservative mechanism governing shell evolution is nodal precession due to Earth’s oblateness. For a near-circular orbit, the first-order secular rate of RAAN drift is formulated

as

$$\dot{\Omega}_{J_2} \approx -\frac{3}{2} J_2 n \left(\frac{R_\oplus}{a} \right)^2 \cos i = -\frac{3\sqrt{\mu_\oplus} J_2 R_\oplus^2}{2a^{7/2}} \cos i \quad (1.24)$$

Nodal regression increases as a decreases and weakens as i approaches polar values. Architectural plane spacing remains viable only when mean-element drift rates remain closely matched across the shell population for a satellite constellation [1, 2, 11].

Linearizing Equation (1.24) about a reference shell makes the sensitivity explicit. For small deviations in a and i ,

$$\frac{\delta \dot{\Omega}}{\dot{\Omega}} \approx -\frac{7}{2} \frac{\delta a}{a} - \tan i \delta i \quad (1.25)$$

where RAAN coherence is therefore highly sensitive to small deviations in mean altitude and inclination because of the $a^{-7/2}$ and $\tan i$ dependences. Even minor discrepancies prevent rigid shell drift and induce progressive shear within the internal plane structure [1, 2, 11].

The primary dissipative mechanism is drag-driven decay of the semi-major axis. The mean motion satisfies

$$n = \sqrt{\frac{\mu_\oplus}{a^3}} \quad (1.26)$$

Thus, the first-order sensitivity is

$$\frac{\delta n}{n} = -\frac{3}{2} \frac{\delta a}{a} \quad (1.27)$$

A satellite at a lower altitude completes revolutions more rapidly and accordingly advances in the along-track phase relative to higher-altitude counterparts. Differential-drag phasing exploits this effect, where increasing drag accelerates altitude loss, and increases mean motion, advancing the phase until the desired spacing is reached [44, 43]. The same sensitivity also underlies in-plane spacing degradation caused by uncontrolled variations in ballistic coefficient, projected area, or attitude history, even when initial slotting is precise [1, 2, 44, 43].

Common-mode evolution, characterized by nearly identical secular nodal rates and drag histories, produces collective shell descent and precession while preserving internal topology in a co-moving mean-element frame. Differential evolution arises from small mismatches in semi-major axis, inclination, ballistic coefficient, area variation, or maneuver timing, producing persistent secular-rate differences that induce phase compression, spreading, plane shear, and regrouping as visualized in Figure 1.9. The relevant analytical boundary is set by secular-rate mismatch, because geometric alignment at a single epoch provides no guarantee of continued shell coherence [1, 2, 11, 15].




Figure 1.9: Differential J_2 -Driven Nodal Regression and Shell-Plane Shear

Secular drift mechanisms are subsequently operationally useful during deployment and re-configuration. Differential nodal regression allows satellites in drifting orbits at different altitudes to accumulate passive RAAN separation before final plane insertion. Similarly, differential drag governs in-plane phase management for satellite constellations [44, 43, 45]. In both cases, secular drift functions as a controllable design resource within the operational design space.

An on-station shell, therefore, presents an intrinsically coupled control problem. Altitude adjustments introduced for phasing, drag management, or reconfiguration simultaneously modify orbital period and nodal drift rate. Altitude thus acts as a coupled control input that influences both along-track and RAAN evolution. Effective shell maintenance consequently requires a secular-rate control perspective that supersedes independent tolerances on instantaneous orbital elements [1, 2, 11, 46, 45].

1.1.7 Relative Motion, Phasing, and Station-Keeping

The transition from orbit propagation to inter-satellite geometry changes both the principal state representation and the control objective [3, 15, 11]. Relative flight may be separated into two operational regimes. Local rendezvous and close formation emphasize bounded geometry, maneuver design, and collision avoidance for separations that are narrow compared with the orbital radius. Large-fleet slot maintenance, by contrast, emphasizes shell allocation, along-track phasing, plane alignment, secular drift, and persistent control over long time horizons. Although both regimes are subject to the same environmental perturbations, the preferred coordinates and control formulations differ substantially [47, 48, 15].

Linear Relative Motion

The Hill-Clohessey-Wiltshire (HCW) equations provide the canonical linearized framework for local relative motion about a circular reference orbit [47, 3]. In a chief-centered RTN frame, let (x, y, z) denote the deputy position relative to the chief in the radial, along-track, and cross-track directions,

respectively. For a circular chief orbit with mean motion n , the deputy dynamics satisfy

$$\ddot{x} - 2n\dot{y} - 3n^2x = u_x, \quad \ddot{y} + 2n\dot{x} = u_y, \quad \ddot{z} + n^2z = u_z \quad (1.28)$$

Here, u_x , u_y , and u_z denote control accelerations resolved in the same chief-centered local frame. The linearity, time invariance, and closed-form solvability of the HCW equations establish the system as a standard model for rendezvous, inspection, and local guidance near circular chief orbits [47, 3]. However, atmospheric drag, Earth's oblateness, and the chief eccentricity all introduce effects that the nominal HCW equations do not capture.

These restrictions limit the HCW equations to near-field operations since, as relative separation, chief eccentricity, or propagation horizon increases, the system ceases to represent the long-horizon geometry relevant to perturbed shell maintenance and slot control [47, 48, 3].

Elliptical-Reference Relative Motion

The Tschauner-Hempel family extends the circular-reference framework by parameterizing the equations with the chief anomaly, which introduces periodic coefficient variation and makes the relative dynamics explicitly dependent on the chief's position along its elliptical orbit [49, 48]. The Yamanaka-Ankersen (YA) state transition matrix (STM) provides a compact and widely used linear time-varying solution for arbitrary elliptical chief orbits with eccentricity in the interval $0 \leq e < 1$ [48, 3].

From this perspective, the HCW equations are a limiting special case within a broader theory of relative motion. Chief eccentricity alters the structure of the dynamics by introducing anomaly-dependent coefficient variation, thereby removing time invariance from the linearized relative dynamics. For larger eccentricities or longer transfers, local Cartesian analysis becomes progressively less descriptive of the evolving relative geometry, and formation analysis is more naturally shifted toward anomaly-parameterized formulations or mean relative orbital elements to preserve long-horizon geometric fidelity and orbital-structure insight [48, 15, 3].

Passive Safety Through Eccentricity-Inclination Separation

A passively safe configuration maintains a nonzero minimum separation ($d_{\min}$) without requiring active intervention. Within the ROE frame, this is achieved by phasing the relative eccentricity and inclination vectors so that the in-plane closest approach does not occur simultaneously with zero cross-track separation, thereby preventing the relative trajectory from collapsing into a true close approach at the same orbital phase. The resulting geometric offset establishes a deterministic

safety margin embedded in the relative trajectory itself [50, 15].

This construction adds a geometric safety layer that is independent of probabilistic uncertainty quantification and provides a deterministic robustness criterion built directly into the relative-orbit design. For satellite constellations and distributed formations, that distinction permits a dynamic separation standard tied to the geometry of the relative motion itself rather than one based solely on probabilistic miss-distance assessment [50, 15].

Differential Drag Phasing

In LEO, differential drag serves as the primary propellantless actuator for along-track phasing. Attitude maneuvers modify the projected drag area and, therefore, the drag acceleration, increasing semi-major axis decay, lowering altitude, increasing mean motion, and advancing the phase relative to low-drag reference satellites [44, 43]. Within the ROE frame, differential drag acts primarily through δa and $\delta \lambda$, which is predominantly applicable for commissioning, slot acquisition, and medium-scale phasing.

Differential drag is inherently limited to in-plane actuation, provides negligible cross-track authority, consumes altitude as a control resource, and depends strongly on thermospheric variability, collectively with precise attitude control [44, 43]. The primary utility lies in coarse fleet phasing over expansive scales for LEO, and it does not replace full rendezvous capability or complete relative-geometry control.

Low-Thrust and Impulsive Station-Keeping

When propulsion is available, station-keeping and reconfiguration rely on continuous low-thrust or impulsive control. Continuous low-thrust distributes actuation over multiple orbits to counter drag, secular drift, and shell distortion, thereby supporting long-horizon maintenance. Impulsive control, in contrast, enables discrete reconfiguration, rapid retargeting, or fuel-efficient transfers between relative geometries. Within the ROE frame, both approaches act on the gradually varying state variables that define the relative orbit, with instantaneous Cartesian deviations serving primarily as local operational coordinates [3, 15, 51].

This perspective defines the main control objectives where altitude control modifies a , thereby affecting both satellite lifetime and shared mean motion. Relative phase is governed primarily by δa and $\delta \lambda$. Plane control through i and Ω is generally more propulsively expensive than along-track maneuvering. Full relative-geometry control, therefore, requires direct manipulation of the e and i components of $\delta \alpha$. Differential drag remains practically valuable for large LEO satellite constellations, but only as a limited along-track actuation mechanism within a broader

station-keeping framework [3, 15, 43, 51].

1.1.8 Conjunctions, Debris Feedback, and Disposal Constraints

Operational conjunction risk arises from the consolidated influences of state uncertainty, persistent debris, and failures in post-mission disposal [52, 53, 6, 54, 7, 8]. Local trajectory proximities escalate into broader sustainability challenges when repeated encounters, fragmentation events, and unsuccessful disposal operations alter the debris environment and thereby adjust future conjunction rates [6, 54, 55, 56]. In large LEO satellite constellations, encounter mechanics, collision-risk assessment, maneuver decision logic, debris proliferation, and disposal processes thus operate as coupled elements within a single operational framework [7, 8, 21, 9, 10].

Encounter Geometry in the B -Plane

Short-duration conjunction geometry is defined at the time of closest approach (TCA) in the B -plane, or encounter plane, which is orthogonal to the instantaneous relative velocity vector ($\mathbf{v}_{\text{rel}}$) [52, 53]. The relative miss vector projected onto this plane represents the effective spatial geometry of the encounter at closest approach. Under the short-encounter approximation, the inherently three-dimensional event is reduced to a two-dimensional problem involving the projected mean miss vector, the projected covariance, the hard-body interaction region, and the time remaining until TCA [52, 53, 57].

This reduction is operationally essential, since for sufficiently rapid encounters local linearization provides an accurate representation of the collision geometry, and B -plane constructs directly support probability-of-collision estimation and maneuver design [52, 53, 57].

Collision Probability Under Gaussian Covariance

Collision probability is the probability mass for which the projected relative-position uncertainty intersects the hard-body region. Let $\boldsymbol{\mu}_B$ denote the projected mean miss vector, $\mathbf{P}_B$ denote the projected covariance in the B -plane, and the hard-body region be represented by a circle of radius R_{HBR} [52, 53, 57]. Thus, the formulation can be denoted as

$$P_c = \iint_{\|\boldsymbol{\rho}_B\| \leq R_{\text{HBR}}} \frac{1}{2\pi\sqrt{\det \mathbf{P}_B}} \exp\left[-\frac{1}{2}(\boldsymbol{\rho}_B - \boldsymbol{\mu}_B)^\top \mathbf{P}_B^{-1}(\boldsymbol{\rho}_B - \boldsymbol{\mu}_B)\right] d\rho_{B,1} d\rho_{B,2} \quad (1.29)$$

Equation (1.29) defines the Gaussian hard-body incursion probability on the encounter plane. This formulation is appropriate when relative motion is locally linear, relative velocity is effectively

constant over the encounter interval, and the projected covariance at TCA adequately characterizes uncertainty [52, 53, 57].

Disposal Mechanics in LEO

Disposal in LEO is mechanically altitude-dependent, governed by ballistic response and available control authority [7, 8, 21, 9, 10]. Passive disposal relies on atmospheric drag for post-mission disposal (PMD) and is most effective in lower LEO, where operational altitudes or modest perigee reductions can enable re-entry within policy-relevant timescales. Passive drag functions as a valuable contingency mechanism and a complement to active disposal when required [7, 8, 21, 12].

At higher LEO altitudes, passive removal becomes prohibitively gradual, so propulsion is generally required to lower perigee or transfer the satellite into a more dissipative orbit. In persistence-dominated shells, disposal reliability becomes a critical satellite constellation design parameter since inoperative satellites can remain in orbit long enough to alter future conjunction rates [6, 54, 55, 7, 8, 10].

Drag augmentation increases the effective cross-sectional area and, in turn, increases atmospheric drag, thereby reducing orbital lifetime. This engineered modification of ballistic response accelerates entry into a more dissipative regime and is best interpreted as an enabling mechanism for disposal rather than as a separate disposal class [1, 2, 23, 8].

Standards and Regulations as Mechanical Constraints

Debris standards and disposal regulations impose constraints on physically admissible orbital behavior. The U.S. Government Orbital Debris Mitigation Standard Practices (ODMSP), NASA debris standards, and the FCC five-year disposal rule place explicit constraints on post-mission lifetime, LEO persistence, breakup prevention, and disposal capability [7, 8, 21, 9, 10, 58]. In that sense, policy acts through mechanics, as guidelines constrain the allowable orbital evolution and the feasible design space of satellite constellations.

This policy structure is especially applicable for large LEO satellite constellations. The 2019 ODMSP revision added explicit treatment of large satellite constellations and proximity operations, while the FCC five-year rule imposes a stricter disposal timeline than the earlier twenty-five-year benchmark for FCC-licensed satellites operating in, or passing through, LEO [7, 10, 58]. Changes in the permitted PMD interval consequently recalibrate requirements for disposal reserve, shell altitude, dependence on atmospheric drag, and the viability of drag augmentation. The regulatory framework thus defines part of the admissible mechanical design space for satellite constellation deployment [7, 8, 21, 9, 10, 58].

1.1.9 Modeling Hierarchy, Public Data Products, and the Scientific Gap

The central epistemic question is how substantively structural information can feasibly be recovered from public orbital data alone. Framing the question in terms of operator-grade products clarifies the scientific gap in which the recoverable content of a dataset may be determined by state semantics, propagation theory, data provenance, and the rigor of uncertainty treatment [16, 17, 20, 18, 19]. A satellite constellation state may appear as a theory-tagged mean-element catalog product, a covariance-bearing state or ephemeris product, or a numerically integrated trajectory [20, 1, 2].

Propagation Hierarchy

The propagation hierarchy defines a fidelity ladder through state semantics and admissible inference tasks, as shown in Figure 1.10. General-perturbation (GP) models, which are essential to TLE and OMM workflows, operate on theory-tagged mean-element states [16, 17, 20, 18, 19]. The practical value of SGP4 lies in its model consistency with the GP mean-element ecosystem, which can be used for operational satellite studies [16, 17, 19].

An intermediate regime is provided by semi-analytical propagators, which evolve averaged or mean states while treating short-period dynamics analytically and reserving numerical integration for slower secular and long-period behavior. This architecture increases force-model fidelity beyond catalog propagation while preserving computational tractability for long-horizon analysis and large satellite constellations [29, 30, 33, 59, 31].

The highest-fidelity regime is full numerical integration, in which force models, geophysical constants, atmospheric and solar drivers, third-body perturbations, maneuvers, and stochastic terms can be selected explicitly. This regime enables the strongest physical interpretation, but also depends on model assumptions and operator-specific inputs that are typically absent from public catalog products [1, 2, 3, 20].




Figure 1.10: Propagation Hierarchy for LEO Propagation

Public Catalogs

Public catalogs remain indispensable for architectural and trend-level analysis supporting shell identification, plane population mapping, approximate phase-distribution analysis, altitude-evolution tracking, deployment-corridor reconstruction, and anomaly screening [16, 17, 20, 18, 19]. In this role, the catalog serves as a universally accessible input for structural inference at scale across populations too large for object-by-object scrutiny, distinct from operator telemetry [4, 5].

Drag-sensitive analyses are also feasible when posed within a catalog-consistent observational framework. Shell contraction, comparative decay trends across subpopulations, and population-level drag-response variability can be inferred when the analysis respects the theory-dependent semantics of mean-element catalog states rather than being treated as osculating orbital trajectories [16, 17, 20, 19, 12].

The limitations become decisive once covariance realism, maneuver prediction, or satellite-specific force metadata are required. Public element sets do not provide operator-grade maneuver histories, attitude records, or projected-area histories, and generally do not support rigorous object-

specific uncertainty propagation or post-maneuver prediction using public data alone [20, 17, 18]. The scientific gap, therefore, lies in the boundary between population-scale structural inference and operator-grade state knowledge.

Richer Orbit Products

Orbit Data Messages (ODMs) formalize the distinction between public catalogs and richer orbit products. Within the ODM family, the Orbit Mean-Elements Message (OMM) supports compact, theory-tagged mean-element exchange, the Orbit Parameter Message (OPM) supports epoch-state exchange with ancillary propagation data, the Orbit Ephemeris Message (OEM) supports higher-fidelity trajectory histories, and the Orbit Comprehensive Message (OCM) extends the exchange to include covariance, physical properties, maneuver information, perturbation-model content, and related metadata. These distinctions are scientifically substantive since they determine which state semantics, uncertainty structures, and force-model assumptions are preserved in the exchanged product [20].

A practical demarcation, therefore, emerges between the classes of inference supported by public catalogs and those requiring richer exchange products. Public catalogs remain well-suited to architectural and trend-level inference, including shell structure, deployment patterns, approximate phase organization, and broad decay characterization. By contrast, covariance-aware conjunction analysis, maneuver reconstruction, high-precision uncertainty propagation, and inference requiring object-specific force-model or configuration metadata require richer products that preserve trajectory history, uncertainty structure, and ancillary dynamical information [20, 17, 52, 53].

Limitations of Public Catalogs

The primary limitation of public catalogs is theory dependence. TLE data is a compressed mean-element fit constructed for compatibility with GP propagation, especially SGP4 [16, 17, 19]. This design yields substantial utility for catalog-scale structural analysis, but it also constrains the set of physically valid dynamical interpretations.

The second limitation is the absence of uncertainty quantification. Deterministic public states support screening and exploratory analysis, but do not independently provide the covariance information required for operator-grade, covariance-based collision-risk assessment [20, 52, 53, 57]. In practical terms, the uncertainty structure required for rigorous probability-of-collision estimation and high-confidence operational decision support is not present in the public catalog product itself.

The third limitation is the absence of satellite-physics metadata. Public catalogs do not preserve

attitude state, projected-area history, maneuver intent, drag area, solar-radiation-pressure area, or object-specific nonconservative-force parameters. These quantities remain latent, aggregated, or structurally uncertain when inference is attempted from orbital history alone [20, 17, 18].

The fourth limitation concerns the information regime under which the data products are constructed. Public state histories and operator-grade ephemerides may refer to the same satellite, but are generated under fundamentally different observational, filtering, and metadata conditions. The difference is therefore one of information content and accuracy [20, 16, 17].

Object-by-Object Analysis Failure at Mega-Constellation Scale

Object-level analysis becomes inadequate at the mega-constellation scale due to population size, internal substructure, and marked temporal heterogeneity. Manual interpretation becomes impractical since shells, planes, launch batches, phasing lanes, control modes, and anomalous subpopulations cannot be robustly identified by inspecting satellites one at a time. Even within a single shell, satellites can occupy different drag states, maneuvering regimes, and lifecycle phases, which further complicates detailed object-by-object interpretation [4, 5, 1, 2].

Scientific Gap

The literature on constellation geometry, perturbation theory, conjunction assessment, and disposal is well developed individually, but an integrated public-data workflow for LEO satellite constellation inference remains absent [13, 1, 2, 52, 7]. The first gap is the lack of a standard, scalable methodology for discovering constellation substructure from public data without relying on operator labels or rigid shell taxonomies. The second gap is the absence of an established public-data framework for estimating satellite physical parameters in a manner consistent with theory-annotated catalog observables and high-fidelity force models. The third gap is the absence of a unified inference chain in which discovered substructure informs subsequent physical inference and uncertainty treatment.

The present work addresses these gaps through a sequential workflow. The first stage identifies latent satellite families from public orbital histories through reduced representations and data-driven decomposition. The second stage develops forecasts of future thermospheric activity to improve drag-consistent propagation. The third stage performs family-specific physical inference focused on effective drag dynamics. The fourth stage performs family-specific propulsion inference through thruster-parameter estimation in catalog-consistent models. The final stage incorporates the recovered family structure and inferred physical parameters into recommendations for satellite constellation-level optimization.

1.2 Starlink Satellite Constellation

1.2.1 Gen 1 Starlink Case Study

This case study employs the orbital mechanics, perturbation environment, modeling hierarchy, and public catalog constraints developed for a large LEO satellite constellation whose licensed architecture and fleet-scale dynamics are sufficiently structured to support empirical inference from public TLE archives [1, 2, 16, 17, 20]. The Gen 1 Starlink satellite constellation is adopted as the model system because of the scale, multi-shell organization, drag-sensitive operating altitudes, and sustained operational activity that generate a public record abundant enough to expose population-level morphology, shell transitions, and re-entry behavior [60, 61, 4, 62, 63, 64]. Historical TLE records for 3024 Starlink satellites, obtained through the public Space-Track archive, provide the empirical basis for the analysis [18].

Scale and Analytical Advantages

The Gen 1 Starlink satellite constellation is analytically valuable since the low-altitude architecture is explicitly licensed, and the realized fleet evolves at a scale sufficient for structure to emerge directly from the public archive. The principal low-altitude divisions are centered near 540 km at 53.2° , 550 km at 53.0° , 570 km at 70.0° , and 560 km at 97.6° [61]. The mid-inclination divisions dominate the deployed population, whereas the higher-inclination divisions exhibit distinct secular evolution and geometrically differentiated coverage. This stratification supports quantitative analysis of shell occupancy, division separation, inclination-dependent secular drift, and lifecycle variance without requiring the fundamental satellite constellation architecture to be reconstructed from unstructured observations alone [61, 4].

In smaller satellite constellations, deployment, residence, and retirement appear primarily as discrete satellite-level events. In the Gen 1 Starlink satellite constellation, the same processes also emerge as coherent population-scale patterns distributed across shells and over time. Across the archival record, collective shell filling, differentiated responses to perturbations, and overlapping phases of injection, orbit raising, service residence, relocation, and disposal can be observed, indicating that lifecycle dynamics are expressed at the level of structured behavior in the aggregate constellation state [65, 63, 64, 66].

The operational regime is especially informative since atmospheric drag strongly shapes orbital evolution and operational timescales at the aforementioned altitudes [61, 67, 62, 66]. Service-shell occupancy therefore coexists with frequent orbit raising, station-keeping, controlled lowering, and replacement activity, much of which effectuates recognizable signatures in the public

catalog. Deployment commonly proceeds from lower insertion altitudes before transfer toward service shells, and the public record correspondingly exhibits orbit raising, shell residence, inter-shell relocation, and retirement across a statistically significant portion [67, 65, 63, 68]. Since the signatures recur across thousands of objects, the satellite constellation supports inference of architectural structure, effective physical behavior, and persistent patterns of operator intervention from aggregate dynamics. The licensed Gen 1 Starlink satellite constellation design is documented through FCC non-geostationary orbit (NGSO) fixed-satellite service (FSS) filings, whereas the public archive records time-indexed GP/TLE states for individual catalog objects [60, 61, 18, 19]. The two information layers, licensed design and public catalog history, are used jointly while remaining methodologically distinct.

Maintaining that separation is methodologically imperative, as the present study is explicitly framed as a public-data case study and does not attempt an exhaustive operational census or an attribution of internal operator intent. The objective is to identify satellite constellation-level behaviors that can be recovered from the public catalog archives, establish which inferences remain physically supported at the GP observation layer, and define the limits imposed by current public catalog products [18, 19, 16, 17, 20]. Accordingly, the scope includes architectural discovery, lifecycle interpretation, and division-specific secular diagnostics that remain observable under catalog-level state definitions and metadata constraints. The case study excludes covariance-grade conjunction assessment, maneuver reconstruction at operator fidelity, and satellite-resolved force modeling, all of which require proprietary ephemerides, covariance products, or satellite-specific metadata that are not present in standard public catalog releases [20, 17, 68].

Catalog Semantics and Essential Caveats

Public TLE records, and GP products more generally, are mean-element constructs designed for use with the SGP4 propagator and thus do not constitute definitive truth states [16, 17, 20, 19]. Their analytical value is substantial, but it remains conditioned on respecting the assumptions, theory dependence, and limited information content of the GP observation layer. Within that layer, the public catalog archive supports robust identification of shell populations, separation among orbital divisions, secular trend analysis, and reconstruction of orbit-raising or orbit-lowering histories from the time evolution of the published element sets. The public catalog archive does not provide access to operator-maintained ephemerides, maneuver logs, covariance information, or satellite-specific force-model parameters such as deployed area, attitude history, or commanded thrust, and those missing quantities bound the precision and physical specificity of any inference drawn solely from public GP data [18, 20, 17].

1.2.2 Orbital Architecture

A rigorous characterization of the constellation is necessary for both historical accuracy and dynamical interpretation. The licensed Gen 1 Starlink architecture is a systematically engineered LEO network organized into distinct divisions separated by altitude, inclination, and orbital-plane allocation [60, 61, 69, 70, 67]. These parameters are dynamically consequential. Shell altitude governs atmospheric drag, orbital-decay timescale, and station-keeping burden; inclination determines latitudinal access and strongly influences secular nodal regression; and the number and distribution of orbital planes determine spatial tessellation and temporal phasing [1, 2, 61].

Shell Structure

The Gen 1 Starlink architecture consists of four principal service divisions centered at 540 km and 53.2°, 550 km and 53.0°, 570 km and 70.0°, and 560 km and 97.6° [61, 67]. In the licensed Gen 1 configuration, the 550 km, 53.0° division contains 72 orbital planes with 22 satellites per plane, the 540 km, 53.2° division also contains 72 orbital planes with 22 satellites per plane, the 570 km, 70.0° division contains 36 orbital planes with 20 satellites per plane, and the 560 km, 97.6° polar division is split between 6 orbital planes with 58 satellites per plane and 4 additional orbital planes with 43 satellites per plane [61]. The licensed shell altitudes specify center altitudes for bounded operational altitude ranges, so the operational populations remain shell-centered while still exhibiting bounded altitude dispersion and continuous turnover driven by deployment cadence, replenishment, relocation, and end-of-life removal [61, 70, 67, 66].

The Gen 1 Starlink satellite constellation is therefore characterized by exclusive use of low-altitude operational shells. All principal service divisions are centered within the 540–570 km range, whereas the earlier authorized architecture placed the original 4,425-satellite system at altitudes of approximately 1110–1325 km [60, 69, 61, 67]. Operation in this lower-altitude regime increases sensitivity to atmospheric drag, accelerates dissipative evolution, shortens passive deorbit timescales, and imposes a persistent requirement for orbit raising and station-keeping, thereby placing the constellation firmly within the LEO dynamical regime [1, 2, 70, 67, 62, 66].

Planes, Phasing, and Division Scale

The organization of orbital planes further stratifies the shell architecture. The principal 53.0° division comprises 72 planes, and the 53.2° division likewise comprises 72 planes. The 70.0° division contains 36 planes, while the 97.6° division is partitioned into two smaller groupings in the revised Gen 1 authorization, namely 6 planes and 4 planes [61]. This geometric partitioning

produces a spatially explicit inertial tiling structured by the coupled relationships among shell altitude, plane count, and inter-plane phase organization.

Intentional plane partitioning introduces substantial anisotropy across the constellation divisions. High plane densities generate persistent and comparatively intricate spatial lattices, whereas lower plane counts yield sparser high-latitude components with more specialized geometric roles. Architectural nonuniformity is therefore intrinsic to the licensed design, and homogenized analytical treatments would suppress physically important structure in the plane-level organization of the system [61, 69, 70].

Mid-Inclination and High-Latitude

The two mid-inclination divisions form the dynamical core of the Starlink system, containing most of the deployed population in closely related low-altitude prograde shells [61, 4]. These divisions dominate operational occupancy, shell maintenance, and much of the recurring orbit-raising and relocation activity visible in the public record. Their nearby but distinct nominal inclinations and shell altitudes, at 53.0° and 53.2° , imply measurably different responses to drag, phasing evolution, and secular drift. Treating them as a single shell would therefore obscure meaningful dynamical structure [61, 1, 2].

The 70.0° and 97.6° divisions extend the system into higher-latitude regimes, but do so with smaller populations and distinct orbital character. The 70.0° division remains prograde and shares much of the drag-sensitive operational environment of the mid-inclination divisions. The 97.6° division is slightly retrograde by inclination and occupies a geometrically distinct near-polar regime. At the level of licensed static architecture, these higher-latitude divisions probe contrasting dynamical environments within a single integrated satellite constellation [61, 1, 2].

1.2.3 Catalog-Space Geometry

The Gen 1 Starlink satellite constellation is physically organized as a shell-based architecture in orbital space, but public analysis is conducted in catalog space. Since operator ephemerides, maneuver logs, and satellite-resolved osculating states are not publicly available, the observable geometry must be reconstructed from GP and TLE records, which encode structures that remain interpretable at the constellation-division scale [16, 17, 20, 18, 19]. The resulting public-data geometry remains division-resolved in three-dimensional orbital space, as shown in Figure 1.11. The visualization provides a direct schematic bridge between the licensed shell architecture and the catalog-derived geometric structure used throughout the remainder of the case study.




Figure 1.11: Combined Three-Dimensional Geometry of the Principal Gen 1 Starlink Divisions

Direct Catalog Coordinates

Published mean elements, particularly the RAAN and mean anomaly, provide the most direct basis for initial geometric orientation in the public archive [16, 17, 1, 18, 19]. These quantities are encoded directly in catalog products and preserve immediate geometric interpretability. RAAN defines inertial plane orientation, while mean anomaly indexes along-track phase within the SGP4-compatible mean-element description. Although the pair (Ω, M) does not determine the full orbital state, it reveals the distribution of satellites across planes and the corresponding phase occupancy within each shell [1, 2, 16, 17].

Direct use of RAAN and mean anomaly is, therefore, methodologically important at the start of the analysis. More abstract embeddings and heavily derived diagnostics introduce additional modeling assumptions that are unnecessary for the first descriptive pass. By contrast, (Ω, M) exposes plane-phase geometry directly in catalog coordinates and provides an immediate observational description of the publicly visible Starlink system [16, 17, 18, 19].

Plane Orientation and Along-Track Structure

Within a shell-resolved constellation, RAAN and mean anomaly exhibit non-uniform structure shaped by plane count, shell inclination, deployment chronology, and maintenance regime. A combined RAAN-versus-mean-anomaly representation, therefore, expresses the constellation as a division-conditioned distribution over inertial plane orientation and along-track phase. At this level, the archive already reveals a structured plane-resolved system with distinct division densities, occupancy patterns, and secular-drift signatures [61, 16, 17, 18, 19].




Figure 1.12: Combined RAAN versus Mean Anomaly Geometry of the Principal Gen 1 Starlink Divisions

Figure 1.12 provides the first global geometric view of that archive-level structure rather than a complete dynamical interpretation. The dense mid-inclination divisions appear as populated plane-phase lattices, whereas the 70.0° and 97.6° divisions appear sparser and exhibit visibly distinct plane-occupation and drift morphology. In this way, the visualization links the licensed shell architecture to the historical orbital morphology developed in the sections that follow [61, 1, 2].

Vector Geometry and Near-Circular Structure

Orbital-vector representations provide a complementary geometric framework [1, 2, 3]. The vectors ($\mathbf{r}$) and ($\mathbf{v}$) determine instantaneous translational kinematics, the angular-momentum vector $\mathbf{h}$ defines orbital-plane orientation, the nodal vector ($\mathbf{n}$) specifies the line of nodes, and the eccentricity vector ($\mathbf{e}$) encodes apsidal direction and magnitude. This representation places all divisions in a common Euclidean embedding and suppresses some of the visual discontinuities associated with angular wraparound in scalar element coordinates. In the present context, the goal is geometric legibility rather than dimensionality reduction. The vectors $\mathbf{h}$ and $\mathbf{n}$ separate the divisions through their plane-orientation structure, while $\mathbf{e}$ remains concentrated near the origin because the operational archive is dominated by near-circular orbits, as shown in Figure 1.13. That concentration indicates that inclination, semi-major axis, RAAN, and secular drift govern most of the large-scale division structure examined in the subsequent analysis, whereas apsidal diagnostics remain secondary [1, 2, 61].




Figure 1.13: Combined Orbital Vector Representations of the Principal Gen 1 Starlink Divisions

1.2.4 Aggregate Historical Morphology

Following the initial geometric characterization, inspection of the public archive reveals a first-order morphological structure that is already visible before clustering, latent embedding, or other higher-level inference. In the Gen 1 Starlink system, the large-scale division architecture remains persistently identifiable through time. Aggregate archival data separate division membership, altitude-energy stratification, and low-eccentricity operating domains, so the primary constellation skeleton is recoverable from direct or minimally transformed catalog observables [61, 16, 17, 18, 19].

Division Structure

Aggregate inclination history provides the most direct analytical view of division structure because inclination is a cataloged quantity rather than a latent coordinate inferred through secondary modeling [16, 17, 18, 19]. In the Gen 1 archive, inclination appears in discrete and persistent bands aligned with the licensed shell divisions instead of filling a continuous domain [61]. Global inclination distributions further show that these bands persist over extended intervals, indicating a durable division structure rather than a transient deployment artifact.




Figure 1.14: Aggregate Inclination of the Gen 1 Starlink Satellite Constellation

As shown in Figure 1.14, the aggregate inclination history separates the principal divisions into persistent bands aligned with the licensed shell inclinations. The observable multi-division inclination structure shows that the archive is organized in a physically meaningful way prior to unsupervised analysis. Subsequent clustering can sharpen internal substructure and isolate transitional populations, but it does not create the principal constellation architecture. The aggregate

archive already contains a historically persistent division system whose core organization is directly encoded in catalog coordinates [61, 4, 18, 19].

Aggregate Lifecycle Coordinate

Inclination separates division membership, while semi-major axis provides the most effective aggregate coordinate for lifecycle progression [16, 17, 1]. In the present setting, the semi-major axis is obtained from the catalog mean motion through the standard Keplerian relation

$$a = \left(\frac{\mu_{\oplus}}{n^2} \right)^{1/3} \quad (1.30)$$

which preserves close fidelity to the catalog layer while yielding a physically interpretable proxy. At the aggregate scale, the history of a functions as a global lifecycle map that captures deployment, service residence, and retirement more effectively than a sequence of raw element values.




Figure 1.15: Aggregate Semi-Major Axis of the Gen 1 Starlink Satellite Constellation

The aggregate semi-major-axis morphology in Figure 1.15 exhibits two dominant structures. Persistent dense bands correspond to the service-shell altitudes specified by the licensed architecture [61]. Below these bands, more diffuse populations represent transitional regimes associated with insertion, orbit raising, delayed shell entry, lowering, and retirement. These two structures coexist throughout the historical record. A mature low-altitude megaconstellation is therefore characterized by simultaneous service-shell occupancy and continuous flux through lower-altitude transitional regimes rather than by a static on-station manifold [65, 63, 66].

Low-Eccentricity Conditioning

The aggregate eccentricity history shows that the constellation remains concentrated in the low-eccentricity regime. The operational system is therefore best described as a collection of near-circular shell populations, with only modest eccentricity excursions during particular operational phases. This empirical conditioning constrains the diagnostic choices used in the subsequent sections.

In low-eccentricity regimes, inclination, semi-major axis, RAAN, and division-wise nodal drift retain strong physical interpretability and numerical stability [1, 2]. As eccentricity approaches zero, apsidal diagnostics lose robustness because the line of apsides becomes weakly defined. The aggregate eccentricity history shows that the constellation remains concentrated in the low-eccentricity regime, as shown in Figure 1.16. The operational system is therefore best described as a collection of near-circular shell populations, with only modest eccentricity excursions during particular operational phases. This empirical conditioning constrains the diagnostic choices used in the subsequent sections.




Figure 1.16: Aggregate Eccentricity of the Gen 1 Starlink Satellite Constellation

In low-eccentricity regimes, inclination, semi-major axis, RAAN, and division-wise nodal drift retain strong physical interpretability and numerical stability [1, 2, 61].

1.2.5 Lifecycle Corridors

Gen 1 Starlink exhibits persistent multi-division stratification in the aggregate public archive. Division-resolved partitioning is necessary for identifying occupied manifolds, transfer corridors, and terminal trajectories because the licensed constellation uses distinct shell patterns across inclinations [61, 18, 19]. Each division exhibits distinct density distributions, shell-occupancy structure, and lifecycle morphology. This partitioning gives physical meaning to the internal archive structure and supports higher-resolution operational interpretation [65, 63, 68].

Lifecycle Corridors and Occupied Manifolds

Division-resolved inclination-versus-semi-major axis diagrams function as dynamic lifecycle maps. Inclination identifies division membership, while semi-major axis quantifies orbital size, altitude, and migration through insertion, ascent, residence, and descent phases [1, 16, 17, 18, 19]. Compact manifolds indicate prolonged shell occupancy, whereas broader lower-altitude fans and branching structures indicate transfer operations, deferred entry, relocation, and end-of-life pathways.

This mapping is consistent with published Starlink mission profiles that include low-altitude insertion, orbit raising toward service shells, station-keeping, and controlled deorbit. Filings for the modified 550 km shell indicate insertion generally near 300–350 km, while Starlink-specific orbit-determination studies indicate parking-stage behavior and extended orbit-raising histories that remain visible in public tracking data [70, 67, 65, 68]. The resulting division-resolved structure in (i, a) space provides a public-data signature of deliberate orbital circulation governed by operational transfer, raising, and disposal dynamics.

The 53° Complex

The 53° complex forms the principal dynamical population and is best interpreted as a coupled operational basin composed of the 53.0°/550 km and 53.2°/540 km divisions. Its inclination–semi-major-axis structure contains two dense occupied manifolds connected by broad lower-altitude populations. The persistence, coherence, and repeated traffic through this compound structure are inconsistent with a purely stochastic deployment pattern [61, 65, 63].




Figure 1.17: Inclination versus Semi-Major Axis for the 53° Complex

Figure 1.17 shows the compound lifecycle geometry of the 53° complex in inclination–semi-major-axis space, including the adjacent dense shell manifolds and the lower-altitude corridors associated with transfer, raising, and descent behavior. The licensed architecture places the two largest Gen 1 divisions in adjacent low-altitude mid-inclination shells [61]. As a result, the 53° complex captures the broadest visible range of lifecycle behavior, including frequent orbit raising, replenishment, shell-maintenance dispersion, relocation, and retirement flow. The upper manifolds indicate dense shell residence, while the lower fans and descending branches record circulation through a shared low-altitude state space.




Figure 1.18: Inclination versus Time of the 53° Complex

The inclination history in Figure 1.18 reinforces this interpretation. The two principal branches remain continuously populated over time, indicating temporal persistence beyond the transient launch structure. Expanding lower tails and transitional regions record ongoing entry and exit, while core shell occupancy remains present throughout the archive. This pattern is characteristic of a mature and actively managed division complex [63, 66].

The 70° Division

The 70° division provides a useful comparative reference. Its inclination–semi-major-axis distribution contains a single dominant manifold near the service shell, accompanied by subordinate lower-altitude corridors and fewer extended late-stage departures than the 53° complex. The morphology is correspondingly more compact, more easily interpreted, and still dynamically active [61, 65, 63].




Figure 1.19: Inclination versus Semi-Major Axis for the 70° Division

Figure 1.19 shows the comparatively compact lifecycle corridor of the 70° division, with a dominant service-shell manifold and subordinate lower-altitude transfer structure. The compactness of the 70° division shows that corridor–manifold morphology varies across the satellite constellation. Some divisions remain concentrated in narrower shell-centered regions, and the 70° case provides a useful baseline against which the broader circulation of the 53° complex and the layered structure of the 97.6° division can be compared.




Figure 1.20: Inclination versus Time of the 70° Division

The inclination history in Figure 1.20 remains concentrated near the nominal value, with departures appearing only as secondary features. This temporal compactness confirms the visual impression of relative stability and supports the use of the division as a reference case in subsequent diagnostics [61, 63].

The 97.6° Division

The 97.6° division is less densely populated than the mid-inclination core, and it remains dynamically informative. Division-resolved plots reveal a layered high-latitude population consisting of a dense service-shell region, a secondary lower-population branch, and visible internal connections between these structures. Before clustering, the public archive already shows nontrivial internal organization [61].




Figure 1.21: Inclination versus Semi-Major Axis for the 97.6° Division

Figure 1.21 shows the layered morphology of the 97.6° division in inclination–semi-major-axis space, including the dense near-service-shell population, lower-altitude branches, and transitional connections among these regions. The high-latitude Starlink structure differs dynamically from the mid-inclination divisions. As this division occupies a slightly retrograde high-latitude regime, it evolves under a distinct secular environment. Its division-resolved morphology nevertheless shows persistent internal layering and a clear coexistence of shell residence and lower-altitude corridor structure [61, 1, 2].




Figure 1.22: Inclination versus Time of the 97.6° Division

The inclination history in Figure 1.22 makes this temporal layering explicit. Persistent occupation near the nominal inclination, together with lower branches and transitional connections, indicates enduring internal organization across the archive. These features support treating the 97.6° population as a distinct division for later secular and clustering analysis [61].

1.2.6 Secular Dynamics and Drift

The previously established division-resolved morphology captures the leading-order secular response of each shell to Earth’s oblateness through the J_2 term. Grouped J_2 precession rates therefore provide a rigorous characterization of the division hierarchy. Nodal precession serves as the primary secular diagnostic, and apsidal precession remains supplementary because low-eccentricity geometry reduces the robustness of apsidal observables as $e \rightarrow 0$ [1, 2, 3, 61].

J_2 Nodal Regression

The first-order J_2 nodal rate provides the most informative secular diagnostic for the present analysis, with the formulation written as

$$\dot{\Omega}_{J_2} = -\frac{3}{2}J_2 n \left(\frac{R_{\oplus}}{p_{\text{slr}}} \right)^2 \cos i, \quad p_{\text{slr}} = a(1 - e^2) \quad (1.31)$$

Here, a , e , and i denote semi-major axis, eccentricity, and inclination, respectively. The sign and magnitude of $\dot{\Omega}_{J_2}$ depend jointly on shell altitude and inclination. Lower semi-major axis increases the magnitude of the drift. Prograde shells with $\cos i > 0$ undergo westward nodal regression, while the 97.6° division, for which $\cos i < 0$, undergoes eastward nodal advancement [1, 2, 61].




Figure 1.23: Grouped First-Order J_2 Nodal Precession Rates of the Principal Gen 1 Starlink Divisions

The grouped nodal-precession data in Figure 1.23 confirm this hierarchy. The 53° complex

exhibits the strongest negative regression because it combines low altitude with prograde inclination. The 70° division exhibits weaker negative regression as a consequence of its higher altitude and smaller $|\cos i|$. The 97.6° division exhibits positive nodal drift, consistent with its slightly retrograde inclination. This ordering shows that the observed inclination–semi-major axis structure is dynamically consistent with first-order oblateness theory [1, 2, 61].

Nodal histories further clarify the secular evolution of the shell divisions. Populations experience differential regression as satellites move through raising, residence, and lowering phases. The magnitude of the drift is controlled primarily by orbital size, while the sign is set by inclination. Grouped nodal-rate analysis, therefore, links the static shell configuration to the secular processes governing long-horizon orbital evolution [1, 2].

Apsidal Rates and the Low- e Limitation

Apsidal precession provides a secondary secular check through the first-order rate, formulated as

$$\dot{\omega}_{J_2} = \frac{3}{4} J_2 n \left(\frac{R_\oplus}{p_{\text{slr}}} \right)^2 (5 \cos^2 i - 1) \quad (1.32)$$

This expression separates the 53° complex from the 70° and 97.6° divisions through the inclination-dependent factor $5 \cos^2 i - 1$ [1, 2]. However, it is less robust than nodal precession as a population-level diagnostic for the present Starlink archive.

In near-circular orbits, the line of apsides becomes weakly defined as $e \rightarrow 0$, which makes apsidal quantities more sensitive and less stable than nodal quantities in aggregate analysis [1, 2, 3]. The grouped apsidal-rate results therefore serve mainly as a broad consistency check against J_2 secular theory while simultaneously illustrating the reduced reliability of apsidal diagnostics in low-eccentricity populations.

Figure 1.24 provides the corresponding apsidal-rate comparison across the principal divisions. The separation among the grouped rates is consistent with the inclination-dependent factor in the first-order J_2 expression, while the low-eccentricity operating regime limits the diagnostic weight assigned to apsidal structure in the later interpretation.




Figure 1.24: Grouped First-Order J_2 Apsidal Precession Rates of the Principal Gen 1 Starlink Divisions

1.2.7 Maturity, Turnover, and Inferred Operational State

The division-wise secular structure indicates that the Gen 1 Starlink constellation has reached a regime of managed orbital regulation. The remaining question is whether the public archive also encodes operational maturity. In low-altitude megaconstellations, maturity is expressed through the simultaneous persistence of shell residence, shell exit, and recurring phase structure associated with raising, occupation, lowering, and decay, all of which can leave observable signatures in historical tracking data [63, 65, 68, 66]. The present analysis adopts a conservative interpretation: these signatures are treated as public-data diagnostics of operational state and are not used to reconstruct proprietary operator timelines or internal decision logic.

Shell Compactness and Persistent Exit

The shell-width metric serves as a proxy for shell compactness or occupancy spread by summarizing the distribution of the population within an orbital corridor at each epoch. This abstraction is useful because it highlights temporal changes in operational concentration. Narrow intervals indicate concentrated residence near the service-shell manifold, while broadened intervals indicate the superposition of shell occupancy with raising, relocation, lowering, or retirement traffic [63, 66].




Figure 1.25: Shell Compactness/Occupancy Spread Chronology

Figure 1.25 shows the time evolution of this occupancy-spread metric across the principal shell groupings, making visible the alternation between compact shell residence and broadened transitional intervals. This distinction is operationally important because constellation maturity is characterized by compact operational occupancy together with persistent dispersive traffic through adjacent altitude bands. In that sense, the Gen 1 archive records the transition from deployment-centered growth to a continuously cycled industrial orbital system [63, 66].

State counts and disposal-onset chronology provide an additional diagnostic of maturity by distinguishing archival intervals dominated by insertion and shell growth from intervals that also contain persistent retirement-related activity. Early epochs are primarily associated with insertion, raising, and active-shell growth, whereas later epochs display sustained retirement signatures with continued shell occupancy. The appearance and persistence of disposal-like states in the census, therefore, mark the onset of turnover [70, 66].




Figure 1.26: Sustainability-State Counts Chronology

Figure 1.26 summarizes the population-level evolution of active-like, disposal-like, and drifting-like states, showing the transition from growth-dominated occupancy toward a regime in which shell residence and retirement activity coexist.




Figure 1.27: Disposal-Onset Chronology

Disposal-onset chronology is best interpreted as a population-level rate process, with episodic bursts and sustained onset periods indicating that retirement activity is organized at the fleet scale

and leaves a coherent historical signature in the public archive. Although detailed scheduling, prioritization, and command logic remain inaccessible, the aggregate structure of retirement remains observable [70, 66]. Figure 1.27 isolates the chronology of disposal-onset events, providing a sharper temporal diagnostic of fleet turnover than the aggregate state-count census alone.

Chapter 2

High-Fidelity Orbital Simulator Development

2.1 Propagation Architecture, State Definition, and Numerical Strategy

This section defines the numerical and state-space architecture of the orbital simulation framework. The propagated variables, frame and epoch conventions, temporal discretization principles, and adaptive integration strategy are specified before the perturbation models are introduced. A single Cartesian initial-value problem (IVP) is used for coast, drag-perturbed, and low-thrust cases, so each additional physical effect enters the same mathematical propagation structure through the force model without requiring a reformulation of the governing initial-value problem [1, 2, 3].

2.1.1 Simulation Scope and Propagated State

The orbital simulation is posed as a Cartesian first-order IVP in which the Sun, Moon, and satellite are stored in a common propagated state description. Perturbing effects are incorporated at the acceleration level, so increases in force-model fidelity do not require force-model-specific changes to the state architecture [1, 2, 3]. Two propagated-state configurations are used. The coast formulation advances an 18-dimensional state containing the Cartesian position and velocity components of the Sun, Moon, and satellite. The low-thrust formulation augments the same translational structure with the satellite mass, thereby yielding a 19-dimensional variable-mass state. The translational equations retain the same Cartesian form under thrusting, while propulsion modifies the dynamics

through an added acceleration term and a single scalar mass-depletion equation.

$$\mathbf{y}(t) = \begin{cases} \left[\mathbf{r}_{\odot}^T, \mathbf{v}_{\odot}^T, \mathbf{r}_{\zeta}^T, \mathbf{v}_{\zeta}^T, \mathbf{r}_{\text{sat}}^T, \mathbf{v}_{\text{sat}}^T \right]^T \in \mathbb{R}^{18}, & \text{coast propagation} \\ \left[\mathbf{r}_{\odot}^T, \mathbf{v}_{\odot}^T, \mathbf{r}_{\zeta}^T, \mathbf{v}_{\zeta}^T, \mathbf{r}_{\text{sat}}^T, \mathbf{v}_{\text{sat}}^T, m_{\text{sat}} \right]^T \in \mathbb{R}^{19}, & \text{low-thrust propagation} \end{cases} \quad (2.1)$$

Each body contribution consists of Cartesian position and velocity components in the common propagation frame. The mass state appears only in variable-mass thrusting cases. The governing dynamics are written in first-order form as

$$\dot{\mathbf{y}} = \mathbf{f}_{\text{prop}}(t, \mathbf{y}; \Theta_{\text{sim}}, \mathcal{E}(t)) \quad (2.2)$$

where Θ_{sim} collects fixed simulation and force-model parameters, and $\mathcal{E}(t)$ denotes time-dependent environmental inputs such as third-body ephemerides, terrestrial rotation, and atmospheric density. This formulation makes the structural organization of the orbit simulation explicit. Gravitational, drag, SRP, and low-thrust contributions are assembled within a unified right-hand side (RHS), preserving a common propagation interface across distinct physical regimes [1, 2, 3].

The propagated state provides the principal organizational structure of the framework. Coast propagation, variable-mass thrust propagation, single-satellite runs, and ensemble simulations all inherit this common state-space architecture. Improvements in physical fidelity can therefore be introduced through changes to the constituent force models or to the thrust scheduling logic, without altering the underlying numerical formulation.

2.1.2 Frames, Epochs, and Time Variables

Frame and epoch conventions are chosen to support accurate long-arc propagation under rotating-Earth dynamics, ephemeris-dependent forcing, and terrestrial-frame environmental evaluations [71, 1, 2]. Earth is treated as the central body, and the translational state is propagated in a geocentric inertial frame consistent with the celestial reference convention adopted by the orbit simulation, while Earth-fixed evaluations are carried out through the corresponding terrestrial transformation chain [71, 1]. The initial epoch t_0 , provided as a calendar date in UTC, is converted internally to the corresponding Julian date JD_0 together with the Earth-rotation quantity required by the selected terrestrial-to-celestial transformation. In the present implementation, this rotational quantity is represented through a Greenwich sidereal angle, while the same role may be expressed in an Earth Rotation Angle formalism within the IERS convention set. These quantities are required by rotating gravity and atmospheric models, ensuring that inertial propagation remains synchronized with force

and geometry evaluations performed in the Earth-fixed frame [71, 1].

Major-body initialization is epoch dependent, with Sun and Moon states evaluated from date-specific ephemeris data. The Sun and Moon state vectors relative to Earth are evaluated for the specified date in both position and velocity and are used to initialize the external-body forcing geometry. Earth-rotation quantities required for frame conversion and terrestrial environmental evaluation are generated from the same epoch definition. This procedure preserves physical consistency in the initial celestial geometry, which is required for third-body perturbations, eclipse determination, and atmosphere-coupled propagation over extended simulation intervals [72, 73, 1].

Time management separates user-facing controls from the internal numerical propagation variable. The simulation horizon is specified in days through t_u , with a final propagation time t_f and a stored output increment Δt_{out} , all of which are converted internally to seconds for segment assembly and solver execution. This separation is necessary because the archived analysis grid and the adaptive integrator step sequence serve different computational roles. The stored grid provides a uniform sampling structure for tracking derived quantities such as semi-major axis, eccentricity, inclination, and satellite mass, while the integrator advances on a variable internal sequence determined by local error control and event localization [74, 75, 76].

This temporal architecture decouples the physical simulation horizon from the internal adaptive step sequence, enabling consistent analysis over long windows without imposing a globally fixed integration step. The same design accommodates secular drift, event-driven termination, and discontinuous schedule boundaries within a single propagation workflow.

2.1.3 Adaptive Integration and Event Handling

The numerical integrator is chosen to match the dominant character of the governing equations over the intended operating regime. The propagated dynamics are high-dimensional and strongly perturbed and are treated as nonstiff within the simulation scenarios considered here. The framework therefore employs DOP853, an explicit eighth-order Runge-Kutta method, through an adaptive IVP solver interface [74, 77, 75, 76]. This choice is well-suited to long-horizon, non-stiff orbital propagation under stringent local tolerances. In segmented dynamic-atmosphere runs, $\text{rtol} = 10^{-10}$ and $\text{atol} = 10^{-12}$ are used in order to tightly control local truncation error and event timing.

Adaptive integration is central to the simulation strategy. The framework resolves cumulative secular effects, phase evolution, and boundary-sensitive transitions without imposing a global step-size restriction. High-order adaptive propagation maintains computational efficiency through dynamically selected internal steps in smooth regions and preserves accuracy near rapid geometric variation, atmospheric updates, and terminal boundaries [74, 75, 77].

Event handling is embedded directly in the propagation architecture. Terminal descending-crossing events are defined at $h_{\text{re}} = 100$ km and $h_{\text{cut}} = 115$ km, corresponding respectively to re-entry and batch-cutoff surfaces, and are triggered only on downward threshold crossings. The global propagation interval is segmented whenever maneuver-schedule boundaries are encountered or when dynamic-atmosphere updates require the loading of new density fields. The full trajectory is therefore assembled as an ordered sequence of locally smooth propagation segments, each initialized from the terminal state of the preceding segment [75].

This segmented, event-driven structure combines adaptive high-order integration with discontinuous schedule logic and day-resolved atmospheric updates. Schedule boundaries, atmospheric transitions, and terminal altitude surfaces act as propagation controls. The propagation loop initializes state and forcing data, advances through adaptive integration within each segment, monitors events and boundary conditions continuously, restarts the solver where required, and returns a uniform archived output set for subsequent physical analysis. This process forms the computational backbone of the simulation framework.

2.2 Initial-Condition Construction and Multi-Satellite Case Assembly

The initialization layer converts raw catalog data into propagation-ready states, supplements each case with epoch-consistent Sun and Moon ephemerides, and constructs either a single trajectory or a multi-satellite ensemble within a unified simulation structure. The architecture is designed from the outset for satellite constellation-oriented case construction, so subsequent propagation stages operate on structured orbital populations as the primary case abstraction [18, 19, 16, 17, 20].

2.2.1 Manual and Catalog-Driven Initialization

The framework supports both manual and catalog-driven initialization since the two pathways serve differing analytical purposes. Manual initialization provides controlled reference cases for development, verification, and force-model isolation, while catalog-driven initialization ingests orbital records and converts those into simulation-ready states for historical or operational population studies. The two pathways converge at a common state-construction layer, so downstream propagation is defined by a single numerical interface instead of source-specific logic.

In the manual pathway, the satellite initial condition may be specified directly in Cartesian form or constructed from a consistent orbital-element description under the frame and epoch conventions

defined for the orbital simulation. In the catalog-driven pathway, public GP records are organized into a structured archive $\mathcal{A}_{\text{cat}}$ containing epoch, inclination, RAAN, eccentricity, argument of perigee, mean anomaly, mean motion, and associated metadata for each candidate satellite [18, 19, 16, 17, 20]. These records are treated as model-dependent mean-element products, such as TLEs or equivalent OMM fields, with their associated GP semantics.

For that reason, the propagation-ready Cartesian initial condition is formed through a model-consistent SGP4 translation. For each selected record $\mathcal{R}_{s,0}$, the framework evaluates the compatible SGP4 state at the record epoch $t_{s,0}$, obtains the corresponding position and velocity in the catalog reference convention, typically TEME for TLE-based initialization, and maps that state into the orbital simulation propagation frame when required [16, 17, 1, 19, 20]. This preserves consistency between the public catalog observation layer and the initialized numerical state supplied to the propagator.

Quantities such as semi-major axis, true anomaly, and specific angular momentum may still be computed during initialization, but are used as derived diagnostics. The fundamental propagated Cartesian state is obtained from the compatible catalog propagation model [1, 2, 16, 17]. This distinction is necessary because public GP records compress observations through a theory-dependent mean-element representation intended for compatible analytical propagation. The initialization layer, therefore, functions as a structured translation interface between catalog-space mean-element products and the orbital simulation state vector $\mathbf{y}_{s,0}$.

Propagation-case selection filters $\mathcal{A}_{\text{cat}}$ by satellite identifier, admissible epoch window, and any mission-specific constraints, then assigns the most recent eligible record on or before the requested start time to each selected satellite. A minimum initialization altitude $h_{\text{min,init}}$ is imposed so that cases are not started below the operational regime of interest or within immediate terminal descent. Each accepted case is converted into an initialized propagated state $\mathbf{y}_{s,0}$, and the accepted satellites define the index set $\mathcal{I}_0$. The assembled case set is written as

$$\mathcal{C}_0 = \{(s, t_{s,0}, \mathbf{y}_{s,0}) : s \in \mathcal{I}_0\} \quad (2.3)$$

which supports either a single-satellite run or an ensemble propagation without any change to the downstream numerical architecture. The initialization stage, therefore, provides the structured interface that connects public orbital records, epoch-dependent Sun and Moon ephemerides, and multi-satellite case assembly to the common propagator state used throughout the simulation framework.

2.2.2 Epoch-Conditioned Major-Body Initialization

External-body geometry is initialized at the same epoch that defines each satellite’s initial state, so the celestial and terrestrial components of the force model begin from a common temporal reference. For each initialized case, the framework constructs an epoch-conditioned external-body package, represented as

$$\mathcal{E}_{s,0}^{\text{ext}} = \{\mathbf{r}_{\odot,s,0}, \mathbf{v}_{\odot,s,0}, \mathbf{r}_{\zeta,s,0}, \mathbf{v}_{\zeta,s,0}, JD_{s,0}, \theta_{\oplus,s,0}\} \quad (2.4)$$

where the Sun and Moon states are evaluated relative to Earth at the initialization epoch $t_{s,0}$ and expressed in the orbit propagation frame and unit system. Earth remains the central body of the propagated satellite dynamics, so terrestrial rotation enters through epoch-conditioned rotation quantities rather than through an independently propagated Earth state vector. This initialization defines the geometry required for third-body perturbations, SRP, eclipse determination, and terrestrial-frame environmental evaluation [72, 73, 71, 1, 2].

The ephemeris layer is designed for repeated simulation rather than one-time preprocessing. Queries keyed by body, reference frame, and epoch are stored in an in-memory cache $\mathcal{K}_{\text{eph}}$, with optional persistence to disk for reuse across repeated runs. This design reduces redundant ephemeris evaluation during ensemble studies, parameter sweeps, and repeated date-conditioned analyses. The initialization stage, therefore, produces a reusable epoch-conditioned data object.

Epoch logic also supplies the temporal quantities required by rotating-Earth force models. Once the initialization date is fixed, the framework computes the corresponding Julian date and Earth-rotation quantity, implemented in the present framework through $\theta_{\oplus,s,0}$ and compatible with either Greenwich sidereal time or Earth Rotation Angle conventions, and propagates these values into the gravity, frame-conversion, and atmospheric components [71, 1]. Initialization thus combines the satellite state, external-body geometry, and terrestrial rotation data into a single epoch-specific configuration. This consistency is important for long-arc propagation because small initialization mismatches can contaminate subsequent force evaluations, event timing, and case-to-case comparisons.

2.2.3 Single-Satellite and Ensemble Case Assembly

A single archived satellite is treated as a particular case within a broader initialized population. Single-satellite propagation corresponds to the special case $N_{\text{ens}} = 1$, while ensemble propagation uses the same state construction logic over a larger initialized index set $\mathcal{I}_0$. In both settings, the

propagation-ready case set is assembled as

$$C_0 = \left\{ \left(s, t_{s,0}, \mathbf{y}_{s,0}, \mathcal{E}_{s,0}^{\text{ext}} \right) : s \in \mathcal{I}_0 \right\} \quad (2.5)$$

so downstream propagation operates on a common case description regardless of population size. This architecture eliminates the need for a separate numerical workflow for satellite constellation studies and preserves a common physical interpretation across one-satellite and many-satellite runs.

The ensemble pathway is governed by initialization-window limits, filtering logic, and execution controls without any change to the propagated state definition. Catalog records are restricted to an admissible epoch window, each accepted satellite is paired with an epoch-conditioned external-body package, and the resulting case set is dispatched for batch execution subject to a fixed worker cap N_{worker} . The transition from nominal propagation to many-satellite analysis is therefore implemented through changes in case count, scheduling, and resource allocation, with the mathematical formulation fixed.

A central consequence of this design is that each trajectory retains the identity of its initialized archived object from the start of the simulation. Satellite identifier, source epoch, initialized Cartesian state, and epoch-conditioned environmental data are attached at case construction and carried into the propagation output. The resulting trajectory products can therefore be interpreted as physically dated continuations of specific archived objects, which is essential for division-level comparisons, phase-dependent forcing analysis, and batch diagnostics across large satellite sets.

2.3 Earth Gravity and Conservative Perturbation Modeling

The conservative forcing layer is formulated by combining Earth’s central attraction, a truncated spherical-harmonic representation of the terrestrial gravitational field, and differential third-body accelerations from the Sun and Moon within a single Cartesian acceleration assembly. Each contribution is derived from a potential or potential difference, persists over the full propagation interval, and acts on the same propagated state. This structure is appropriate for long-arc orbit propagation because it preserves a unified force-model hierarchy while allowing progressively higher fidelity within the same numerical framework [1, 2, 3, 37].

2.3.1 Central Gravity and Truncated Spherical Harmonics

The terrestrial conservative acceleration is decomposed as

$$\mathbf{a}_{\text{grav}}(\mathbf{r}, t) = \mathbf{a}_0(\mathbf{r}) + \mathbf{a}_{\text{harm}}(\mathbf{r}, t) \quad (2.6)$$

where the central term is

$$\mathbf{a}_0(\mathbf{r}) = -\frac{\mu_{\oplus}}{r^3} \mathbf{r} \quad (2.7)$$

with $r = \|\mathbf{r}\|$. This contribution defines the underlying two-body motion, while the harmonic correction accounts for the nonspherical terrestrial mass distribution that drives the dominant conservative deviations from Keplerian dynamics in LEO [1, 2, 37].

To represent the nonspherical field, the framework uses an EGM2008-based geopotential model truncated to

$$n_{\text{max}}^{\text{grav}} = m_{\text{max}}^{\text{grav}} = 50 \quad (2.8)$$

This truncation is a modeling choice for the present orbital simulation rather than a property of EGM2008 itself. The full EGM2008 model extends to much higher degree and order, but a degree-and-order-50 representation retains the dominant low-altitude conservative structure needed here while remaining computationally manageable for repeated and ensemble propagation [37, 1, 2].

In the terrestrial frame, the gravitational potential is written as

$$U_{\oplus}(r, \phi, \lambda) = \frac{\mu_{\oplus}}{r} \left[1 + \sum_{n=2}^{n_{\text{max}}^{\text{grav}}} \left(\frac{R_{\oplus}}{r} \right)^n \sum_{m=0}^{\min(n, m_{\text{max}}^{\text{grav}})} \bar{P}_{nm}(\sin \phi) (\bar{C}_{nm} \cos m\lambda + \bar{S}_{nm} \sin m\lambda) \right] \quad (2.9)$$

where ϕ and λ are the geocentric latitude and longitude in the terrestrial frame, $\bar{C}_{nm}$ and $\bar{S}_{nm}$ are fully normalized spherical-harmonic coefficients, and $\bar{P}_{nm}$ denotes the fully normalized associated Legendre functions [37, 71, 1]. The corresponding Earth-fixed gravitational acceleration is obtained from the potential gradient,

$$\mathbf{a}_{\text{grav}}^{\text{ITRS}} = \nabla U_{\oplus} \quad (2.10)$$

and the inertial-frame acceleration used by the propagator is then

$$\mathbf{a}_{\text{grav}}(t) = Q(t) R(t) W(t) \mathbf{a}_{\text{grav}}^{\text{ITRS}}(t) \quad (2.11)$$

In this form, the harmonic field is fixed in the terrestrial frame but becomes explicitly time dependent in the inertial propagation frame through Earth rotation [71, 1].

Over the altitude range relevant to LEO satellite constellations, the nonspherical terrestrial field sets the principal conservative structure beyond central attraction, governs secular precession of orbital geometry, and establishes the background trajectory on which drag, SRP, and thrust act. The Earth gravity model, therefore, serves as the conservative backbone of the full force hierarchy used in the orbital simulation [1, 2, 37].

2.3.2 Third-Body Gravitational Forcing

The Sun and Moon are treated as the dominant external gravitational perturbers in the present Earth-centered propagation model. Their geocentric states are supplied to the force model from epoch-conditioned ephemerides, so the third-body geometry remains consistent with the orbital simulation time standard, celestial-terrestrial frame handling, eclipse logic, and SRP geometry [72, 73, 71, 1, 2]. Let the retained third-body set be

$$\mathcal{B}_{3b} = \{\odot, \mathbb{C}\} \quad (2.12)$$

For each $b \in \mathcal{B}_{3b}$, with geocentric position $\mathbf{r}_b(t)$, gravitational parameter μ_b , and satellite position $\mathbf{r}(t)$, the differential third-body acceleration is written as

$$\mathbf{a}_{3b}^{(b)}(\mathbf{r}, t) = \mu_b \left(\frac{\mathbf{r}_b(t) - \mathbf{r}}{\|\mathbf{r}_b(t) - \mathbf{r}\|^3} - \frac{\mathbf{r}_b(t)}{\|\mathbf{r}_b(t)\|^3} \right) \quad (2.13)$$

The first term is the direct attraction of body b on the satellite, and the second is the indirect term that removes the corresponding acceleration of the geocenter. This yields the physically relevant geocentric tidal forcing and preserves consistency with the Earth-centered equations of motion [1, 2].

The total third-body contribution is

$$\mathbf{a}_{3b}(\mathbf{r}, t) = \sum_{b \in \mathcal{B}_{3b}} \mathbf{a}_{3b}^{(b)}(\mathbf{r}, t) \quad (2.14)$$

Because $\mathbf{r}_b(t)$ varies with epoch, this forcing is explicitly time dependent even though each contribution remains non-dissipative and derives from a gravitational potential difference. The epoch-conditioned initialization of $(\mathbf{r}_{\odot}, \mathbf{v}_{\odot})$ and $(\mathbf{r}_{\mathbb{C}}, \mathbf{v}_{\mathbb{C}})$ ensures that third-body forcing, eclipse geometry, terrestrial rotation handling, and radiation-pressure geometry are evaluated from a common celestial configuration [72, 73, 71]. Over long LEO arcs, the lunar term is weak instantaneously but cumulative, while the solar term contributes both direct third-body forcing and the geometric

reference required by the subsequent SRP model. The Sun and Moon are therefore retained as the dominant external conservative bodies in the present orbital simulation [1, 2].

2.3.3 Conservative-Force Assembly in the Propagated Equations

The total conservative acceleration acting on the satellite is assembled as

$$\mathbf{a}_{\text{cons}} = \mathbf{a}_{\text{grav}} + \mathbf{a}_{3\text{b}} = \mathbf{a}_0 + \mathbf{a}_{\text{harm}} + \sum_{b \in \mathcal{B}_{3\text{b}}} \mathbf{a}_{3\text{b}}^{(b)} \quad (2.15)$$

Expanding the retained third-body set gives

$$\mathbf{a}_{\text{cons}} = \mathbf{a}_0 + \mathbf{a}_{\text{harm}} + \mathbf{a}_{3\text{b}}^{(\odot)} + \mathbf{a}_{3\text{b}}^{(\lrcorner)} \quad (2.16)$$

This decomposition preserves physical transparency and numerical consistency. The terrestrial gravity field sets the dominant conservative structure of the orbit, while the solar and lunar terms provide the time-dependent external gravitational environment. Each contribution is evaluated directly at the acceleration level and accumulated within the common Cartesian right-hand side of the propagator [1, 2, 37].

All conservative terms remain active throughout the propagation arc and therefore require no event logic or discontinuous scheduling. Their evaluation is local in time, but their accumulated influence can alter nodal evolution, apsidal evolution, and phase history over long horizons, especially when later coupled with drag, SRP, or low-thrust segments. The terrestrial geopotential and third-body fields therefore form a single conservative layer that establishes the physical baseline for all subsequent nonconservative perturbation models [1, 2, 3].

2.4 Solar Radiation Pressure, Eclipse Geometry, and Effective Optical Response

The SRP component of the orbital simulation framework is formulated as a coupled nonconservative forcing layer that combines photon momentum transfer, Sun-line visibility, and an effective optical-response surrogate. In LEO, SRP is typically weaker than atmospheric drag in the along-track direction, but it remains dynamically relevant over long arcs because it acts persistently, depends on the Sun-satellite geometry, and is modulated sharply at eclipse ingress and egress [1, 2, 78]. The framework, therefore, represents SRP as an explicit acceleration term whose magnitude is controlled

by illumination geometry, enabling force-consistent propagation across full-light, penumbral, and umbral conditions.

2.4.1 Solar Radiation Pressure Force Formulation

The direct SRP model uses the Sun-satellite distance $r_\odot$, the Sun-line unit vector $\hat{\mathbf{s}}$, the reference SRP at 1 AU, and an effective optical-response factor represented by $C_R A_{\text{SRP}}/m$. With $\mathbf{r}$ denoting the geocentric satellite position and $\mathbf{r}_\odot$ the geocentric Sun position, the Sun-satellite distance and line-of-sight direction are

$$r_\odot = \|\mathbf{r}_\odot - \mathbf{r}\|, \quad \hat{\mathbf{s}} = \frac{\mathbf{r}_\odot - \mathbf{r}}{r_\odot} \quad (2.17)$$

The SRP acceleration is then written as

$$\mathbf{a}_{\text{SRP}} = -\chi_{\text{ecl}} P_\odot \left(\frac{\text{AU}}{r_\odot} \right)^2 \left(\frac{C_R A_{\text{SRP}}}{m} \right) \hat{\mathbf{s}} \quad (2.18)$$

where $P_\odot$ is the reference SRP at 1 AU, C_R is the effective reflectivity coefficient, A_{SRP} is the effective SRP area, m is satellite mass, and $\chi_{\text{ecl}} \in [0, 1]$ is the eclipse factor [1, 2, 78]. This form separates inverse-square dilution, effective optical response, and eclipse modulation while retaining a compact acceleration-level model suitable for large-scale propagation.

The SRP term is evaluated directly within the Cartesian equations of motion and is accumulated with gravity, drag, and thrust contributions in the common propagated right-hand side. This arrangement preserves consistency across illumination states and ensures that SRP transitions are handled as part of the continuous force model within the common propagation architecture [1, 2].

2.4.2 Eclipse Geometry and Partial Shadowing

Shadowing is incorporated explicitly because the SRP magnitude depends on the visible fraction of the solar disc. The framework characterizes eclipse geometry through the apparent angular radii of the Sun and Earth as seen from the satellite, together with the angular separation of their apparent centers [1, 2]. These quantities are defined as

$$\psi_\odot = \arcsin\left(\frac{R_\odot}{r_\odot}\right), \quad \psi_\oplus = \arcsin\left(\frac{R_\oplus}{r}\right), \quad \gamma_{\text{ecl}} = \arccos\left(-\frac{\mathbf{r} \cdot \hat{\mathbf{s}}}{r}\right) \quad (2.19)$$

where $R_\odot$ and $R_\oplus$ are the solar and terrestrial reference radii used by the visibility model. These angles determine the illumination regime. Full illumination occurs when

$$\gamma_{\text{ecl}} \geq \psi_\oplus + \psi_\odot \quad (2.20)$$

Complete umbra occurs when the apparent solar disc is fully obscured by Earth, which, in the present spherical-disc construction, is represented by

$$\psi_\oplus \geq \psi_\odot \quad \text{and} \quad \gamma_{\text{ecl}} \leq \psi_\oplus - \psi_\odot \quad (2.21)$$

Penumbra occupies the intermediate overlap regime,

$$|\psi_\oplus - \psi_\odot| < \gamma_{\text{ecl}} < \psi_\oplus + \psi_\odot \quad (2.22)$$

The eclipse factor is assigned as

$$\chi_{\text{ecl}} = \begin{cases} 1, & \text{full illumination} \\ 0, & \text{complete umbra} \\ \chi_{\text{pen}}(\psi_\odot, \psi_\oplus, \gamma_{\text{ecl}}) \in (0, 1), & \text{penumbra} \end{cases} \quad (2.23)$$

where χ_{pen} is obtained from the visible fraction of the apparent solar disc under circular-disc overlap geometry [1, 2]. This continuous treatment is important for long-arc propagation because it avoids artificial force discontinuities across eclipse ingress and egress while still suppressing SRP completely in the umbra. The eclipse geometry, therefore, enters the orbital simulation as a direct dynamical modifier of the SRP acceleration.

2.4.3 Effective Optical Response and Projected-Area Surrogates

The orbital simulation does not resolve panel-level optical behavior or attitude-dependent surface exchange in full detail. Instead, it adopts an effective optical-response surrogate in which unresolved reflection, absorption, reradiation, and projected-area effects are consolidated into the parameters C_R and A_{SRP} . The resulting quantity $C_R A_{\text{SRP}}/m$ governs the translational sensitivity of the satellite to solar photon momentum transfer [1, 2, 78, 79].

The area A_{SRP} represents an effective projected area relative to the Sun line, and C_R functions as a surrogate momentum-transfer coefficient rather than as a unique material property. These

parameters, therefore, characterize the translational optical response relevant to orbit propagation, not a detailed reconstruction of satellite surface physics. Depending on the analysis context, these may be treated as fixed inputs, calibrated quantities, or phase-dependent effective parameters in later inference or identification stages [78, 79].

Within the present framework, this surrogate treatment preserves the correct scaling structure of SRP acceleration, supports explicit eclipse modulation, and remains computationally efficient for single-satellite and ensemble propagation. The SRP layer thus provides a physically interpretable nonconservative forcing model with explicit visibility geometry and reduced optical parameterization, while avoiding attitude-resolved surface complexity in the long-arc orbital simulation [1, 2, 78, 79].

2.5 Thermospheric Drag Coupling and Effective Aerodynamic State

Within the orbital simulation framework, the dissipative aerodynamic layer provides the principal mechanism of nonconservative orbital energy loss in LEO. Atmospheric drag drives secular decay of semi-major axis, reduces orbital energy, and amplifies sensitivity to environmental variability during later mission phases [1, 2, 12, 34, 35, 36]. The simulation architecture supports both a static reference atmosphere and a dynamic, space-weather-responsive pathway within a common Cartesian propagation framework, so aerodynamic forcing can be examined either under controlled baseline conditions or under higher-fidelity environmental variability.

2.5.1 Relative Velocity and Drag Acceleration

The drag formulation is referenced to the neutral atmosphere rather than to inertial orbital motion alone. Let $\mathbf{r}$ and $\mathbf{v}$ denote the inertial position and velocity, respectively, and let $\boldsymbol{\omega}_\oplus$ denote the terrestrial rotation vector. Under the standard co-rotating-atmosphere assumption, the relative flow velocity is

$$\mathbf{v}_{\text{rel}} = \mathbf{v} - \boldsymbol{\omega}_\oplus \times \mathbf{r} \quad (2.24)$$

Aerodynamic forcing, therefore, depends on the satellite velocity relative to the rotating thermosphere rather than on inertial speed alone [1, 2, 12]. The quadratic drag acceleration is written as

$$\mathbf{a}_{\text{drag}} = -\frac{1}{2} \rho \left(\frac{C_D A_D}{m} \right) \|\mathbf{v}_{\text{rel}}\| \mathbf{v}_{\text{rel}} = -\frac{\rho}{2 BC} \|\mathbf{v}_{\text{rel}}\| \mathbf{v}_{\text{rel}} \quad (2.25)$$

where ρ is the local thermospheric mass density, C_D is the effective drag coefficient, A_D is the effective drag area, m is satellite mass, and $BC = m/(C_D A_D)$ is the ballistic coefficient. This form preserves the standard dependence on atmospheric density, effective aerodynamic response, and the square of the relative speed, while remaining directly compatible with the Cartesian propagation architecture [1, 2, 12, 34, 35].

This representation localizes uncertainty in a transparent way. Environmental variability enters primarily through ρ , while geometry, surface interaction, and unresolved attitude effects are absorbed into $C_D A_D/m$, or equivalently into BC . Drag is therefore treated as a forward model of effective translational aerodynamic response rather than as a resolved gas-surface interaction reconstruction, which is appropriate for long-arc propagation settings where only the aggregate aerodynamic effect is practically observable [1, 2, 38, 39, 12].

2.5.2 Static Reference Atmosphere

The static reference-atmosphere pathway provides an efficient baseline for propagation, benchmarking, and sensitivity analysis. In the present framework, it is implemented as a piecewise scale-height surrogate based on the USSA76 reference profile over the altitude range from 0 to 1000 km. Within each atmospheric layer, density is evaluated through exponential interpolation,

$$\rho(h) = \rho_i \exp\left[-\frac{h - h_i}{H_i}\right], \quad h_i \leq h < h_{i+1} \quad (2.26)$$

where h is altitude above $R_\oplus$, h_i is the lower boundary of the active layer, ρ_i is the reference density at that boundary, and H_i is the corresponding scale height. This construction yields a smooth and computationally inexpensive density surrogate suitable for large ensemble runs and controlled numerical comparisons [80, 1].

Within the simulation architecture, the static atmosphere serves as a methodological benchmark. Because the density field is fixed, it isolates aerodynamic and numerical behavior from transient solar-cycle, geomagnetic, and short-timescale thermospheric variability. It is therefore useful for algorithm development, debugging, and sensitivity baselining, whereas the dynamic atmosphere pathway supplies the environmental fidelity required for operational or space-weather-aware studies. NASA’s U.S. Standard Atmosphere 1976 documentation describes the model as an idealized, steady-state representation of the atmosphere to 1000 km, which is consistent with its use here as a controlled reference background [80].

2.5.3 Dynamic Atmosphere Interface

The dynamic atmosphere pathway employs precomputed Earth-fixed density grids derived from NRLMSIS-family model evaluations, introducing thermospheric variability through temporally resolved, spatially gridded density fields rather than through a closed-form analytic climatology [34, 35, 81]. The active density field is represented on WGS84 geodetic coordinates $(\varphi_g, \lambda_g, h_g)$, with grid dimensions

$$n_{\text{lat}} = 37, \quad n_{\text{lon}} = 72, \quad n_h = 98 \quad (2.27)$$

covering geodetic latitude from -90° to 90° , longitude from 0° to 355° , and geodetic altitude from 115 to 600 km at 5° and 5 km increments respectively. The stored field is total mass density in kg m^{-3} , sampled at four UTC epochs per day. Denoting the corresponding daily density snapshots by $\mathcal{G}_\rho^{(\ell)}$, the snapshot epochs are

$$t_\ell^\rho \in \{00, 06, 12, 18\} \text{ UTC}, \quad \ell = 1, \dots, 4 \quad (2.28)$$

During runtime, the propagated inertial satellite state is transformed into the terrestrial frame and converted to WGS84 geodetic coordinates before querying the active density fields. The density supplied to the drag law is then evaluated as a spatiotemporally interpolated field,

$$\rho_{\text{dyn}} = \rho_{\text{dyn}}(\varphi_g, \lambda_g, h_g, t) \quad (2.29)$$

obtained by spatial interpolation within the active Earth-fixed grids together with temporal interpolation between adjacent UTC snapshots. When propagation crosses a daily boundary, the interpolation stencil includes the first snapshot of the subsequent day so that density remains continuous across the handoff. This architecture preserves synchronization between terrestrial rotation, geodetic position, and the precomputed atmosphere backend [82, 71, 35, 81].

Since the dynamic atmosphere is loaded from discrete daily files or daily grid packages, the propagation interval is segmented at atmosphere-file boundaries whenever a change in the active density set is required. Each segment inherits the terminal state of the previous segment, while the associated atmosphere backend is refreshed for the new date interval. The dynamic pathway, therefore, integrates Earth-fixed density loading, inertial-to-terrestrial state conversion, and time-resolved interpolation within the same propagation framework used by the force model.

2.5.4 Effective Aerodynamic State

The aerodynamic response of the satellite is represented through an effective translational surrogate rather than through panel-resolved free-molecular gas-surface interaction physics. In the present framework, the drag law depends on the parameters C_D , A_D , and m only through the combination

$$\frac{C_D A_D}{m} = \frac{1}{BC} \quad (2.30)$$

where $BC = m/(C_D A_D)$ is the ballistic coefficient. This parameterization preserves the correct drag-force scaling while consolidating unresolved attitude effects, projected-area variability, accommodation uncertainty, and residual density-model discrepancy into a single orbit-effective aerodynamic response [1, 2, 38, 39, 12, 36].

Under the variable-mass low-thrust formulation, the effective aerodynamic state evolves even when C_D and A_D are held fixed, because propellant consumption changes the mass entering the drag law. In that case,

$$BC(t) = \frac{m(t)}{C_D A_D}, \quad \frac{C_D A_D}{m(t)} = \frac{1}{BC(t)} \quad (2.31)$$

The drag model, therefore, remains structurally identical in coast and thrust phases, while the aerodynamic response becomes state dependent through the propagated mass. This treatment is physically cleaner and more interpretable than prescribing an arbitrary time-varying $C_D A_D(t)$ history, because the evolving drag sensitivity is tied directly to the propagated satellite mass state [1, 3].

Within the orbital simulation, aerodynamic uncertainty is addressed by combining the backend density ρ_{dyn} , the local relative flow $\mathbf{v}_{\text{rel}}$, and the effective response parameters (C_D, A_D, m) to close the drag law. The resulting aerodynamic layer is computationally efficient, transparent at the force-model level, and applicable to both controlled reference propagation and space-weather-responsive dynamic simulations over extended arcs [1, 2, 12, 34, 35, 36].

2.6 Phase-Conditioned Low-Thrust Scheduling and Variable-Mass Dynamics

The propulsion model applies schedule-driven, phase-conditioned low thrust that is dynamically coupled to satellite mass, so maneuvering is represented through finite-duration thrust and coast intervals instead of impulsive abstractions. This structure is appropriate for long-horizon LEO simulation because it preserves the temporal organization of orbit raising, shell maintenance,

and disposal lowering while remaining compatible with adaptive numerical propagation, variable atmospheric forcing, and the common Cartesian state-space formulation [3, 1, 83, 84].

2.6.1 Maneuver Phase Taxonomy and Schedule Construction

The framework partitions low-thrust activity into three operational classes: insertion or raising, operational shell maintenance, and disposal lowering. Each schedule segment is assigned a phase label together with the propulsion quantities required by the equations of motion, including duty cycle, thrust magnitude, specific impulse, and thrust-direction sign. For schedule segment κ , the effective continuous-equivalent thrust is defined as

$$T_{\text{eff},\kappa} = d_\kappa \bar{T}_\kappa \quad (2.32)$$

where $d_\kappa \in [0, 1]$ is the duty cycle and $\bar{T}_\kappa$ is the mean thrust magnitude associated with that segment. This construction replaces unresolved high-frequency firing detail with a continuous-equivalent thrust level that is appropriate for long-arc translational propagation [3, 1, 83].

For each satellite, phase intervals are imported as time-stamped records, restricted to the active propagation window, and sorted chronologically into a piecewise control schedule

$$\mathcal{U} = \{[\tau_0, \tau_1), [\tau_1, \tau_2), \dots, [\tau_{N_{\text{sch}}-1}, \tau_{N_{\text{sch}}})\} \quad (2.33)$$

where N_{sch} is the number of schedule segments. Each interval is assigned a maneuver-phase label σ_κ and the corresponding control parameters. Explicit coast segments are inserted wherever no thrusting phase is active, ensuring that the control history covers the full propagation interval without ambiguity. The result is a complete chronological control realization instead of a sparse list of isolated maneuver times.

This schedule layer translates observationally inferred or externally specified phase timing into a propagator-ready control object. The numerical integrator therefore advances on a fully prepared piecewise control history, and the equations of motion read the active segment parameters directly without reconstructing maneuver logic during propagation.

2.6.2 Along-Track Thrust Direction and Mass-Depletion Law

During an active thrust segment, acceleration is applied in the local along-track direction of the instantaneous orbital frame. Let

$$\hat{\mathbf{r}} = \frac{\mathbf{r}}{\|\mathbf{r}\|}, \quad \hat{\mathbf{h}} = \frac{\mathbf{r} \times \mathbf{v}}{\|\mathbf{r} \times \mathbf{v}\|} \quad (2.34)$$

where $\mathbf{r}$ and $\mathbf{v}$ are the inertial satellite position and velocity. The tangential unit vector is then

$$\hat{\mathbf{t}} = \frac{\hat{\mathbf{h}} \times \hat{\mathbf{r}}}{\|\hat{\mathbf{h}} \times \hat{\mathbf{r}}\|} \quad (2.35)$$

This direction aligns the applied thrust with the instantaneous along-track direction and therefore targets orbital-energy modification in a direct way [3, 1]. Orbit raising and shell maintenance use the positive along-track direction, while disposal lowering uses the negative direction. That choice is encoded through a segment-level thrust-direction sign

$$\eta_\kappa \in \{-1, 0, +1\} \quad (2.36)$$

where $\eta_\kappa = 0$ corresponds to a coast segment. The thrust acceleration in schedule segment κ is then

$$\mathbf{a}_{T,\kappa} = \eta_\kappa \frac{T_{\text{eff},\kappa}}{m} \hat{\mathbf{t}} \quad (2.37)$$

This form separates schedule-level decisions about timing, sign, and magnitude from the state-dependent geometry that defines the local thrust direction.

Mass depletion is propagated explicitly whenever thrust is active. Using the standard specific-impulse relation, the mass state obeys

$$\dot{m} = -\frac{T_{\text{eff},\kappa}}{g_0 I_{\text{sp},\kappa}} \quad (2.38)$$

where $I_{\text{sp},\kappa}$ is the specific impulse assigned to schedule segment κ and

$$g_0 = 9.80665 \text{ m s}^{-2} \quad (2.39)$$

This variable-mass treatment is required because both low-thrust acceleration and aerodynamic response depend on the instantaneous satellite mass. As propellant is consumed, the thrust acceleration increases for fixed $T_{\text{eff},\kappa}$, and the drag sensitivity changes through the evolving area-to-mass

ratio, or equivalently through the evolving ballistic coefficient [83, 84, 1]. A dry-mass constraint $m \geq m_{\text{dry}}$ is enforced so that thrusting and propellant depletion terminate once the available propellant budget is exhausted.

The resulting propulsion layer preserves structural consistency across thrust and coast phases. The same propagated state is used throughout, while the schedule determines whether the active segment contributes thrust acceleration and mass depletion or simply reduces to coast dynamics. This design produces a propulsion model that is computationally efficient, physically interpretable, and directly compatible with phase-resolved LEO satellite constellation analysis.

2.6.3 Segment Boundaries and Coast Intervals

The low-thrust simulation advances through an ordered sequence of propagation segments whose boundaries are determined by schedule transitions, atmosphere-backend handoffs, and the global start and end times of the simulation window. Let the complete ordered boundary set be

$$\mathcal{T}_{\text{bnd}} = \left\{ \tau_0^{\text{bnd}}, \tau_1^{\text{bnd}}, \dots, \tau_{N_{\text{bnd}}-1}^{\text{bnd}} \right\} \quad (2.40)$$

where $\tau_0^{\text{bnd}} = t_0$, $\tau_{N_{\text{bnd}}-1}^{\text{bnd}} = t_f$, and the interior elements collect all schedule-change times, explicitly inserted coast-interval boundaries, and atmosphere-update breakpoints induced by the density-grid architecture. These atmosphere transitions occur at fixed UTC boundaries because the active interpolation stencil and file loading logic change when the simulation crosses daily or snapshot-dependent backend limits. The boundary set, therefore, provides a unified temporal partition of maneuver changes, coast intervals, atmosphere handoffs, and global propagation limits.

Segmented propagation is required because the low-thrust control law is piecewise defined, even though the satellite dynamics remain continuous within each open interval. On each segment

$$\mathcal{S}_q = [\tau_q^{\text{bnd}}, \tau_{q+1}^{\text{bnd}}) \quad (2.41)$$

the active control parameters and atmosphere backend are held fixed in the sense required by the local force-model definition, and the terminal state of segment $\mathcal{S}_q$ is inherited as the initial state of segment $\mathcal{S}_{q+1}$. Explicit coast labeling is retained throughout the schedule partition instead of being inferred from missing thrust labels. As a consequence, every subinterval of the propagation window is assigned a well-defined control regime, either thrusting or coasting, and the numerical integrator advances only through fully specified model states.

This boundary logic is directly compatible with adaptive numerical propagation. Continuous-

thrust maneuvers are represented as finite-duration force-model intervals, while start and stop times introduce control discontinuities that are handled through segmentation and state inheritance. The same strategy accommodates atmosphere-backend changes without requiring a reformulation of the governing equations on either side of a boundary [75, 74].

2.6.4 Thrust-Integrated Diagnostics

The schedule-driven propulsion layer produces diagnostics that are physically interpretable for individual cases and comparable across large ensembles. The cumulative delivered impulse is

$$J_{\text{tot}} = \int_{t_0}^{t_f} T_{\text{eff}}(t) dt \quad (2.42)$$

and the total thrust-on duration is

$$\Delta t_{\text{on}} = \int_{t_0}^{t_f} \mathbb{I}_{\text{thrust}}(t) dt \quad (2.43)$$

where $\mathbb{I}_{\text{thrust}}(t)$ is the indicator of an active thrust segment. Because the schedule is explicitly phase-labeled, phase-resolved impulse diagnostics may also be formed as

$$J_{\sigma} = \int_{t_0}^{t_f} \mathbb{I}_{\sigma}(t) T_{\text{eff}}(t) dt \quad (2.44)$$

with $\mathbb{I}_{\sigma}(t)$ selecting the maneuver class σ . These quantities summarize the magnitude, timing, and phase allocation of control actions over the full trajectory without requiring post hoc reconstruction from the propagated state history [83, 84].

When a segment-level thruster-efficiency surrogate is specified, the framework also evaluates an electrical input-power proxy. Using the standard jet-power relation together with the definition of thruster efficiency and the specific-impulse mass-flow law, the segment-level power surrogate is written as

$$P_{\text{in},\kappa} = \frac{T_{\text{eff},\kappa} g_0 I_{\text{sp},\kappa}}{2 \eta_{\text{thr},\kappa}} \quad (2.45)$$

where $\eta_{\text{thr},\kappa}$ is the effective thruster efficiency assigned to schedule segment κ . Time integration then yields the cumulative energy surrogate, represented as

$$E_{\text{in}} = \int_{t_0}^{t_f} P_{\text{in}}(t) dt \quad (2.46)$$

This quantity provides a mission-level integral diagnostic that links thrust scheduling to energetic demand in a compact and interpretable form. A full satellite power-system model would require additional bus, storage, thermal, and power-processing dynamics beyond the present reduced propulsion surrogate [83].

Together with initial mass, final mass, and propellant expenditure, the impulse, thrust-on time, and energy surrogates provide a concise integral summary of how the phase-conditioned control schedule shapes the propagated trajectory over the analysis window.

2.7 Parallel Batch Propagation and Output Products

This section formalizes the extension of the propagation framework from a single initialized case to large ensembles and specifies the trajectory products retained for subsequent analysis. The architecture operates natively in batch mode, with default settings `num_simulations = 32` and `max_parallel_workers = 12`, so ensemble propagation is treated as a primary operating mode rather than as an auxiliary wrapper around the single-satellite solver. Methodological consistency is maintained by preserving a common force model, event structure, schedule logic, and export protocol across all propagated cases [74, 75, 76, 85].

2.7.1 Ensemble Execution and Worker Orchestration

The ensemble layer treats each initialized case in C_0 as an independent propagation task. Once the satellite’s initial state, epoch-conditioned external-body data, atmospheric inputs, and maneuver schedule are assembled, the selected cases are dispatched as a bounded batch of size N_{sim} , with at most N_{worker} propagations executing concurrently. The batch architecture, therefore, separates physical case definition from computational orchestration: each task inherits the same dynamical model and numerical rules, while the worker pool governs only resource allocation and concurrency [86].

This execution model provides reproducible per-case inputs and a uniform output contract across the ensemble, while avoiding implicit changes to the governing equations as the population size grows. The worker pool does not alter the propagated physics; it limits the number of simultaneously active cases and collects completed results into a common batch result set $\mathcal{R}_{\text{batch}}$. In this sense, scaling from one satellite to many satellites is implemented through bounded parallel task dispatch rather than through any reformulation of the propagation problem [86, 85].

A unified retention rule is also imposed at the batch level. In addition to case-specific terminal events, the common cutoff altitude $h_{\text{cut}} = 115$ km defines the lower retention boundary for ensemble

products. This condition excludes deeply decayed trajectories from the nominal batch archive and preserves comparability among retained cases by truncating output before the onset of strongly nonuniform terminal descent behavior.

2.7.2 Stored State and Orbital Element Products

The output architecture is hierarchical. For each propagated case s , the retained state history is stored as

$$\mathcal{Y}_s = \left\{ \left(t_j^{\text{out}}, \mathbf{y}_s(t_j^{\text{out}}) \right) \right\}_{j=0}^{N_s^{\text{out}}-1} \quad (2.47)$$

where t_j^{out} is the stored output grid and N_s^{out} is the number of retained samples for case s . In variable-mass runs, $\mathcal{Y}_s$ includes the propagated mass state together with the Cartesian translational state. Derived orbital-element products are then formed from the retained Cartesian history, denoted as

$$\mathcal{A}_s^{\text{orb}} = \left\{ \left(t_j^{\text{out}}, \boldsymbol{\alpha}_{\text{COE},s}(t_j^{\text{out}}) \right) \right\}_{j=0}^{N_s^{\text{out}}-1} \quad (2.48)$$

thereby combining the numerical robustness of Cartesian propagation with the interpretability of orbital-element diagnostics required for long-arc physical analysis [1, 2].

In addition to retained histories, each case produces a compact diagnostic summary $\mathcal{D}_s$ containing integral and terminal quantities such as final epoch, final mass, terminal altitude status, cumulative impulse, thrust-on duration, and any energy surrogate or event counts evaluated during propagation. The exported case-level product is therefore written as

$$\mathcal{O}_s = \{ \mathcal{Y}_s, \mathcal{A}_s^{\text{orb}}, \mathcal{D}_s \} \quad (2.49)$$

and the ensemble archive is the collected set of such outputs over the completed batch [85, 87].

Trajectory retention is configurable in order to balance archival completeness against storage cost. In the default batch configuration, compact case-level summaries and sampled state products are prioritized over exhaustive retention of all internal solver data, which would be unnecessarily expensive for dense long-arc ensemble runs. Full trajectory retention remains available when detailed inspection of individual event sequences, force-model behavior, or schedule transitions is required. The output layer is therefore designed to preserve the quantities needed for downstream scientific interpretation, calibration, and cross-case comparison without exporting every intermediate numerical artifact generated during propagation [85, 87].

Chapter 3

Data Decomposition and Constellation Clustering

3.1 Constellation Decomposition as a Family-Discovery Problem

This analysis advances from descriptive satellite constellation morphology to a data-driven coarse-graining of the catalog population. The emphasis shifts from individual satellite records to the collective ensemble of public mean-element states $\{\bar{\alpha}_s(t) : s \in \mathcal{I}_{\text{cat}}\}$ evaluated at a common epoch, with $N_{\text{cat}} = |\mathcal{I}_{\text{cat}}|$, interpreted as an engineered and regulated distribution in orbit-geometry space [16, 17, 88, 20, 4]. In this section, a family denotes an inferred cluster-level geometric and operational regime within that snapshot. This usage is distinct from the term division, which remains reserved for the earlier Starlink grouping by inclination, altitude, and plane configuration. The methodology replaces object-by-object inspection with geometry-aware partitioning of the observable population through reduced structural feature vectors $\mathbf{z}_s$ collected in $\mathbf{Z}$, allowing subsequent analysis to address persistent families, transitional populations, and off-nominal states while preserving the catalog as a structured orbital population rather than an undifferentiated cloud [89, 90]. Figure 3.1 illustrates the progression from public orbital data to reduced geometry, density-based family extraction, and operational interpretation.

3.1.1 Large Constellations Require Family-Wise Decomposition

Large satellite constellations manifest as controlled dynamical ensembles in orbital-element space, with observed structure shaped by shell design, plane spacing, launch logistics, secular perturba-

tions, drag-induced dispersion, and station-keeping maneuvers [13, 14, 1, 2, 4, 5]. A single-epoch catalog snapshot superposes nominal service states, commissioning phases, transfer corridors, disposal pathways, and sporadic anomalies. Under these conditions, object-by-object inspection does not resolve the mesoscopic structure that governs shell occupancy, plane organization, maintenance burden, and conjunction exposure, and it does not scale to the complexity present in the population [4, 5].

The limitations of the public catalog intensify this mismatch. In the absence of operator truth states, maneuver logs, and covariance-resolved trajectories, the catalog supplies model-dependent mean-element products whose interpretation is constrained by the underlying observation and propagation assumptions [16, 17, 88, 20]. Scientific analysis, therefore, benefits from focusing on persistent structure within the observable population rather than attempting to reconstruct every individual satellite history from incomplete public products. Family-wise decomposition serves that purpose by mapping the snapshot into a reduced structural representation $\mathbf{Z}$, extracting recurrent families $f \in \mathcal{F}$, and preserving transitional or off-nominal members that do not belong to any persistent geometric regime.

This dimensionality reduction is enabled by timescale separation. Along-track phase evolves rapidly and often acts as a nuisance variable for structural classification, whereas shell altitude, inclination band, orbital-plane orientation, and shape descriptors evolve slowly and are directly regulated [1, 2, 3]. An effective decomposition, therefore, retains variables that govern persistent geometry while suppressing or marginalizing variables that primarily introduce ambiguity. The resulting map from $\mathbf{Z}$ to extracted families $\mathcal{F}$, labels f_s , and any transitional remainder $\mathcal{I}_{\text{tr}}$ preserves operationally significant transients and enables streamlined interpretation.

3.1.2 Satellite Families

A satellite family is defined here as a persistent and operationally meaningful region in reduced orbit-geometry space. Let $\mathbf{z}_s$ denote the reduced structural feature vector for satellite s , let $d_{\mathcal{Z}}$ denote the similarity metric imposed on that representation, and let $\mathcal{G}_{\text{nn}}$ denote the induced neighborhood graph used to evaluate local adjacency and density connectivity. A family $f \in \mathcal{F}$ is then a recurrent component of the common-epoch catalog population whose members

$$\mathcal{I}_f = \{ s \in \mathcal{I}_{\text{cat}} : f_s = f \}$$

remain coherent under the specified representation, metric, and density criteria [89, 90]. Operational meaning is assigned through downstream use. In the present thesis, that meaning includes shell-level

interpretation, plane-level monitoring, deployment assessment, anomaly triage, and family-level physical inference. Satellites belong to the same family when their states occupy a common geometric regime whose coherence is sufficiently strong to support one or more of those tasks. The analysis, therefore, treats a family as an empirically recovered controlled orbital regime, with evidentiary status conditioned by the feature representation, neighborhood graph, density model, and validation target.

This framework interprets the satellite constellation as a controlled and perturbed dynamical system observed through a lossy public measurement channel. Design establishes a lattice over slowly varying orbital coordinates such as altitude, inclination, and nodal orientation. Environmental forcing, secular drift, and vehicle-dependent drag deform that lattice over time, while control actions act to preserve service geometry within operational tolerances [1, 2, 3]. The public catalog records the resulting population through mean-element products of limited fidelity [16, 17, 88, 20]. Within this observation model, families correspond to high-density basins in the reduced geometry of slow variables, whereas low-density bridges, sparse filaments, and isolated members indicate transfer states, weakly occupied transition corridors, or anomalous behavior [91, 92, 93]. Satellites that do not participate in a persistent regime may be retained in the transitional set $\mathcal{I}_{\text{tr}}$ or, when the density-based extraction explicitly rejects them, in the outlier set $\mathcal{I}_{\text{noise}}$.

This dynamical interpretation constrains the role of latent space. The embedding is used as a chart for neighborhood structure induced by the observed orbit geometry, with physical interpretation assigned through preserved adjacency, cluster persistence, and orbital diagnostics [94, 95]. The primary structures of interest are therefore the adjacency relations encoded by $\mathcal{G}_{\text{nn}}$, the associated density stratification, and the persistent connected components that survive under the chosen clustering rule. A family is characterized by persistent local adjacency and density connectivity in the induced geometry, conditioned on a representation that preserves the orbital invariants most relevant for structural classification [91, 92, 93].

The family concept must satisfy both empirical and constructive requirements. Empirically, families serve as public-data proxies for distinct dynamical and operational regimes, including shells, planes, and stable sub-populations within planes. Constructively, families provide a reduced taxonomy that makes large satellite constellations tractable for monitoring, diagnosis, and downstream inference. This thesis adopts both roles while keeping them conceptually distinct. The recovered partition is presented as a task-appropriate scientific taxonomy induced by the observable data and the chosen similarity geometry [89, 90].

3.1.3 Validation Logic for Labeled, Semi-Labeled, and Unlabeled Cases

Validation cannot be anchored to a single truth criterion across all benchmark satellite constellations considered in this work. The evidentiary standard must follow the strength and granularity of the available prior information. The methodology, therefore, adopts a three-tier validation logic [90].

The first tier is a labeled benchmark for which GPS functions as the strongest positive control because the plane-level architecture is publicly documented with high specificity [96]. The decomposition is expected to recover the canonical six-plane organization while accommodating practical departures such as slot expansion, unequal slot occupancy, and surplus operational vehicles. Validation at this tier is explicitly external. Published architecture supplies the reference against which the structural fidelity of family recovery can be assessed.

The second tier is a semi-labeled architectural benchmark for which OneWeb supplies public shell-level and plane-level priors, yet those priors are less definitive than the GPS benchmark because the operational architecture has evolved and the cataloged substructure is not fully resolved by public documentation alone [97]. Validation in this setting is conditioned on architectural concordance rather than exact family labels. The decomposition should recover plane-consistent organization, support interpretable substructure within that organization, and remain stable under the perturbations introduced by deployment chronology, drag modulation, and residual within-plane heterogeneity.

The third tier is an unlabeled discovery setting using Starlink, as public disclosures provide shell-level priors and broad architectural constraints, but do not provide authoritative family-level annotation for the internal substructure targeted in this thesis [61, 4]. Validation, therefore, emphasizes structural coherence, robustness, and physical interpretability. Effective decomposition should align with shell-scale organization, resolve persistent within-shell structure, and separate recurrent operational populations from transitional and marginal states. The absence of family labels is part of the scientific problem itself, so the taxonomy must be inferred from the observation layer rather than read from an external label source.




Figure 3.1: Conceptual Workflow for Constellation Decomposition as Satellite Family Discovery

This three-tier framework assigns a distinct methodological role to each benchmark satellite constellation. GPS tests direct recovery of a well-specified public architecture. OneWeb tests robustness under partial prior specification and evolving large satellite constellation geometry. Starlink is the principal discovery case, where shell-level priors constrain the admissible interpretation space but do not determine the family taxonomy in advance [96, 97, 61, 4]. The validation sequence, therefore, proceeds from strong external architectural reference to partial architectural concordance to unlabeled inference within physically delimited constraints.

3.2 Public Data and Orbit-Geometry Representation

Family-discovery methodology begins with public GP catalog products rather than operator truth states, covariance-resolved trajectories, or maneuver histories. The observable input is therefore compressed and model-conditioned. Its content depends on catalog-generation assumptions, reference-epoch choices, and the semantics of the associated propagation model [16, 17, 88, 20]. The present framework accepts those constraints explicitly, synchronizes the catalog to a common analysis epoch t^* , reconstructs a standardized orbit-geometry representation $\mathbf{g}_s$ for each retained

satellite, and defines the similarity geometry used later for family discovery. The objective of this stage is structural inference in public-data space. Full-fidelity orbit determination lies outside its scope.

3.2.1 TLE-Derived Public Data and Epoch Synchronization

The workflow processes public GP element sets, implemented here as TLE records drawn from the catalog archive [16, 17, 88, 20]. The raw archive is defined as

$$\mathcal{D}_{\text{raw}} = \{\text{TLE}_j(t_j)\}_{j=1}^{N_{\text{raw}}} \quad (3.1)$$

where each record is associated with its catalog epoch t_j . The synchronized working set is then defined by temporal filtering around the target epoch $t^\star$,

$$\mathcal{D}(t^\star, \Delta t) = \{\text{TLE}_j(t_j) \in \mathcal{D}_{\text{raw}} \mid |t_j - t^\star| \leq \Delta t\} \quad (3.2)$$

From this filtered set, one representative record is retained for each satellite in the common-epoch snapshot, thereby inducing the index set $\mathcal{I}_{\text{cat}}$ and the retained population size $N_{\text{cat}} = |\mathcal{I}_{\text{cat}}|$. For the present analysis, $t^\star = 2025-04-10$ and $\Delta t = 0.5$ days. This constitutes a single-epoch structural analysis with tolerance-based synchronization rather than a multi-epoch dynamical clustering framework.

Synchronization is necessary because large satellite constellations are cataloged at non-uniform epochs. Without explicit epoch control, geometric differences caused by shell structure, plane separation, and deployment state are superposed with artifacts caused solely by timestamp offsets. The synchronization window reduces this temporal inconsistency, although it does not remove dependence on the public propagation model. A TLE remains a model-specific mean-element product rather than a covariance-resolved Cartesian truth state [16, 17, 88]. Any reconstructed state derived from synchronized TLEs must therefore be interpreted as a diagnostic representation of the public catalog record.

A deterministic preprocessing map Φ_{orb} converts each synchronized record into a standardized orbit-geometry representation,

$$\mathbf{g}_s = \Phi_{\text{orb}}(\text{TLE}_s(t_j)), \quad s \in \mathcal{I}_{\text{cat}} \quad (3.3)$$

and the resulting dataset is collected row-wise as

$$\mathbf{G} = \begin{bmatrix} \mathbf{g}_1^\top \\ \mathbf{g}_2^\top \\ \vdots \\ \mathbf{g}_{N_{\text{cat}}}^\top \end{bmatrix} \quad (3.4)$$

In implementation, the catalog quantities n , e , i , Ω , ω , and M are transformed into a common inertial geometric proxy by recovering a Keplerian semi-major-axis proxy from mean motion, solving Kepler’s equation, and rotating perifocal position and velocity vectors into the chosen inertial frame [1, 98, 99]. The resulting $(\mathbf{r}_s, \mathbf{v}_s)$ states and their derived quantities provide a common ambient geometry for downstream representation learning. Their provenance remains that of synchronized public mean-element products, so the reconstructed states are used here as Keplerized diagnostic embeddings of the catalog record rather than as inputs to precision orbit determination [16, 17].

3.2.2 Limitations of Raw Classical Elements in Dense Constellations

Raw classical elements retain strong descriptive value, but their utility depends on population density and on the scale of the analysis [1, 2, 98, 99]. At coarse resolution, shell stratification is often visible directly in semi-major axis and inclination, producing distinct altitude and inclination bands. These bands are informative for shell-level interpretation and remain useful diagnostic coordinates throughout the analysis. Dense near-circular satellite constellations, however, exhibit substantial projection overlap after global mixing of shells and planes, so pairwise plots of $(a, e, i, \Omega, \omega, \nu)$ do not reliably expose fine family structure.

Three mechanisms drive this limitation. First, angular periodicity complicates similarity calculations. Coordinates such as Ω , ω , and ν wrap on 2π , so naive Euclidean differences can distort genuine angular proximity [1, 98]. Second, near-circular regimes degrade the interpretability of periapsis-based coordinates. As $e \rightarrow 0$, the periapsis direction becomes poorly conditioned, and the descriptive content of ω and ν correspondingly weakens [1, 2, 98]. Third, dense LEO shells produce geometric crowding. Large numbers of satellites share nearly identical a , e , and i while differing mainly through plane orientation, phasing history, deployment chronology, or small control-box dispersions. These effects compress distinct operational sub-populations into overlapping projections.

The family-discovery pipeline, therefore, does not cluster directly in raw classical-element

space. It instead operates on the standardized vector-valued orbit-geometry representation $\mathbf{g}_s$, which is designed to preserve slow structural variables, regularize periodic coordinates, and reduce singular sensitivity in near-circular regimes. Classical elements remain indispensable for post hoc physical interpretation and for verifying that recovered families are consistent with orbital mechanics, especially through explanatory combinations such as (a, Ω) and (i, Ω) . Their role in this chapter is consequently diagnostic and interpretive, while the primary extraction of families proceeds in the derived orbit-geometry space [1, 2, 98, 99].

3.2.3 Reduced Orbital Parameter Vector

This section employs five orbit-geometry vectors evaluated at the common epoch $t^\star$. The inertial position and velocity vectors are written as

$$\mathbf{r} = \begin{bmatrix} r_x & r_y & r_z \end{bmatrix}^\top, \quad \mathbf{v} = \begin{bmatrix} v_x & v_y & v_z \end{bmatrix}^\top \quad (3.5)$$

The specific angular-momentum vector is derived from these quantities as

$$\mathbf{h} = \mathbf{r} \times \mathbf{v} \quad (3.6)$$

The eccentricity vector is given by

$$\mathbf{e} = \frac{1}{\mu_\oplus} (\mathbf{v} \times \mathbf{h}) - \hat{\mathbf{r}}, \quad \hat{\mathbf{r}} = \frac{\mathbf{r}}{\|\mathbf{r}\|} \quad (3.7)$$

The nodal vector is defined with respect to the inertial z -axis unit vector $\hat{\mathbf{k}}$ as

$$\mathbf{n} = \hat{\mathbf{k}} \times \mathbf{h} \quad (3.8)$$

At this stage, the geometric interpretation is made explicit. Since $\mathbf{n}$ is obtained from the cross product of $\hat{\mathbf{k}}$ and $\mathbf{h}$, its magnitude is

$$\|\mathbf{n}\| = \|\mathbf{h}\| \sin i \quad (3.9)$$

Therefore, $\mathbf{n}$ vanishes for equatorial orbits, while its direction specifies the ascending node within the equatorial plane. The vector $\mathbf{h}$ defines the normal to the orbital plane, and $\mathbf{e}$ encodes both periapsis direction and scalar eccentricity [1, 2, 98, 99].

These vectors are concatenated to form the Reduced Orbital Parameter Vector (ROPV) for

satellite s ,

$$\mathbf{x}_{\text{ROPV},s} = \begin{bmatrix} \mathbf{r}_s^\top & \mathbf{v}_s^\top & \mathbf{e}_s^\top & \mathbf{h}_s^\top & \mathbf{n}_s^\top \end{bmatrix}^\top \in \mathbb{R}^{15} \quad (3.10)$$

Across the synchronized population, this induces the feature map

$$\Phi_{\text{ROPV}}: \mathbb{R}^6 \rightarrow \mathbb{R}^{15}, \quad \boldsymbol{\alpha}_{\text{COE}} = \begin{bmatrix} a & e & i & \Omega & \omega & \nu \end{bmatrix}^\top \mapsto \mathbf{x}_{\text{ROPV}} \quad (3.11)$$

and the stacked population matrix

$$\mathbf{X}_{\text{ROPV}} = \begin{bmatrix} \mathbf{x}_{\text{ROPV},1}^\top \\ \mathbf{x}_{\text{ROPV},2}^\top \\ \vdots \\ \mathbf{x}_{\text{ROPV},N_{\text{cat}}}^\top \end{bmatrix} \quad (3.12)$$




Figure 3.2: Reduced Orbital Parameter Vector Structure for the Global Starlink Population

Figure 3.2 shows the induced vector geometry of the ROPV for the synchronized Starlink population, including the radius, velocity, angular-momentum, eccentricity, nodal, and perifocal minor-axis vector projections. These projections make visible the differing geometric roles of each vector block before standardization and embedding. The map is not intended to be state-minimal, since a six-parameter orbital chart is sufficient for Keplerian state determination. Its purpose is instead to express orbit-plane orientation, energy scale, and shape through explicit Euclidean coordinates that are more compatible with neighborhood-based unsupervised learning [1, 2, 98, 99, 89]. In this setting, “reduced” refers to a reduction of coordinate pathologies and operational representational burden for downstream analysis, not to a reduction in ambient dimension.

3.2.4 Local Metric Geometry

Given the feature map Φ_{ROPV} , the central issue is the similarity structure induced in $\mathbb{R}^{15}$. The Jacobian

$$\mathbf{J}_{\text{ROPV}}(\boldsymbol{\alpha}_{\text{COE}}) = \frac{\partial \Phi_{\text{ROPV}}}{\partial \boldsymbol{\alpha}_{\text{COE}}} \in \mathbb{R}^{15 \times 6} \quad (3.13)$$

characterizes the local mapping from classical-element space to feature space. In combination with a feature-space weighting matrix $\mathbf{M}_{\text{ROPV}}$, it induces the local pullback metric

$$\boldsymbol{\Gamma}_{\text{ROPV}}(\boldsymbol{\alpha}_{\text{COE}}) = \mathbf{J}_{\text{ROPV}}(\boldsymbol{\alpha}_{\text{COE}})^\top \mathbf{M}_{\text{ROPV}} \mathbf{J}_{\text{ROPV}}(\boldsymbol{\alpha}_{\text{COE}}) \quad (3.14)$$

For a small perturbation $d\boldsymbol{\alpha}_{\text{COE}}$, the squared local distance between nearby orbits is approximated by

$$\Delta^2(\boldsymbol{\alpha}_{\text{COE}}, \boldsymbol{\alpha}_{\text{COE}} + d\boldsymbol{\alpha}_{\text{COE}}) \approx d\boldsymbol{\alpha}_{\text{COE}}^\top \boldsymbol{\Gamma}_{\text{ROPV}}(\boldsymbol{\alpha}_{\text{COE}}) d\boldsymbol{\alpha}_{\text{COE}} \quad (3.15)$$

This construction clarifies the representational role of the ROPV. It defines a local similarity geometry designed so that Euclidean neighborhoods in feature space better align with shared orbital structure. The resulting metric is therefore a task-driven geometric design for manifold learning and density-based clustering, rather than a claim that an orbit possesses fifteen independent dynamical degrees of freedom [100, 101].

This distinction is significant because the induced geometry is highly sensitive to block scaling. The practical workflow, therefore, standardizes the ROPV coordinates componentwise before neighborhood construction,

$$\tilde{\mathbf{x}}_{\text{ROPV},s} = \mathbf{D}_{\text{ROPV}}^{-1/2} (\mathbf{x}_{\text{ROPV},s} - \hat{\boldsymbol{\mu}}_{\text{ROPV}}) \quad (3.16)$$

where $\hat{\boldsymbol{\mu}}_{\text{ROPV}}$ is the empirical sample-mean vector and $\mathbf{D}_{\text{ROPV}}$ is the diagonal empirical variance matrix. This transformation yields a standardized feature space with zero-mean, unit-variance coordinates across heterogeneous vector blocks. Full whitening would additionally remove cross-coordinate correlations and is therefore not implied by diagonal variance scaling alone [101, 102]. The standardized vectors $\tilde{\mathbf{x}}_{\text{ROPV},s}$ constitute the immediate Euclidean representation from which later reduced structural features are derived.

3.3 Neighborhood-Preserving Embedding and Density-Based Family Extraction

The orbit-geometry representation developed above imposes a similarity structure on the synchronized catalog snapshot, but it does not by itself produce an interpretable family decomposition. Extracting such a structure requires a lower-dimensional chart in which neighborhood relations inherited from the standardized ROPV geometry are preserved sufficiently well for density-connected formations to be identified. The present framework, therefore, adopts a sequential pipeline. UMAP constructs a neighborhood-preserving latent chart from the standardized ROPV input, and HDBSCAN then operates on that chart to extract persistent density structure without imposing rigid parametric assumptions on family shape [94, 95, 103, 92, 93]. Within this pipeline, the embedding stage is treated as a geometric approximation problem whose fidelity must be evaluated before density extraction is attempted. The selected latent coordinates are retained together with the synchronized catalog indices $\mathcal{I}_{\text{cat}}$, so that downstream family labels f_s and any noise assignments $\mathcal{I}_{\text{noise}}$ can be remapped directly to the original orbital records and interpreted in the full orbit-geometry space.

3.3.1 Neighborhood-Preserving Latent Charting

Let the standardized input matrix be

$$\tilde{\mathbf{X}}_{\text{ROPV}} = \begin{bmatrix} \tilde{\mathbf{x}}_{\text{ROPV},1}^\top \\ \tilde{\mathbf{x}}_{\text{ROPV},2}^\top \\ \vdots \\ \tilde{\mathbf{x}}_{\text{ROPV},N_{\text{cat}}}^\top \end{bmatrix} \in \mathbb{R}^{N_{\text{cat}} \times 15} \quad (3.17)$$

where each row is the componentwise standardized ROPV feature vector for one satellite in the synchronized catalog snapshot. UMAP does not seek global isometry in this space. It instead constructs a weighted neighborhood graph from local distances and then optimizes a low-dimensional chart whose fuzzy adjacency relations approximate those of the input graph [94, 95]. The resulting coordinates therefore encode neighborhood structure induced by the orbit-geometry representation, rather than direct physical orbital variables.

For each satellite $s \in \mathcal{I}_{\text{cat}}$, let $\mathcal{N}_k^{\text{in}}(s)$ denote the k -nearest-neighbor set in the standardized input space under the chosen metric, and let $d_{ss'}^{\text{in}}$ denote the corresponding pairwise distance to neighbor s' . UMAP assigns a local connectivity offset ρ_s^{UMAP} and a local smooth scale σ_s^{UMAP} , yielding the directed fuzzy neighborhood memberships

$$p_{s'|s}^{\text{UMAP}} = \exp\left(-\frac{\max(0, d_{ss'}^{\text{in}} - \rho_s^{\text{UMAP}})}{\sigma_s^{\text{UMAP}}}\right), \quad \sum_{s' \in \mathcal{N}_k^{\text{in}}(s) \setminus \{s\}} p_{s'|s}^{\text{UMAP}} \approx \log_2 k \quad (3.18)$$

These directed memberships are symmetrized through the probabilistic t -conorm,

$$w_{ss'}^{\text{UMAP}} = p_{s|s'}^{\text{UMAP}} + p_{s'|s}^{\text{UMAP}} - p_{s|s'}^{\text{UMAP}} p_{s'|s}^{\text{UMAP}}, \quad w_{ss'}^{\text{UMAP}} \in [0, 1] \quad (3.19)$$

which defines the weighted undirected neighborhood graph used by the embedding stage [94, 95].

The low-dimensional latent chart is represented by coordinates ($\mathbf{y}_s \in \mathbb{R}^{d_{\text{emb}}}$, $d_{\text{emb}} = 3$), stacked row-wise as

$$\mathbf{Y}_{\text{UMAP}} = \begin{bmatrix} \mathbf{y}_1^\top \\ \mathbf{y}_2^\top \\ \vdots \\ \mathbf{y}_{N_{\text{cat}}}^\top \end{bmatrix} \quad (3.20)$$

UMAP equips this latent chart with the differentiable similarity kernel

$$q_{ss'}^{\text{UMAP}} = \left(1 + a_{\text{UMAP}} \|\mathbf{y}_s - \mathbf{y}_{s'}\|_2^{2b_{\text{UMAP}}}\right)^{-1} \quad (3.21)$$

and determines the coordinates by minimizing the fuzzy-set cross-entropy objective

$$\mathcal{L}_{\text{UMAP}} = \sum_{s \neq s'} \left[w_{ss'}^{\text{UMAP}} \log \frac{w_{ss'}^{\text{UMAP}}}{q_{ss'}^{\text{UMAP}}} + \left(1 - w_{ss'}^{\text{UMAP}}\right) \log \frac{1 - w_{ss'}^{\text{UMAP}}}{1 - q_{ss'}^{\text{UMAP}}} \right] \quad (3.22)$$

so that low-dimensional affinities reproduce the input fuzzy graph as closely as possible [94, 95].

Two user-controlled hyperparameters dominate the geometric behavior of this stage: `n_neighbors`

and `min_dist`. The first fixes the neighborhood scale used to construct the input graph. Small values emphasize highly local structure and can fragment the geometry, whereas larger values incorporate broader manifold context and suppress fine-scale fragmentation. The second controls how tightly the embedding is permitted to pack points together. Lower values encourage compact local condensations, which can be advantageous for density-based clustering, while larger values preserve broader topological spread and resist artificial clumping [94, 95].




Figure 3.3: UMAP-HDBSCAN Grid Search Space

Figure 3.3 summarizes the coupled hyperparameter search used to select the embedding and clustering configuration. The left column records trustworthiness over the UMAP grid, the middle column records DBCV over the HDBSCAN grid, and the right column visualizes the combined latent-density response induced by candidate parameter settings. Interpretation of the latent chart must remain disciplined. Neither the axes nor the Euclidean distances in $\mathbf{Y}_{\text{UMAP}}$ carry direct orbital-mechanical meaning. The scientifically relevant objects are the preserved neighborhood relations, their density stratification, and the stability of these structures under parameter perturbation and repeated optimization. The latent chart is therefore used as a computational atlas for orbit-similarity geometry rather than as a reduced orbit state [94, 95]. In implementation, the retained embedding $\mathbf{Y}_{\text{UMAP}}$ is exported together with the corresponding synchronized catalog ordering, ensuring that all later family assignments remain traceable to the original orbit-geometry records.

3.3.2 Trustworthiness

Because UMAP depends on both hyperparameter choice and stochastic optimization, the embedding stage requires an explicit fidelity criterion before clustering is performed. The present framework uses trustworthiness as the primary selection metric because it directly evaluates preservation of local neighborhood structure, which is the geometric property most critical for subsequent density-based family extraction [104, 102]. Trustworthiness penalizes points that appear among the nearest neighbors in the embedding, even though these were not comparably near in the original standardized feature space. The penalty increases with the rank discrepancy in the input geometry.

Let $\mathcal{N}_k^{\text{emb}}(s)$ denote the k -nearest neighbors of satellite s in the embedding, and let $r_{\text{in}}(s, s')$ denote the rank of satellite s' relative to s in the standardized input geometry. The trustworthiness score at neighborhood scale k is

$$T_{\text{trust}}(k) = 1 - \frac{2}{N_{\text{cat}}k(2N_{\text{cat}} - 3k - 1)} \sum_{s \in \mathcal{I}_{\text{cat}}} \sum_{s' \in \mathcal{N}_k^{\text{emb}}(s) \setminus \mathcal{N}_k^{\text{in}}(s)} (r_{\text{in}}(s, s') - k) \quad (3.23)$$

for $0 \leq T_{\text{trust}}(k) \leq 1$ with the standard admissibility condition $k < N_{\text{cat}}/2$. Values approaching unity indicate that the embedding introduces few spurious low-dimensional neighbors. Lower values indicate stronger distortion of local adjacency and therefore weaker suitability for density-based family extraction [104, 102].




Figure 3.4: Trustworthiness Evaluation Workflow

Figure 3.4 formalizes the trustworthiness calculation as a rank-based comparison between neighborhoods in the standardized ROPV input space and neighborhoods in the UMAP embedding. This workflow makes explicit that embedding selection is governed by local-neighborhood fidelity before density extraction is applied. In this framework, trustworthiness is interpreted strictly as a local-rank fidelity diagnostic. It is not a measure of global geometric faithfulness, nor does it certify physical interpretability of the embedding axes. The relevance arises from the fact that HDBSCAN acts on local density organization in the latent chart. The workflow therefore evaluates a grid of $n_neighbors$ and min_dist settings, computes $T_{trust}(k)$ for each candidate embedding at the neighborhood scale used for graph construction, and retains only the highest-fidelity chart for density-based family extraction. The selected $\mathbf{Y}_{UMAP}$ is then passed forward to the clustering stage, and all resulting family labels f_s and noise assignments $\mathcal{I}_{noise}$ are remapped to the synchronized catalog snapshot through the preserved satellite index ordering.

3.3.3 Hierarchical Density Extraction

Following the construction of the neighborhood-preserving latent chart, the next stage isolates the persistent density-connected structure. HDBSCAN is well-suited to this setting because it does not impose convexity, equal-variance assumptions, or a fixed global density threshold [91, 103, 92, 93]. This is important for satellite constellation analysis, where recovered families may appear as compact components, elongated manifolds, shell-conditioned arcs, or fragmented operational corridors. The hierarchical density model is therefore more compatible with the geometry of orbit-similarity structure than centroid-based partitions.

Let $\mathbf{y}_s \in \mathbb{R}^{d_{\text{emb}}}$ denote the retained UMAP coordinate of satellite s , and let

$$k_{\text{core}} = \text{min_samples} \quad (3.24)$$

denote the HDBSCAN neighborhood parameter used to define local density. The core distance of satellite s is the distance from $\mathbf{y}_s$ to its k_{core} -th nearest neighbor in the retained embedding,

$$d_{\text{core}}(s) = d(\mathbf{y}_s, \mathbf{y}_{(k_{\text{core}})}(s)) \quad (3.25)$$

where $d(\cdot, \cdot)$ is the Euclidean metric in the retained latent chart. For any pair of satellites $s, s' \in \mathcal{I}_{\text{cat}}$, the mutual-reachability distance is then

$$d_{\text{mreach}}(s, s') = \max\{d_{\text{core}}(s), d_{\text{core}}(s'), \|\mathbf{y}_s - \mathbf{y}_{s'}\|_2\} \quad (3.26)$$

This transformed distance expands sparse regions and suppresses short bridging edges between weakly related components, thereby making the subsequent hierarchical extraction more robust to noise and tenuous connectivity.

HDBSCAN conceptually applies single-linkage clustering in the mutual-reachability geometry. In practice, this is achieved by constructing the minimum spanning tree $\mathcal{T}_{\text{mst}}$ of the mutual-reachability graph, converting that tree into a single-linkage hierarchy $\mathcal{H}_{\text{SL}}$, and then condensing the hierarchy under the minimum-cluster-size constraint `min_cluster_size` [103, 92, 93]. The condensed tree $\mathcal{T}_{\text{cond}}$ retains only those branchings for which both descendants satisfy the minimum-size requirement. When one side of a split falls below that threshold, the interpretation is that points fall out of a persistent cluster rather than that a new cluster has formed.

The condensed representation is most naturally expressed on the inverse-distance density scale

$$\lambda = \frac{1}{\varepsilon} \quad (3.27)$$

where ε is the mutual-reachability distance threshold. For a candidate cluster C in the condensed tree, let $\lambda_{\text{birth}}(C)$ denote the density level at which the branch is born and $\lambda_{\text{death}}(C)$ the level at which it terminates through a valid split. For each satellite $s \in C$, let $\lambda_s(C)$ denote the density level at which that point leaves the branch. The cluster stability is then

$$\text{Stab}(C) = \sum_{s \in C} (\lambda_s(C) - \lambda_{\text{birth}}(C)) \quad (3.28)$$

This quantity is the discrete excess-of-mass stability functional used to select persistent branches from the condensed hierarchy. Branches with large stability persist across broad density ranges and support a larger accumulated point mass. Under the standard excess-of-mass selection rule, the retained flat clustering is the non-overlapping set of branches that maximizes total stability over the tree.

Persistent branches correspond to stable structural components in the local orbit-similarity geometry under density refinement. Operational families are expected to exhibit such persistence, whereas transfer states, disposal trajectories, or anomalous satellites typically appear as weak branches or are excluded from the noise set $\mathcal{I}_{\text{noise}}$. The resulting HDBSCAN partition is therefore treated as a data-driven candidate family decomposition induced by density persistence in the latent chart.

3.3.4 Density-Based Clustering Validation

The second-stage model-selection criterion must be compatible with a density-connected cluster structure. Classical internal indices are designed primarily for centroidal or approximately globular partitions and are therefore poorly matched to the irregular cluster families recovered by HDBSCAN [105, 106, 107]. This framework accordingly employs the Density-Based Clustering Validation (DBCv) index, a relative validity measure developed specifically for density-based clustering [107]. DBCv evaluates whether each recovered cluster is internally supported by high density and externally separated by density valleys in the same density-connected geometry used by the clustering analysis.




Figure 3.5: Density-Based Clustering Validation Workflow

Figure 3.5 summarizes the validation logic used after HDBSCAN extraction, including label preprocessing, valid-cluster checks, mutual-reachability graph construction, minimum-spanning-tree evaluation, cluster-level density sparseness, inter-cluster density separation, and coverage-adjusted global scoring. Let the retained flat clustering be

$$\mathcal{P}_{\text{HDB}} = \{C_1, \dots, C_L\} \quad (3.29)$$

with any non-assigned objects collected in $\mathcal{I}_{\text{noise}}$. For a cluster C_i , the density sparseness $\text{DSC}(C_i)$ is defined as the largest internal edge weight in the density-connected tree associated with that cluster. For a pair of clusters C_i and C_j , the density separation $\text{DSPC}(C_i, C_j)$ is the smallest

density-connected linkage joining the two clusters. The cluster-level validity index is then

$$V_{\text{DBCV}}(C_i) = \frac{\min_{j \neq i} \text{DSPC}(C_i, C_j) - \text{DSC}(C_i)}{\max\{\min_{j \neq i} \text{DSPC}(C_i, C_j), \text{DSC}(C_i)\}}, \quad V_{\text{DBCV}}(C_i) \in [-1, 1] \quad (3.30)$$

A large positive value indicates that cluster C_i is internally denser than the density valleys that separate it from competing clusters. Negative values indicate that the proposed cluster is not strongly supported by the ambient density structure.

The global DBCV score is taken as the size-weighted average of the cluster-level validities over the full synchronized population,

$$\text{DBCV}(\mathcal{P}_{\text{HDB}}) = \sum_{i=1}^L \frac{|C_i|}{N_{\text{cat}}} V_{\text{DBCV}}(C_i) \quad (3.31)$$

This normalization is useful in the present application because it penalizes partitions that discard large fractions of the synchronized population as noise. A clustering that achieves strong within-cluster density structure only by labeling substantial portions of the satellite constellation as unassigned will therefore receive a reduced global score. High positive DBCV values indicate that the recovered partition is supported by the density organization of the embedded orbit-geometry representation, whereas values near zero or below zero indicate weak or unstable family structure.

Within the present workflow, DBCV serves as the second-stage selection criterion after trustworthiness has been used to select the latent chart. The embedding stage first ensures that local neighborhood structure is preserved well enough for density analysis to be meaningful. The clustering stage then evaluates candidate HDBSCAN settings through $\text{DBCV}(\mathcal{P}_{\text{HDB}})$, retaining the partition that best balances internal density coherence, external density separation, and controlled treatment of noise. The selected clustering is finally remapped to the synchronized catalog ordering, yielding family labels f_s and the associated family member sets $\mathcal{I}_f$ for downstream physical interpretation.

3.4 Positive-Control Recovery of Known Architecture

The benchmark phase evaluates whether the family-discovery pipeline recovers structural organization imposed by externally documented satellite constellation architectures. The criterion, therefore, extends beyond visual plausibility in the latent chart and instead requires agreement between the recovered density hierarchy, the low-dimensional projection, and the classical-element diagnostics on one hand, and the publicly documented orbital architecture on the other [94, 95,

92, 93]. In this chapter, GPS and OneWeb serve as positive controls with different evidentiary strengths. GPS provides the strongest external benchmark because its core architecture is publicly documented at the plane-and-slot level and remains geometrically regular [108, 109, 96]. OneWeb provides a more demanding architectural control because its public configuration retains a clear shell and plane organization while exhibiting greater density, larger population size, and more complex embedded geometry [97]. In the present benchmark phase, the OneWeb reference is therefore the public shell-and-plane architecture, with satellite-by-satellite external labels unavailable at the required granularity.

3.4.1 GPS as an External Benchmark

GPS is the strongest positive control in this chapter because its public architecture is both explicit and geometrically regular [108, 109, 96]. At the clustering scale relevant here, the principal objective is recovery of the plane structure. The selected pipeline achieves that objective: the condensed tree exhibits six persistent branches across the density scale λ , the UMAP embedding separates into six compact and well-isolated groups, and the classical-element diagnostics show narrow RAAN bands at nearly constant semi-major axis and inclination. The condensed-tree structure in Figure 3.6 provides the first validation diagnostic by showing that six branches persist under density refinement.




Figure 3.6: GPS Condensed Tree




Figure 3.7: GPS UMAP Clusters

The corresponding UMAP projection in Figure 3.7 shows the same six-family structure as compact, well-separated latent components. This agreement between the condensed hierarchy and the low-dimensional chart indicates that the recovered GPS families are supported by both density persistence and neighborhood-preserving geometry [94, 95, 92, 93].




Figure 3.8: GPS Orbital Elements Pairplot

Figure 3.8 verifies the latent clustering against classical orbital elements by showing that the recovered families occupy distinct RAAN groupings while remaining concentrated within a shared medium-altitude orbital regime. Semi-major axis and inclination define the shared medium-altitude shell, but the plane structure is specified through nodal orientation and therefore appears as six discrete RAAN groupings in classical-element space [108, 109, 96, 1]. This behavior is consistent with the ROPV-to-UMAP-to-HDBSCAN pipeline. The representation preserves orbital-

plane geometry, the embedding preserves local neighborhood structure, and the density extractor identifies connected components without imposing a fixed cluster count.

The GPS benchmark also supplies the operational definition of success used throughout this chapter. The goal is recovery of stable, compact, and plane-consistent families that reproduce the documented architecture at clustering-relevant resolution. Public documentation defines a six-plane baseline satellite constellation with expandable slots and additional non-baseline satellites, so minor count asymmetries in a synchronized snapshot are expected [108, 109, 96]. Accordingly, the criterion is architectural and topological fidelity, with count-level asymmetry treated as an expected property of the synchronized operational snapshot.

3.4.2 OneWeb as an Architecture-Conditioned Benchmark

OneWeb provides a more stringent benchmark because its public architecture remains regular while appearing less discretely separated in the recovered geometry. Within the near-polar shell, the analysis identifies twelve families, which is consistent with the current public Phase 1 architecture of the 87.9° population [97]. Interpretation nevertheless requires greater caution than in the GPS case because there is no authoritative public family-by-family label set for the synchronized catalog snapshot, and because the present architecture emerged through an FCC-approved reconfiguration from the earlier eighteen-plane design [97]. Benchmarking is therefore conditioned on the documented shell-and-plane architecture under partial external labeling.




Figure 3.9: OneWeb Condensed Tree

Figure 3.9 shows that the OneWeb hierarchy contains a larger and more intricate set of persistent branches than the GPS benchmark, reflecting plane-consistent density structure within a denser near-polar shell.




Figure 3.10: OneWeb UMAP Clusters

The corresponding latent geometry in Figure 3.10 separates the twelve recovered OneWeb families into loops, arcs, elongated components, and compact groups. This structure indicates that the recovered plane-level organization is expressed through nonlinear local manifolds instead of uniformly spherical clusters [94, 95, 92, 93].




Figure 3.11: OneWeb Orbital Elements Pairplot

Classical-element diagnostics in Figure 3.11 support the same interpretation. Semi-major axis, eccentricity, and inclination remain tightly concentrated within the near-polar shell, whereas RAAN provides the dominant plane-level separability. The public architecture is therefore visible in classical coordinates, although the latent chart more clearly resolves the nonlinear family geometry [97, 1].

The recovered OneWeb geometry is qualitatively distinct from the GPS benchmark. The

condensed tree exhibits a stable but more intricate hierarchy, with branches that are more numerous and less cleanly isolated than in the six-plane GPS case. The UMAP embedding still separates the dominant components, but these often appear as loops, arcs, or elongated manifolds instead of compact point clouds. This pattern indicates persistent intra-family structural variation within a documented plane architecture instead of convergence to approximately spherical clusters. The OneWeb case, therefore, provides a demanding intermediate benchmark for the pipeline, because successful recovery must remain robust under nonconvex and manifold-like latent geometry of the same general type encountered later in shell-conditioned Starlink discovery [97, 94, 92].

Classical-element diagnostics support the same interpretation. Semi-major axis, eccentricity, and inclination remain tightly concentrated within the near-polar shell, whereas RAAN provides the dominant plane-level separability. The public architecture is therefore still present, but it is embedded in a denser and less discretely separated local geometry than in GPS. This is precisely the regime in which neighborhood-preserving embedding and density-based extraction are preferable to methods that presuppose convex clusters or homogeneous within-cluster variance [105, 106, 107]. The OneWeb result is accordingly interpreted as architecture-conditioned recovery of plane-consistent structure in a geometrically nonlinear regime.

3.5 Starlink as a Shell-Conditioned Discovery Problem

Benchmark results establish that the pipeline recovers externally documented plane architectures under both strongly labeled and architecture-conditioned regimes. Starlink introduces a qualitatively different inference problem. Public filings define shell-level parameters, including altitude bands, inclination regimes, and nominal plane counts, while leaving family-level labels unavailable at the granularity represented in synchronized catalog snapshots [61, 4]. The analytical objective, therefore, shifts from supervised recovery to shell-conditioned unsupervised discovery under weak architectural priors [90].

This setting requires shell-conditioned analysis. The Starlink satellite constellation is a superposition of multiple design strata, each defined by a distinct orbital regime. When treated globally, these strata dominate the similarity geometry and suppress the finer-scale structure that defines within-shell organization. When stratified, the apparent global irregularity resolves into a limited set of recurring geometric regimes with consistent topological signatures.

3.5.1 Shell-Conditioned Analysis

Global clustering of the Starlink constellation conflates two structurally distinct scales. The first is shell-level separation, governed by altitude and inclination. The second is within-shell organization, governed by orbital-plane orientation, occupancy distribution, deployment history, and station-keeping behavior [61, 4, 1, 2]. A single unsupervised partition across all shell weights the metric toward inter-shell separability, which is large in magnitude, and thereby suppresses the contrast required to resolve intra-shell families.

This distinction has both physical and statistical foundations. The authorized Gen 1 Starlink architecture includes multiple shells with distinct inclinations and altitude ranges, notably near 53° , 53.2° , 70° , and 97.6° , with orbital altitudes concentrated in the 540–570 km regime [61, 4]. These shells correspond to different coverage objectives and operational regimes, and therefore constitute a stratified population composed of several distinct but internally structured subpopulations.

Within this stratified system, the low-inclination shells near 53° and 53.2° are densely populated and exhibit behavior consistent with Walker-type constellation geometry, producing continuous and highly sampled manifold structure in the latent chart [13, 14, 61]. The 70° shell represents an intermediate regime, with reduced population density and less uniform geometric organization. The 97.6° polar shell exhibits a more discretized configuration associated with its distinct coverage role and deployment structure [61].

Applying a global clustering model to this heterogeneous population imposes an implicit assumption of geometric homogeneity that is inconsistent with the underlying system design. The dominant inter-shell separations then saturate the similarity metric, masking the finer-scale variations that encode operational structure within each shell.

Shell-conditioned decomposition removes this degeneracy. Let

$$\mathcal{I}_{\text{cat}} = \bigcup_{\ell \in \mathcal{L}_{\text{shell}}} \mathcal{I}^{(\ell)} \quad (3.32)$$

denote a partition of the synchronized catalog index set into shell-conditioned subsets, where $\mathcal{L}_{\text{shell}}$ indexes the distinct inclination-altitude regimes. The family-discovery pipeline is then applied independently to each subset $\mathcal{I}^{(\ell)}$, yielding shell-specific family sets

$$\mathcal{F}^{(\ell)} = \{f_s : s \in \mathcal{I}^{(\ell)}\}. \quad (3.33)$$

This decomposition removes large-scale inter-shell offsets from the similarity geometry and allows the embedding and density-extraction stages to operate at the intrinsic geometric scale of each shell

[94, 95, 92, 93].

Empirical results demonstrate that this stratified approach reveals distinct topological regimes. The low-inclination shells produce densely sampled continuous manifolds in the latent space. The intermediate shell produces fragmented or partially connected structures. The polar shell produces strongly discrete and well-separated components. These regimes are interpreted as signatures of underlying differences in satellite constellation design, deployment sequencing, sampling density, and operational control. The stratified structure, therefore, constitutes the central scientific result of the Starlink analysis: apparent global irregularity resolves into a small number of recurring geometric archetypes when conditioned on shell-level architecture.

3.5.2 The 53.0° Shell

The 53.0° shell provides the clearest example of a densely sampled continuous latent manifold in the shell-resolved analysis. In the UMAP embedding restricted to this shell, the population forms an annular, horseshoe-like trajectory with coherent local sectors and nearly uninterrupted global support. The resulting geometry indicates structured local organization within a continuous support, so the recovered families are interpreted as density-supported sectors of a common shell manifold [94, 95].

The condensed-tree diagnostic supports this interpretation. The hierarchy contains branches with broadly comparable density levels and similar persistence ranges on the λ scale, without a single branch dominating the shell. This pattern is consistent with distributed persistence, in which local density-supported sectors are recoverable while the global shell remains connected through gradual transitions in the latent geometry. In this regime, HDBSCAN acts as a coarse-graining operator on a structured continuum, yielding operationally useful families even though the underlying population exhibits gradual continuity across the shell [92, 93].

The classical-element diagnostics further support the continuous-manifold interpretation. Inclination remains narrowly concentrated near the shell value, eccentricity is small, and RAAN spans the angular domain with dense occupancy. The dominant visual structure is therefore a sampled progression through nodal orientation. The authorized architecture of this shell, with many orbital planes at approximately constant altitude and inclination, makes this latent geometry physically plausible [61, 1, 2]. The main analytical value of the embedding is that it resolves local sector structure in a regime where raw classical-element projections are saturated.

3.5.3 The 53.2° Shell

The 53.2° shell exhibits a topology closely related to the 53.0° shell, with an even smoother annular organization in the retained latent chart. The population forms a high-occupancy ring-like manifold with reduced visual fragmentation and weaker local discontinuity. The recovered family structure is therefore best interpreted as local sectorization of a continuous shell geometry, with family boundaries reflecting density variation along the manifold [94, 95].

The condensed tree reflects the same continuity. Its branch structure is smoother and less jagged than that of the 53.0° shell, indicating weak discretization at the clustering scale. Local sectors remain statistically recoverable, and the hierarchy supports a continuous-shell interpretation with weak evidence for a few sharply separated operator-level families. HDBSCAN therefore provides a useful operational partition of a dense manifold, while the scientific interpretation remains that the shell is fundamentally continuous at the resolution of the synchronized catalog snapshot [92, 93].

Classical-element diagnostics emphasize why the embedding is necessary. Semi-major axis, eccentricity, and inclination are tightly concentrated, while RAAN covers the angular domain broadly. The public architecture authorizes this shell as a densely populated 53.2° configuration with many orbital planes, making a ring-like latent geometry consistent with the expected nodal progression [61]. UMAP represents a high-density neighborhood graph in which adjacent planes and nearby local sectors are connected through continuous similarity relationships.

3.5.4 The 70.0° Shell

The 70.0° shell forms an intermediate topological regime between the low-inclination annular shells and the discrete polar shell. In the shell-resolved UMAP embedding, the population fragments into islands, filaments, and small loop-like structures separated by visible gaps. This makes the 70.0° shell the most fragmented of the nonpolar regimes and indicates discontinuous population support within the shell-conditioned geometry [94, 95].

The condensed tree is correspondingly uneven. The branch structure is jagged, with persistence distributed across separated local components. The hierarchy still encodes meaningful density structure, and a fragmented-support interpretation better matches the observed separated local components. HDBSCAN therefore identifies families that correspond to separated occupancy regions within the shell, while also preserving the possibility that some objects represent transitional or sparsely sampled segments of the population [92, 93].

The documented shell architecture supports this physical interpretation. The 70° shell has fewer authorized planes than the low-inclination 53°-class shells, and the observed catalog pop-

ulation is less uniformly filled [61]. Reduced plane count, local occupancy gaps, deployment chronology, station-keeping variation, and mixtures of operational and transitional states can all generate separated embedded components. The recovered geometry should therefore be interpreted as partial shell fragmentation, with cluster count treated as a density-geometry outcome distinct from a direct architectural plane count.

3.5.5 The 97.6° Shell

The 97.6° shell is the most discrete regime in the shell-resolved analysis. The UMAP embedding contains a small number of compact components separated by large gaps, and the corresponding condensed tree displays a few dominant branches with stronger persistence than those observed in the lower-inclination shells. This is the case most closely aligned with a classical disconnected-cluster configuration in the latent geometry [92, 93].

This result requires cautious interpretation. Public authorization specifies two 97.6° polar subconfigurations, one with 6 planes and one with 4 planes [61]. A discrete low-count embedding is therefore expected. The number of recovered components in a synchronized public catalog snapshot remains distinct from the operator-intended plane architecture. Occupancy, partial deployment, local mergers under the ROPV similarity geometry, shell-specific sampling density, and catalog incompleteness may cause the recovered partition to represent a subset, aggregation, or distorted projection of the authorized structure.

The polar result, therefore, confirms topological discreteness while providing limited evidence for an exact plane-count reconstruction. The 97.6° regime is the most strongly separated shell in the retained latent geometry, and its compact components are consistent with a low-count polar architecture. Precise correspondence between recovered families and operator-defined plane groups remains model-dependent and should be evaluated through supporting diagnostics in classical-element space, especially RAAN concentration, inclination compactness, and family-level membership counts [61, 1, 2].

3.5.6 Cross-Shell Comparison

Shell-resolved analysis reframes the global Starlink clustering as a structured comparative examination of distinct latent topologies. The 53.0° shell forms a densely sampled annular manifold characterized by distributed persistence and moderate sectorization. The 53.2° shell displays a smoother annular architecture, with reduced discretization and increased uniformity. The 70.0° shell diverges from these low-inclination regimes, manifesting as a fragmented ensemble of islands,

loops, and short filaments. The 97.6° shell culminates this sequence as a compact, highly discrete, multi-component structure.




Figure 3.12: Starlink Condensed Tree

Figure 3.12 summarizes the shell-conditioned HDBSCAN hierarchies and shows that each Starlink shell possesses a distinct persistence structure. The low-inclination shells exhibit distributed branch persistence, the 70.0° shell exhibits fragmented branching, and the 97.6° shell exhibits a more discrete hierarchy.




Figure 3.13: Starlink UMAP Clusters

Figure 3.13 provides the corresponding shell-conditioned latent charts. The 53.0° and 53.2° shells form annular manifolds with sector-like families, the 70.0° shell separates into fragmented island and filament structures, and the 97.6° shell resolves into compact components with strong inter-component separation.




Figure 3.14: Starlink Orbital Elements Pairplot

The shell-conditioned pair plots in Figure 3.14 ground the latent families in classical orbital-element diagnostics. Semi-major axis and inclination confirm shell membership, eccentricity remains concentrated in the near-circular regime, and RAAN supplies the principal angular coordinate through which within-shell organization becomes physically interpretable. Shell-conditioned labels represent local taxonomies tailored to distinct geometric regimes. This distinction enables interpretation of the forthcoming global remapping as a synthesis across heterogeneous shell-level structures of the satellite constellation.

Starlink comprises a family of shell-conditioned clustering problems, each defined by distinct intrinsic geometries and appropriate coarse-graining scales. Low-inclination shells yield persistent

sectors within near-continuous supports. The 70° shell reveals components embedded in a partially fragmented manifold, while the polar shell presents strongly discrete, density-supported elements. Starlink is therefore described as exhibiting shell-dependent latent topology, spanning continuous annular manifolds, fragmented intermediates, and discrete polar configurations [61, 4, 94, 95, 92, 93].

3.6 Global Satellite Synthesis

Shell-resolved analysis indicates that Starlink possesses heterogeneous latent geometries across shell-conditioned regimes. The current objective is to integrate these local findings into a coherent global taxonomy while maintaining the shell-conditioned structure established by the authorized shell architecture and the preceding shell-resolved decomposition [61, 4]. The synthesis process consists of four stages. First, the procedure establishes the remapping of shell-local labels onto a continuous global index. Next, the resulting global UMAP embedding is interpreted. Subsequently, centroid-level diagnostics assess whether latent coordinates can be reduced to a classical-element surrogate. Finally, classical-element pair plots are used diagnostically to constrain interpretation while leaving family inference to the latent-density pipeline.

3.6.1 Local-to-Global Remapping

The shell-resolved clustering stage produces a collection of local taxonomies that must be remapped before population-level synthesis. The synchronized Starlink population is decomposed into shell-conditioned subsets

$$\mathcal{I}_{\text{cat}} = \bigcup_{\ell \in \mathcal{L}_{\text{shell}}} \mathcal{I}^{(\ell)} \quad (3.34)$$

with each shell ℓ carrying a local label set

$$\mathcal{C}_{\text{loc}}^{(\ell)} = \{-1, 0, 1, \dots, L_\ell\} \quad (3.35)$$

In this context, negative labels indicate HDBSCAN noise within the local shell analysis [92, 93]. Global synthesis requires embedding these shell-local labels within a unified index space to enable joint enumeration, integrated visualization, and cross-shell comparison of family-level results.

The implementation performs this by defining, for each shell ℓ , a local-to-global map

$$\phi_\ell: \mathcal{C}_{\text{loc}}^{(\ell)} \rightarrow \mathbb{N}_0 \quad (3.36)$$

Every non-noise local label is assigned a unique positive global identifier, while all local noise labels are mapped to the reserved global value 0. For any satellite $s \in \mathcal{I}^{(\ell)}$ with local label $c_s^{(\ell)} \in C_{\text{loc}}^{(\ell)}$, the remapped global label is

$$g_s = \phi_\ell(c_s^{(\ell)}) = \begin{cases} 0, & c_s^{(\ell)} < 0 \\ \gamma_\ell(c_s^{(\ell)}), & c_s^{(\ell)} \geq 0 \end{cases} \quad (3.37)$$

The resulting global index remains continuous across shell-conditioned groups and is ordered according to the numerical sorting of the shell labels.

Remapping serves an administrative role while preserving the logic of shell-conditioned discovery. The global family collection $\mathcal{F}_{\text{glob}}$ supports population-level statistics, centroid summaries, cross-shell comparisons, and integrated visualization without collapsing local structure into a new clustering framework. The global taxonomy is derived from shell-resolved analyses and retains the original shell-conditioned clustering decisions.

3.6.2 Global UMAP Structure

The remapped global UMAP embedding $\mathbf{Y}_{\text{glob}}$ appears as a stratified assembly of shell-conditioned local families with shared population-level organization. The global plot reveals arcs, loops, short filaments, and compact islands, each situated within larger neighborhoods whose internal organization reflects the previously established shell-resolved structures. Several families maintain loop-like geometries, some form short linear strands, and others concentrate into compact, high-density components. This arrangement is consistent with expectations when heterogeneous shell-level manifolds are integrated while preserving local topology [94, 95].




Figure 3.15: Global Starlink UMAP Clusters

Figure 3.15 visualizes the remapped global family taxonomy as an integrated atlas of shell-conditioned structures. The figure preserves the local shell geometries while enabling population-level comparison across the full synchronized Starlink snapshot. Intermediate-scale organization predominates in the structural landscape. The global embedding preserves shell-conditioned neighborhoods and maintains distinct local families within these regions. The global latent space serves as an atlas of related local manifolds with shell-conditioned structure retained in a shared coordinate representation [94, 95].

Shell-resolved results corroborate this interpretation. The 53.0° and 53.2° shells yield annular and weakly discretized local sectors, the 70° shell produces fragmented island-like components, and the 97.6° shell forms compact discrete groups. The global embedding integrates these inherited structures into a unified coordinate chart, preserving their distinct topologies. Consequently, the remapped UMAP plot visualizes the cross-shell taxonomy inherited from shell-resolved clustering [61, 4, 94].

3.6.3 Centroid-Level Feature Diagnostics

Centroid-level correlation analysis assesses whether the global latent coordinates re-express a single classical orbital element. For each non-noise global family $g \in \mathcal{F}_{\text{glob}}$, let

$$\mathcal{I}_g = \{s \in \mathcal{I}_{\text{cat}} : g_s = g\} \quad (3.38)$$

denote its member set, and let $\bar{\mathbf{b}}_g$ denote the corresponding centroid-level diagnostic feature vector assembled from selected classical elements and latent-coordinate summaries. Collecting these centroids gives

$$\mathbf{B}_{\text{cent}} = \begin{bmatrix} \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_2^\top \\ \vdots \\ \bar{\mathbf{b}}_G^\top \end{bmatrix}, \quad \mathbf{R}_{\text{cent}} = \text{corr}(\mathbf{B}_{\text{cent}}) \quad (3.39)$$

where $\text{corr}(\cdot)$ denotes the Pearson product-moment correlation matrix, used here as a linear association diagnostic with no causal interpretation assigned [110, 111]. The strongest correlations are observed among physical shell descriptors, including semi-major axis, eccentricity, inclination, and argument of perigee. UMAP coordinates demonstrate only weak to moderate linear associations with any individual physical feature. Among these, UMAP₂ exhibits the clearest, though still moderate, relationship to RAAN. Other latent coordinates display minimal association and do not serve as proxies for single elements.




Figure 3.16: Centroid-Level Correlation for the Starlink Satellite Constellation

Figure 3.16 shows the centroid-level correlation structure between classical orbital descriptors and latent coordinates. The absence of a dominant one-to-one alignment between any UMAP coordinate and a single classical element indicates that the global embedding retains composite neighborhood information from the vector-valued representation. This outcome reflects the pipeline architecture, which relies on vector-valued representations and prioritizes local neighborhood fidelity over variance maximization or one-dimensional interpretability [94, 95, 101]. The correlation result should therefore be interpreted as a diagnostic constraint on axis-level interpretation, with nonlinear embedding geometry evaluated through neighborhood preservation and density structure [110, 111].

Apparent coupling between eccentricity and argument of perigee provides insufficient support for treating argument of perigee as a reliable explanatory variable for the near-circular Starlink pop-

ulation. In low-eccentricity regimes, perigee-based angles are subject to coordinate pathologies, numerical instability, and dependence on conventions [1, 2, 98]. Their presence in the correlation figure serves as a diagnostic for differences among remapped family centroids, with limited evidentiary weight for primary physical relevance.

3.6.4 Cluster Orbital Element Diagnostics

Classical-element pair plots retain value as global diagnostics. These plots delineate residual shell-scale separations in raw orbital elements and highlight where latent taxonomic structure succeeds beyond the capabilities of unprocessed plots. Once populations are remapped globally, these plots cease to provide the principal basis for inferring family structure [1, 2].

The reduced pair plot of the largest global families in Figure 3.17, denoted by $\mathcal{F}_{\text{top}}$, provides maximal interpretive clarity for the main analysis. Shell stratification persists distinctly in semi-major axis and inclination, while RAAN captures much of the remaining intrafamily separability. This grounds remapped families in established orbital mechanics. Nevertheless, the plot exhibits substantial crowding relative to the latent chart. Overlapping low-eccentricity and near-common-inclination features persist, precluding unequivocal extraction of global families from raw elements alone [1, 2, 98].




Figure 3.17: Global Starlink Orbital Elements Pairplot

3.7 Interpretation, Limits, and Methodological Boundaries

The preceding sections introduced a data representation, an embedding protocol, and a density-based family extraction method, applying this pipeline to recover benchmark structures and construct a shell-conditioned Starlink taxonomy. The present section imposes interpretive discipline by delineating the semantic scope of the latent chart, identifying limitations arising from the public data product, and specifying the claims substantiated by the current validation framework.

3.7.1 Latent Space

The latent space functions as a coordinate chart for the data-induced neighborhood geometry across synchronized catalog states. Scientifically valid inferences pertain to local adjacencies defined by the ROPV representation, their density stratifications, and stable connected components that persist through the HDBSCAN hierarchy. At this level of resolution, the embedding preserves defensible structural content. Proximity within the latent chart indicates shared neighborhoods in the high-dimensional feature space, while persistent components correspond to families occupying analogous regions in the induced orbit-similarity geometry [94, 95, 92, 93].

The UMAP axes do not possess intrinsic orbital semantics. Point locations on UMAP_1 , UMAP_2 , or UMAP_3 arise from the optimization of fuzzy neighborhood memberships, with no direct projection onto conserved orbital coordinates. Absolute distances require cautious interpretation. Although large separations may indicate distinct geometric regimes, the embedding does not provide a globally isometric representation of orbital-state space. Euclidean distance in the latent chart should not be interpreted as maneuver cost, dynamical separation, or operational interchangeability without independent validation of those correspondences [94, 95].

This interpretive stance is consistent with the aim-relative approach to clustering adopted throughout the analysis. The taxonomy does not assert uniqueness as the natural partitioning of Starlink. It represents a representation- and aim-dependent coarse-graining that preserves local geometry essential for shell-resolved structural identification, family-level interpretation, and subsequent physical inference [89, 90]. The latent chart provides value by organizing the observable population into stable, physically interpretable neighborhood structures validated through benchmarking and diagnostic procedures.

3.7.2 TLE and SGP4 Data Products

This analysis is subject to the constraints of the public catalog product. TLEs are GP mean elements designed for SGP4 propagation and do not include osculating truth states, public covariance data, or maneuver history [16, 17, 88, 20]. State reconstructions are therefore conditioned on the SGP4-compatible mean-element framework and the conventions used in transforming catalog quantities to inertial position and velocity vectors at the analysis epoch.

Model dependence is evident in two primary domains. First, the feature map is conditioned by the mean-element product, and the ROPV geometry summarizes the catalog state as a public mean-element diagnostic. Second, the synchronized dataset constitutes a temporal snapshot, which mitigates some catalog inconsistencies but does not resolve temporal continuity, station-keeping

intent, or transfer history. Consequently, a family identified in this context represents an empirical signature of persistent structure within a mean-element snapshot, with no temporally resolved dynamical-object status [16, 17, 88].

Similar caution is necessary when interpreting classical elements. In dense, near-circular shells, the argument of perigee and the true anomaly are unreliable descriptors, as these become undefined or convention-dependent as eccentricity decreases [1, 2, 98]. This limitation restricts the function of global diagnostic pair plots while preserving their utility as shell-scale diagnostics. These plots remain effective for confirming shell stratification and RAAN organization, whereas primary family interpretation is based on vector-valued representations, neighborhood-preserving embeddings, and density hierarchies.

3.7.3 Hyperparameter Sensitivity, Stochasticity, and Validation

The current validation framework supports claims limited to local neighborhood preservation between the input feature space and the selected embedding. Trustworthiness quantifies this local fidelity and supports interpretation of the embedding as a representation of local density structure [104, 102]. However, trustworthiness does not address global metric fidelity, axis interpretability, or comprehensive hyperparameter invariance, thereby restricting its scope to local validation of latent coordinates.

Condensed-tree and cluster-persistence diagnostics demonstrate stability across varying density thresholds in the embedded space. Persistent branches reveal structural components that remain robust under threshold variation for the chosen representation and metric [103, 92, 93]. However, these diagnostics do not establish direct correspondence with operator labels, dynamical invariants, or distinct physical categories. For benchmark satellite constellations, persistence combined with external architectural congruence provides a strong validation argument. In the case of Starlink, persistence, supported by shell-level priors and internal consistency, provides a provisional evidentiary basis [61, 90].

DBCV serves as an internal validity criterion suited to irregular, density-connected clusters and is better matched to this problem than compactness-based indices designed for centroidal partitions [105, 106, 107]. DBCV evaluates coherence within the same density geometry as the clustering process, but does not independently validate external ground truth. Within this workflow, reported DBCV values are further influenced by implementation choices, including label handling, mutual-reachability tree construction, noise management, and numerical safeguards [107].

Stochasticity and hyperparameter sensitivity limit the current methodology. UMAP’s stochastic optimization requires explicit random seed control, metadata export, and thread management

to improve reproducibility [94, 95]. Nevertheless, dependence remains on parameter selection, feature scaling, and the single-epoch nature of the synchronized catalog. As a result, the strongest claims are structural, supporting a shell-conditioned latent taxonomy that is stable within the current pipeline and interpretable under public architectural priors. Claims regarding uniqueness, temporal persistence, or operator intent require further analysis and larger datasets.

3.7.4 Methodological Limitations

Three principal limitations characterize the current formulation. First, representation dependence results from reliance on the ROPV map, standardized feature scaling, and shell-conditioned pre-processing choices [1, 2, 101]. Second, data-product dependence arises from TLE-derived mean elements, which compress dynamics through SGP4 modeling and exclude covariance and control history [16, 17, 88, 20]. Third, temporal limitations stem from analyzing synchronized snapshots without multi-epoch trajectory linkage, restricting the analysis to structural regimes at a single epoch, and preventing assessment of long-term family persistence.

Chapter 4

Forecast-Conditioned Thermospheric Forcing

4.1 Forecast-Conditioned Thermospheric Forcing for Future Satellite Propagation

The propagation of large LEO satellite constellations requires a thermospheric forcing model that resolves atmospheric-density variability driven by solar radiative activity, geomagnetic energy deposition, local solar time, geodetic latitude, geodetic altitude, seasonal structure, and solar-cycle phase [34, 35, 112, 113, 36]. For long-horizon orbit simulation, the primary modeling requirement is the generation of forecasted solar and geomagnetic driver histories compatible with Naval Research Laboratory Mass Spectrometer and Incoherent Scatter Radar model 2.0 (NRLMSIS 2.0) [35, 81]. The resulting methodology replaces fixed reference-atmosphere assumptions with forecast-conditioned thermospheric density fields evaluated over future intervals relevant to satellite propagation, drag accumulation, shell maintenance, and disposal analysis.

4.1.1 Future Atmospheric Forcing

The developed orbit simulator enables forward propagation through future epochs, which requires the atmospheric model to evolve consistently with forecasted space-weather conditions. Reference atmospheres such as the U.S. Standard Atmosphere 1976 (USSA76) provide baseline vertical stratification, yet omit the temporal evolution of solar-cycle phase, geomagnetic disturbances, and storm-time density enhancement that govern secular drag accumulation in LEO [80, 34, 35]. This omission is acceptable for short-duration reference calculations and becomes restrictive for

extended propagation of actively maintained satellite constellations because density variability is a dominant source of environmental nonstationarity.

This restriction is particularly important for Starlink-like satellite constellations, where semi-major axis decay, shell residence duration, lowering corridors, nodal drift, operational turnover, and disposal chronology are determined by the cumulative density encountered along each orbital path [12, 36, 34, 35]. Accurate simulation over future intervals, therefore, requires thermospheric forcing conditioned on predicted solar and geomagnetic drivers.

Let the merged daily space-weather archive be denoted by $\mathcal{D}_{\text{SW}}$. The multivariate driver vector assembled at epoch t is written as

$$\mathbf{d}_{\text{SW}}(t) = \left[F_{10.7}^{\text{obs}}(t) \quad F_{10.7}^{\text{adj}}(t) \quad F_{10.7}^{\text{URSI}}(t) \quad \bar{F}_{10.7}(t) \quad K_p(t) \quad a_p(t) \quad D_{\text{st}}(t) \quad S_N(t) \right]^{\top} \quad (4.1)$$

where $F_{10.7}^{\text{obs}}$, $F_{10.7}^{\text{adj}}$, and $F_{10.7}^{\text{URSI}}$ denote observed, adjusted, and URSI-D 10.7 cm solar radio flux channels, respectively [112, 114]. The smoothed background solar radio flux term $\bar{F}_{10.7}$ represents the longer-window solar activity input required by the empirical atmosphere model [34, 35, 113]. The geomagnetic channels K_p , a_p , and D_{st} encode planetary geomagnetic activity, NRLMSIS-compatible equivalent-amplitude forcing, and storm-time ring-current disturbance information, while S_N denotes sunspot number as an auxiliary solar-activity descriptor [115, 116, 117, 118, 119, 120].

For an initialization epoch t_0 , a lookback window H_{look} , and a forecast horizon H_{fcst} , the forecasting pipeline maps the historical driver archive into a future driver sequence,

$$\widehat{\mathbf{d}}_{\text{SW}}(t_0 : t_0 + H_{\text{fcst}}) = \mathcal{M}_{\text{RIT}} \left[\mathbf{d}_{\text{SW}}(t_0 - H_{\text{look}} : t_0) \right] \oplus \mathcal{M}_{\text{gm}} \left[\mathbf{d}_{\text{SW}}(t_0 - H_{\text{look}} : t_0) \right] \quad (4.2)$$

where $\mathcal{M}_{\text{RIT}}$ denotes the ResNeXt-iTransformer forecasting map used for solar radio flux prediction, $\mathcal{M}_{\text{gm}}$ denotes the conditional geomagnetic-driver model, and $\oplus$ denotes assembly of the solar and geomagnetic forecast components into a single NRLMSIS-compatible driver sequence [121, 122, 123, 124, 125].

The forecasted drivers are supplied to NRLMSIS 2.0 through an empirical thermosphere map,

$$\widehat{\rho}_{\text{fc}}(t, \varphi_g, \lambda_g, h_g) = \mathcal{M}_{\text{NRLMSIS}} \left(t, \varphi_g, \lambda_g, h_g, \widehat{\mathbf{d}}_{\text{SW}}(t) \right) \quad (4.3)$$

where $\widehat{\rho}_{\text{fc}}$ is the forecast-conditioned thermospheric mass-density field, φ_g is geodetic latitude, λ_g is geodetic longitude, and h_g is geodetic altitude on the World Geodetic System 1984 (WGS84) ellipsoid [35, 81, 82]. The forecast-conditioned density-grid snapshot at density epoch t_{ℓ}^{ρ} is then

written as

$$\mathcal{G}_{\rho, \text{fc}}^{(\ell)} = \widehat{\rho}_{\text{fc}} \left(t_{\ell}^{\rho}, \varphi_g, \lambda_g, h_g \right) \quad \text{evaluated on the prescribed WGS84 geodetic grid} \quad (4.4)$$

These snapshots form the future dynamic-atmosphere input used by the propagation architecture through the same interpolation interface previously denoted by ρ_{dyn} .

4.1.2 Scope of the Present Formulation

Predictive solar and geomagnetic driver histories are developed using an advanced forecasting pipeline and an expanded merged archive that incorporates D_{st} . These drivers are integrated with the official Fortran implementation of NRLMSIS 2.0 to construct future thermospheric density fields [81, 35]. The resulting output is a physically informed neutral-density forcing model designed for orbit simulation. Dynamo modeling, global magnetosphere simulation, and satellite-parameter inference remain outside the scope of this formulation.

The objective is to generate dynamically credible future density histories for satellite propagation over extended timescales. The approach is empirical and forecast-driven. Historical space-weather data are integrated into the multivariate daily archive $\mathcal{D}_{\text{SW}}$. Multi-year solar radio flux forecasts are produced using the ResNeXt-iTransformer framework, while the geomagnetic component is modeled through the conditional process $\mathcal{M}_{\text{gm}}$. These synthesized driver trajectories are provided to NRLMSIS 2.0 to generate forward thermospheric mass-density fields for propagation, ensemble analysis, and later uncertainty-quantification workflows [121, 122, 123, 124, 125, 35, 36].

4.2 Space-Weather Driver Data

Thermospheric forcing is determined by a defined set of solar and geomagnetic observables that control the background neutral-density environment encountered by low Earth orbit (LEO) satellites. This section defines the solar radio flux and geomagnetic driver families incorporated within the forecasting architecture and describes the unified historical archive used to construct them. The discussion emphasizes physical forcing significance, source-aware data integration, preprocessing discipline, and operational compatibility with NRLMSIS 2.0 [112, 34, 35, 81].

Let the raw collection of source-specific space-weather records be denoted by $\mathcal{D}_{\text{SW}}^{\text{raw}}$. After

synchronization onto a complete daily UTC calendar $\mathcal{T}_{\text{SW}}$, the cleaned archive is represented as

$$\mathcal{D}_{\text{SW}} = \{(t, \mathbf{d}_{\text{SW}}(t)) : t \in \mathcal{T}_{\text{SW}}\} \quad (4.5)$$

where $\mathbf{d}_{\text{SW}}(t)$ is the multivariate daily driver vector used by the forecasting model. The preprocessing operation is written as

$$\mathbf{d}_{\text{SW}}^{\text{clean}}(t) = \mathcal{P}_{\text{SW}} [\mathbf{d}_{\text{SW}}^{\text{raw}}(t)] \quad (4.6)$$

where $\mathcal{P}_{\text{SW}}$ denotes the deterministic preprocessing map that removes nonfinite values, interpolates internal gaps, performs edge filling, and enforces admissible geomagnetic-index ranges. The cleaned archive is subsequently identified with $\mathbf{d}_{\text{SW}}(t)$ unless a distinction between raw and processed records is required.

4.2.1 The Solar Radio Flux Driver Family

The primary radiative driver is the 10.7 cm solar radio flux, $F_{10.7}$, measured at 2800 MHz and commonly reported in solar flux units, where $1 \text{ sfu} = 10^{-22} \text{ W m}^{-2} \text{ Hz}^{-1}$. The strong empirical correlation with solar activity establishes $F_{10.7}$ as a reliable proxy for upper-atmosphere heating and solar extreme-ultraviolet forcing in empirical thermosphere modeling [112, 113]. In the context of LEO propagation, $F_{10.7}$ captures the temporal evolution of the solar state that influences thermospheric density across solar-rotation, seasonal, and solar-cycle timescales [34, 35, 36].




Figure 4.1: Historical 10.7 cm Solar Radio Flux

Figure 4.1 shows the multi-decadal evolution of the observed, adjusted, and URSI-D solar radio flux channels, making visible both the recurring solar-cycle envelope and the short-duration

enhancements that influence thermospheric density forcing. Three representations of the solar radio flux record are utilized:

$$\mathcal{S}_{F10.7} = \{F_{10.7}^{\text{obs}}, F_{10.7}^{\text{adj}}, F_{10.7}^{\text{URSI}}\} \quad (4.7)$$

The observed channel $F_{10.7}^{\text{obs}}(t)$ preserves the direct daily measurement. The adjusted channel $F_{10.7}^{\text{adj}}(t)$ corrects the measurement to a standard Sun-Earth distance, thereby isolating intrinsic solar radiative variability from annual distance modulation. The URSI-D channel $F_{10.7}^{\text{URSI}}(t)$ applies the standardized Series D scaling used in geophysical modeling contexts [114, 126]. Although these channels share the same dominant solar-cycle structure, each carries a distinct preprocessing convention. Their joint use reduces sensitivity to a single source representation and permits the forecasting model to separate common low-frequency solar structure from channel-specific scaling and amplitude effects.

The driver vector includes the solar radio flux family as

$$\mathbf{d}_{\text{SW}}(t) \supset \left[F_{10.7}^{\text{obs}}(t), F_{10.7}^{\text{adj}}(t), F_{10.7}^{\text{URSI}}(t), \bar{F}_{10.7}(t) \right]^T \quad (4.8)$$

where $\bar{F}_{10.7}(t)$ denotes the smoothed background solar radio flux term used to represent persistent solar activity in empirical atmosphere forcing [34, 35]. Within the present framework, $F_{10.7}^{\text{adj}}(t)$ is retained as the primary radiative channel for subsequent geomagnetic conditioning and NRLMSIS-compatible forcing, while $F_{10.7}^{\text{obs}}(t)$ and $F_{10.7}^{\text{URSI}}(t)$ provide complementary measurement and modeling representations [112, 114, 126].

4.2.2 Geomagnetic Activity Drivers

Solar radiative forcing is fundamental to the slowly varying thermospheric background, and geomagnetic energy deposition introduces additional storm-time and substorm-time density variability [34, 35, 36]. The forcing archive, therefore, includes the planetary K index K_p , the planetary equivalent-amplitude index a_p , and the disturbance storm time index D_{st} . These variables are collected into the geomagnetic driver family

$$\mathcal{S}_{\text{gm}} = \{K_p, a_p, D_{\text{st}}\} \quad (4.9)$$

The combined representation provides a compact description of the disturbance environment across quasi-logarithmic planetary activity, linearized empirical-atmosphere forcing, and storm-time ring-current structure [115, 116, 117, 127, 118, 128].

The $K_p(t)$ index provides a planetary-scale measure of geomagnetic disturbance derived

from standardized magnetic observatory records [115, 116, 117, 127]. The $a_p(t)$ index supplies a linearized equivalent-amplitude representation derived from K_p , which makes it useful for statistical modeling and empirical density forcing [115, 116, 127]. Within the forecasting pipeline, $a_p(t)$ serves as the primary geomagnetic forecast target because its linear scale supports conditional mean estimation, stochastic residual modeling, and direct construction of NRLMSIS-compatible disturbed forcing scenarios [34, 35, 117].

The $D_{st}(t)$ channel extends the archive by encoding large-scale storm-time ring-current dynamics [118, 128, 119]. Its inclusion improves the physical specificity of the driver state supplied to the forecasting model by adding a disturbance measure associated with global magnetic-storm evolution. Accurate specification of future thermospheric density requires representation of both the persistent radiative background and the dynamic geomagnetic disturbance state. The expanded driver archive is therefore written schematically as

$$\mathbf{d}_{sw}(t) = \left[N_{\odot}^{\text{rot}}(t) \ F_{10.7}^{\text{obs}}(t) \ F_{10.7}^{\text{adj}}(t) \ F_{10.7}^{\text{URSI}}(t) \ \bar{F}_{10.7}(t) \ K_p(t) \ a_p(t) \ S_N(t) \ D_{st}(t) \right]^{\top} \quad (4.10)$$

with source-specific archive columns omitted from this expression when these are unavailable or redundant for a given model configuration.

4.2.3 Historical Driver Archive

The historical driver archive integrates three primary data sources into a unified, daily resolved record prior to model construction. The `SW-All.csv` dataset from CelesTrak provides solar radio flux channels, Bartels-cycle metadata, sunspot number, and geomagnetic summaries [126]. The `dst_index_daily.csv` file from the World Data Center for Geomagnetism, Kyoto supplies the D_{st} index [119, 128]. The daily sunspot archive from the Solar Influences Data Analysis Center (SILSO) at the Royal Observatory of Belgium contributes the total sunspot number $S_N(t)$, which serves as an exogenous descriptor of solar-cycle phase [120]. These datasets are merged onto the complete daily calendar $\mathcal{T}_{sw}$ to establish a temporally uniform state for model input and to eliminate irregular, source-driven missingness before time-series forecasting.

Following the merging process, the daily state is represented by the following archive variables:

$$\{\text{Rotation, Obs, Adj, URSI-D, Kp, ap, Total, Dst}\}$$

The **Rotation** channel is represented by $N_{\odot}^{\text{rot}}(t)$ and encodes a source-specific solar-rotation count. In the CelesTrak convention, this channel corresponds to the Bartels Solar Rotation Number

(BSRN), with the companion variable $N_D(t)$ identifying the day within the Bartels 27-day cycle [126]. In the solar-flux source convention, the reported rotation count is the Carrington Rotation Number (CRN), which is used for long-term solar-flux organization [114]. The Obs, Adj, and URSI-D channels define the radio-flux driver family $\mathcal{S}_{F10.7}$. The Kp and ap channels quantify geomagnetic activity through $K_p(t)$ and $a_p(t)$. The Total channel records the daily sunspot number $S_N(t)$, while Dst provides the disturbance storm time channel $D_{st}(t)$ [126, 120, 117, 127, 119]. This structure balances compactness for long-horizon forecasting with sufficient breadth to represent the principal empirical drivers of thermospheric density.

Preprocessing is intentionally minimal to preserve source interpretability and improve numerical stability. Infinite values are removed. Internal missing entries are filled by linear interpolation, and leading or trailing gaps are handled by edge filling. Geomagnetic indices are clipped to physically admissible ranges before model training, with the admissible ranges defined by the conventional K_p , a_p , and D_{st} index scales and source metadata [127, 117, 119]. The archive is optimized for stable multivariate forecasting, harmonic augmentation, and subsequent NRLMSIS-compatible density generation. Its primary role is to provide a consistent daily forcing state for the construction of forecasted driver trajectories $\hat{\mathbf{d}}_{sw}(t)$, which are then passed to the empirical thermosphere map $\mathcal{M}_{NRLMSIS}$ for forecast-conditioned density generation [126, 117, 120, 35, 81].




Figure 4.2: Historical Daily Sunspot Number

The sunspot number $S_N(t)$, shown in Figure 4.2, serves as an exogenous phase constraint. Its primary utility is to bound large-scale solar evolution across multi-decadal intervals, thereby supporting model performance on extended forecasting horizons [120, 129, 130]. The radio-flux family $\mathcal{S}_{F10.7}$ remains the dominant radiative driver, while $S_N(t)$ stabilizes the learned phase structure and improves representation of long-term solar variability.

4.3 Harmonic Phase Encoding

The historical driver archive exhibits pronounced periodic structure, which is encoded explicitly before model training. Two phase-coordinate families are incorporated into the augmented space-weather driver vector: an annual harmonic basis aligned with the calendar year and a solar-cycle harmonic basis aligned with Schwabe-scale modulation of solar activity. These channels provide physically motivated encodings of established forcing geometry and solar modulation, allowing the forecasting architecture to allocate representational capacity to departures from organized background structure, including solar-rotation harmonics, cycle-to-cycle variability, amplitude asymmetry, and transient storm-time dynamics [131, 132, 112, 133, 130].

Let the augmented driver vector be denoted by

$$\mathbf{d}_{\text{SW}}^{\text{aug}}(t) = \left[\mathbf{d}_{\text{SW}}(t)^\top \quad \mathbf{s}_{\text{ann}}(t)^\top \quad \mathbf{s}_{\text{cyc}}(t)^\top \right]^\top \quad (4.11)$$

where $\mathbf{d}_{\text{SW}}(t)$ is the cleaned multivariate space-weather driver vector, $\mathbf{s}_{\text{ann}}(t)$ is the annual harmonic feature vector, and $\mathbf{s}_{\text{cyc}}(t)$ is the solar-cycle harmonic feature vector. The present subsection defines the annual component, while the solar-cycle component is introduced subsequently.

4.3.1 Annual Harmonic Basis

The first harmonic family encodes annual phase using sine and cosine channels derived from the day-of-year index. Let $n_{\text{doy}}(t) \in \{0, \dots, P_{\text{yr}} - 1\}$ denote the zero-based day-of-year index for sample epoch t , where $P_{\text{yr}} = 365$ days in the present implementation. Let $m_{\text{ann}} \in \{1, \dots, N_{\text{ann}}\}$ denote the annual harmonic order, with $N_{\text{ann}} = 3$. The annual sine and cosine channels are defined by

$$s_{\text{ann},m}^{\text{sin}}(t) = \sin\left(\frac{2\pi m_{\text{ann}} n_{\text{doy}}(t)}{P_{\text{yr}}}\right), \quad s_{\text{ann},m}^{\text{cos}}(t) = \cos\left(\frac{2\pi m_{\text{ann}} n_{\text{doy}}(t)}{P_{\text{yr}}}\right) \quad (4.12)$$

The annual harmonic feature map is therefore written as

$$\Phi_{\text{ann}} : n_{\text{doy}}(t) \mapsto \mathbf{s}_{\text{ann}}(t), \quad (4.13)$$

with

$$\mathbf{s}_{\text{ann}}(t) = \left[s_{\text{ann},1}^{\text{sin}}(t) \quad s_{\text{ann},1}^{\text{cos}}(t) \quad s_{\text{ann},2}^{\text{sin}}(t) \quad s_{\text{ann},2}^{\text{cos}}(t) \quad s_{\text{ann},3}^{\text{sin}}(t) \quad s_{\text{ann},3}^{\text{cos}}(t) \right]^\top \quad (4.14)$$

These six channels form an explicit annual phase coordinate system encoding recurrent structure associated with observation geometry, Sun-Earth distance variation, seasonal sampling

regimes, and the annual organization of geophysical time series [131, 132]. The sine-cosine formulation removes the endpoint discontinuity that would occur if the day of year were supplied as a scalar coordinate, thereby preserving continuity between adjacent calendar endpoints.

Employing three harmonic orders provides sufficient flexibility to represent departures from a purely sinusoidal annual envelope. The fundamental order encodes the dominant annual phase, while the second and third orders provide additional curvature and mild asymmetry in the learned seasonal representation. This structure is useful for multidecadal daily records, where calendar-phase effects can otherwise be absorbed into unrelated driver channels and reduce the interpretability of the learned representation [131, 132].




Figure 4.3: Seasonal Fourier Harmonic Channels

Figure 4.3 visualizes the first three annual sine-cosine harmonic orders used to encode calendar phase. The figure illustrates how the annual basis supplies a continuous periodic representation of seasonal position while avoiding an artificial discontinuity at the year boundary.

4.3.2 Solar-Cycle Harmonic Basis

The second harmonic family encodes phase information associated with Schwabe-scale modulation of solar activity. Let $\tau_{\text{SW}}(t)$ denote the elapsed day index of sample epoch t relative to the first epoch in the merged space-weather archive. Let the nominal solar-cycle period used for harmonic encoding be

$$P_{\text{cyc}} = 11 \times 365 \text{ days} \quad (4.15)$$

This period is a computational phase convention for representing the dominant Schwabe-scale background. Observed solar cycles vary in duration, amplitude, rise time, decay time, and internal morphology [133, 130, 134, 129]. Let $m_{\text{cyc}} \in \{1, \dots, N_{\text{cyc}}\}$ denote the solar-cycle harmonic order, with $N_{\text{cyc}} = 3$. The normalized solar-cycle phase coordinate is defined as

$$\phi_{\text{cyc}}(t) = \frac{\tau_{\text{SW}}(t) \bmod P_{\text{cyc}}}{P_{\text{cyc}}} \quad (4.16)$$

The solar-cycle sine and cosine channels are then defined as

$$s_{\text{cyc},m}^{\sin}(t) = \sin(2\pi m_{\text{cyc}} \phi_{\text{cyc}}(t)), \quad (4.17)$$

$$s_{\text{cyc},m}^{\cos}(t) = \cos(2\pi m_{\text{cyc}} \phi_{\text{cyc}}(t)) \quad (4.18)$$

The solar-cycle harmonic feature map is written as

$$\Phi_{\text{cyc}} : \tau_{\text{SW}}(t) \mapsto \mathbf{s}_{\text{cyc}}(t) \quad (4.19)$$

where

$$\mathbf{s}_{\text{cyc}}(t) = \begin{bmatrix} s_{\text{cyc},1}^{\sin}(t) & s_{\text{cyc},1}^{\cos}(t) & s_{\text{cyc},2}^{\sin}(t) & s_{\text{cyc},2}^{\cos}(t) & s_{\text{cyc},3}^{\sin}(t) & s_{\text{cyc},3}^{\cos}(t) \end{bmatrix}^{\top} \quad (4.20)$$

These six channels constitute a compact coordinate system for representing long-period solar phase. The approximation $P_{\text{cyc}} = 11 \times 365$ days provides a fixed computational scaffold for a process whose observed cycles vary in length and amplitude [133, 130, 134, 129]. The harmonic basis supplies an explicit low-frequency phase representation, situating each daily record within a broad cycle-scale envelope and reducing the burden on finite-window sequence memory to reconstruct slow phase evolution during long-horizon forecasting.




Figure 4.4: Solar Cycle Fourier Harmonic Channels

Figure 4.4 shows the solar-cycle harmonic channels used to encode Schwabe-scale phase information. These channels provide the model with an explicit long-period coordinate system against which departures from nominal cycle morphology can be learned. This architectural approach is essential for multi-year forcing prediction. Explicit solar-cycle harmonics separate the slow background phase from solar-rotation-scale variability and short-period disturbances within the finite input window. Incorporating solar-cycle harmonics directly encodes the slowly varying component, allowing the model to allocate capacity to departures with predictive significance. These channels support representation of onset, ascent, maximum, and decay morphology across extended horizons, thereby reducing dependence on recursive extrapolation for phase structure [133, 130, 134].

This phase-oriented methodology is also applied in the conditional a_p model, which uses solar-cycle harmonics together with daily $F_{10.7}^{\text{adj}}(t)$ and the smoothed background flux $\bar{F}_{10.7}(t)$ to construct the conditional geomagnetic forcing state. Consequently, the forecasting pipeline conditions both radiative and stochastic geomagnetic drivers on explicit low-frequency phase coordinates in addition to short-window driver statistics.

4.3.3 Construction of the Forecast State

The forecasting state integrates measured geophysical drivers with twelve harmonic phase channels. Define the stacked harmonic feature vector as

$$\mathbf{s}_{\text{harm}}(t) = \begin{bmatrix} \mathbf{s}_{\text{ann}}(t)^\top & \mathbf{s}_{\text{cyc}}(t)^\top \end{bmatrix}^\top \quad (4.21)$$

with

$$N_{\text{harm}} = 2N_{\text{ann}} + 2N_{\text{cyc}} = 12 \quad (4.22)$$

The target driver block is written as

$$\mathbf{d}_{\text{tar}}(t) = \begin{bmatrix} F_{10.7}^{\text{obs}}(t) & F_{10.7}^{\text{adj}}(t) & F_{10.7}^{\text{URSI}}(t) & a_p(t) \end{bmatrix}^\top \quad (4.23)$$

while the contextual driver block is written as

$$\mathbf{d}_{\text{ctx}}(t) = \begin{bmatrix} K_p(t) & S_N(t) & D_{\text{st}}(t) & N_{\odot}^{\text{rot}}(t) \end{bmatrix}^\top \quad (4.24)$$

with source-specific terms omitted when a particular archive configuration does not retain them.

The complete forecast-state vector supplied to the forecasting architecture is then

$$\boldsymbol{\xi}_{\text{SW}}(t) = \begin{bmatrix} \mathbf{d}_{\text{tar}}(t)^\top & \mathbf{d}_{\text{ctx}}(t)^\top & \mathbf{s}_{\text{harm}}(t)^\top \end{bmatrix}^\top \quad (4.25)$$

Equivalently, this state is the operational realization of the augmented driver vector $\mathbf{d}_{\text{SW}}^{\text{aug}}(t)$, specialized to the forecast-target and conditioning structure used in this implementation.

This construction is based on physical interpretation. The raw and processed driver channels represent measured or derived geophysical states, while the harmonic channels locate those states within two structured periodic backgrounds. The annual basis regularizes calendar-phase information, and the solar-cycle basis regularizes long-period solar-phase information. Together, these bases provide structured coordinates that help the model contextualize anomalies, storm-time departures, cycle-scale asymmetry, and local distortions in the forcing record [131, 132, 133, 130].

The advantages of this design are most evident at extended forecast horizons. Multi-year prediction performance degrades when the model allocates excessive representational capacity to reconstructing the slow phase from a limited context. Explicit long-period phase representation mitigates this issue, while measured drivers preserve short-period variation. The model, therefore, forecasts the interaction between analytic background coordinates and dynamic departures relevant to future thermospheric forcing.

The resulting state decomposes the forcing record into measured geophysical content and explicit phase geometry. This partition is consistent with downstream propagation objectives. NRLMSIS-compatible forcing depends on future solar and geomagnetic drivers, whose dominant low-frequency organization can be represented analytically even when detailed future realizations remain uncertain [35, 81].

4.4 Hybrid ResNeXt-iTransformer Forecasting Architecture

The forecasting model distinguishes between local temporal feature extraction and cross-channel dependency modeling. Solar radio flux, geomagnetic activity, and harmonic phase coordinates evolve on distinct temporal scales, including solar-rotation modulation, storm-time excursions, annual structure, and the gradual progression of the solar cycle. The architecture addresses this scale heterogeneity by first compressing the daily multivariate record through a multiscale one-dimensional convolutional backbone and then applying inverted self-attention to latent channel representations. This design preserves the forecasting objective of generating future atmospheric drivers for Naval Research Laboratory Mass Spectrometer and Incoherent Scatter Radar model 2.0 (NRLMSIS 2.0), while incorporating the expanded SW-A11 archive and disturbance storm time index D_{st} . The architecture combines aggregated residual transformations from ResNeXt with the variate-token paradigm of iTransformer, thereby supporting long-horizon, multivariate time-series forecasting across heterogeneous space-weather driver channels [121, 135, 136, 122, 137, 123, 125, 124].

4.4.1 Forecast State, Sequence Length, and Targets

The forecast state employs a one-year lookback and a one-year prediction horizon. The implementation indexes the dataset with

$$L_{in} = 365, \quad L_{out} = 365 \quad (4.26)$$

where L_{in} is the input lookback length in daily samples and L_{out} is the forecast length in daily samples. Each supervised training sample q_{seq} maps a daily multivariate input sequence to a daily target sequence,

$$\mathbf{X}_q^{in} = \begin{bmatrix} \boldsymbol{\xi}_{SW}(t_q)^\top \\ \boldsymbol{\xi}_{SW}(t_q + \Delta t_d)^\top \\ \vdots \\ \boldsymbol{\xi}_{SW}(t_q + (L_{in} - 1)\Delta t_d)^\top \end{bmatrix}, \quad (4.27)$$

$$\mathbf{Y}_q^{\text{out}} = \begin{bmatrix} \mathbf{d}_{\text{tar}}(t_q + L_{\text{in}}\Delta t_d)^\top \\ \mathbf{d}_{\text{tar}}(t_q + (L_{\text{in}} + 1)\Delta t_d)^\top \\ \vdots \\ \mathbf{d}_{\text{tar}}(t_q + (L_{\text{in}} + L_{\text{out}} - 1)\Delta t_d)^\top \end{bmatrix} \quad (4.28)$$

where $\Delta t_d = 1$ day, $\xi_{\text{SW}}(t)$ is the forecast-state vector, and $\mathbf{d}_{\text{tar}}(t)$ is the target driver block. Thus,

$$\mathbf{X}_q^{\text{in}} \in \mathbb{R}^{L_{\text{in}} \times N_{\text{feat}}}, \quad \mathbf{Y}_q^{\text{out}} \in \mathbb{R}^{L_{\text{out}} \times N_{\text{tar}}} \quad (4.29)$$

where N_{feat} is the number of input forecast-state channels and N_{tar} is the number of target channels.

This configuration is consistent with the physical and statistical requirements of the forecasting task. A one-year input window captures seasonal phase, solar-rotation-scale modulation, intermediate recurrence, and sufficient local context to position the record within the broader solar-cycle envelope. The one-year output window generates future atmospheric drivers over propagation-relevant timescales and extends the workflow beyond short-term nowcasting [112, 133, 130, 138].

The target block consists of the solar radio flux family together with $a_p(t)$:

$$\mathbf{d}_{\text{tar}}(t) = \begin{bmatrix} F_{10.7}^{\text{obs}}(t) & F_{10.7}^{\text{adj}}(t) & F_{10.7}^{\text{URSI}}(t) & a_p(t) \end{bmatrix}^\top \quad (4.30)$$

so that $N_{\text{tar}} = 4$. In code-level archive terminology, these target columns are

$$\{\text{Obs}, \text{Adj}, \text{URSI-D}, \text{ap}\}$$

The contextual inputs include $K_p(t)$, total sunspot number $S_N(t)$, $D_{\text{st}}(t)$, and the solar-rotation count $N_{\odot}^{\text{rot}}(t)$. The twelve harmonic channels contained in $\mathbf{s}_{\text{harm}}(t)$ augment these measured and derived drivers, resulting in a multivariate state that encodes both empirical forcing and explicit phase information. The data stream therefore incorporates the expanded $\mathcal{D}_{\text{SW}}$ archive, the SW-All records, and the D_{st} channel within the unified daily forcing state [126, 119, 128, 120].

Before model input, each forecast-state channel is normalized using training-set statistics. For feature channel j , let μ_j^{tr} and σ_j^{tr} denote the training-set mean and standard deviation. The normalized forecast state is

$$\tilde{\xi}_{\text{SW},j}(t) = \frac{\xi_{\text{SW},j}(t) - \mu_j^{\text{tr}}}{\sigma_j^{\text{tr}}} \quad (4.31)$$

and the corresponding normalized vector is denoted by $\tilde{\xi}_{\text{SW}}(t)$. This normalization is estimated only from the training partition to avoid leakage of future distributional information into the supervised forecast construction [101, 102].

This approach represents a forcing-state design problem. The input tensor encodes a structured combination of solar radiative state, geomagnetic disturbance, and analytically derived phase coordinates. The network learns departures from organized backgrounds, while explicit phase encoding reduces the need to reconstruct slowly varying phase from finite-window driver histories. This structure is critical for maintaining long-horizon predictive accuracy in multivariate space-weather forcing applications [123, 125, 124, 138].

The learned forecasting map is written as

$$\hat{\mathbf{Y}}_q^{\text{out}} = \mathcal{M}_{\text{RIT}}(\mathbf{X}_q^{\text{in}}) \quad (4.32)$$

where $\mathcal{M}_{\text{RIT}}$ denotes the trained ResNeXt-iTransformer model and $\hat{\mathbf{Y}}_q^{\text{out}}$ is the predicted future target-driver sequence used to assemble $\hat{\mathbf{d}}_{\text{SW}}(t)$.

Figure 4.5 summarizes the full ResNeXt-iTransformer forecasting architecture, including harmonic-augmented solar inputs, multiscale one-dimensional convolutional feature extraction, inverted Transformer encoding, and projection to the solar radio flux forecast sequence.




Figure 4.5: Model Architecture

4.4.2 MultiScale ResNeXt Backbone

The architectural foundation is a one-dimensional MultiScale ResNeXt backbone. ResNeXt extends residual learning by increasing representational capacity through aggregated residual transforma-

tions. In this architecture, cardinality is defined as the number of parallel grouped transformations and serves as a principal design axis alongside depth and width [121]. Let C_{card} denote this cardinality. A generic aggregated residual block can be written schematically as

$$\mathcal{B}_{\text{RNX}}(\mathbf{x}) = \mathbf{x} + \sum_{c=1}^{C_{\text{card}}} \mathcal{T}_c(\mathbf{x}) \quad (4.33)$$

where $\mathcal{T}_c$ denotes the c -th grouped transformation within the residual block.

This configuration produces a multibranch feature extractor that learns diverse local filters without requiring every branch to process the full channel space. Such modularity is well matched to the daily driver archive, which contains heterogeneous local phenomena, including short-period solar modulation, intermittent geomagnetic activity spikes, and channel-specific responses that benefit from differentiated convolutional paths [121, 137].




Figure 4.6: ResNeXt-101 Backbone

Figure 4.6 details the one-dimensional ResNeXt-101 backbone used to transform the normalized daily driver sequence into a compact feature map. The staged bottleneck structure provides progressive temporal abstraction before the reduced representation is passed to the inverted Transformer encoder. The backbone processes the normalized multivariate daily sequence after harmonic augmentation. The paper formulation and accompanying implementation use the ResNeXt-101 stage configuration of $\mathbf{n}_{\text{stage}} = (3, 4, 23, 3)$ with base channel width $W_{\text{base}} = 64$ and cardinality $C_{\text{card}} = 32$. The resulting one-dimensional grouped bottleneck structure applies progressive temporal downsampling and multiscale convolutional filtering [121]. The early layers preserve finer temporal granularity, enabling resolution of short- and mesoscale driver dynamics. Later layers increase the effective receptive field and condense the input sequence into a latent temporal representation.

The ResNeXt feature-extraction map is denoted by Φ_{RNX} . For sample q_{seq} ,

$$\mathbf{H}_{\text{RNX},q} = \Phi_{\text{RNX}}(\mathbf{X}_q^{\text{in}}) \quad (4.34)$$

where $\mathbf{H}_{\text{RNX},q}$ is the latent feature map produced by the convolutional backbone. A 1×1 temporal-channel reducer Φ_{red} then projects the convolutional output into the latent dimension required by the Transformer block:

$$\mathbf{H}_{\text{red},q} = \Phi_{\text{red}} (\mathbf{H}_{\text{RNX},q}), \quad \mathbf{H}_{\text{red},q} \in \mathbb{R}^{L_{\text{red}} \times d_{\text{model}}} \quad (4.35)$$

where L_{red} is the downsampled temporal length and d_{model} is the Transformer latent dimension.

Within the forecasting framework, the convolutional front end performs pre-attentive modeling of local and mesoscale temporal structures. Attention mechanisms capture long-range and cross-variate dependencies effectively after the input has been organized into informative latent representations. The ResNeXt front end, therefore, converts the normalized driver sequence into a compressed feature map whose channels encode multiscale temporal summaries. The downstream inverted-attention block receives a representation that has already been filtered, downsampled, and structured into latent series suitable for multivariate dependency modeling [121, 135, 122, 137].

4.4.3 Inverted Transformer Encoder

The second stage uses an inverted Transformer encoder following the iTransformer framework. In this design, tokenization shifts from timestamp-level aggregation to variate-level embedding. Standard multivariate Transformer formulations often construct each token from all variables observed at a single timestamp. The inverted formulation embeds the temporal history of each variable as a variate token and applies self-attention across these tokens. The resulting attention mechanism directly models multivariate dependence, while the feed-forward modules process each variate-token representation independently [135, 122].

Let the reduced latent sequence produced by the convolutional reducer be $\mathbf{H}_{\text{red},q} \in \mathbb{R}^{L_{\text{red}} \times d_{\text{model}}}$. The inverted embedding map is denoted by Φ_{inv} , with

$$\mathbf{Z}_q^{(0)} = \Phi_{\text{inv}} (\mathbf{H}_{\text{red},q}), \quad \mathbf{Z}_q^{(0)} \in \mathbb{R}^{N_{\text{var}} \times d_{\text{model}}} \quad (4.36)$$

where N_{var} is the number of variate tokens supplied to the encoder after convolutional reduction. Each row of $\mathbf{Z}_q^{(0)}$ represents one latent variate token carrying temporal information extracted by the ResNeXt backbone.

For encoder layer $r \in \{1, \dots, N_{\text{enc}}\}$, the cross-variate attention update is written as

$$\tilde{\mathbf{Z}}_q^{(r)} = \text{LN} \left[\mathbf{Z}_q^{(r-1)} + \text{MHA}_{\text{var}} \left(\mathbf{Z}_q^{(r-1)} \right) \right] \quad (4.37)$$

followed by the variate-wise feed-forward update

$$\mathbf{Z}_q^{(r)} = \text{LN} \left[\tilde{\mathbf{Z}}_q^{(r)} + \text{FFN}_{\text{var}} \left(\tilde{\mathbf{Z}}_q^{(r)} \right) \right] \quad (4.38)$$

Here, $\text{MHA}_{\text{var}}(\cdot)$ denotes multi-head attention (MHA) applied across variate tokens, $\text{FFN}_{\text{var}}(\cdot)$ denotes the shared feed-forward network applied to each variate token, $\text{LN}(\cdot)$ denotes layer normalization (LN), and N_{enc} is the number of stacked encoder layers [135, 139, 122].




Figure 4.7: iTransformer Encoder

Figure 4.7 illustrates the inverted encoder stage, in which temporally informed latent variates are embedded as tokens and processed through cross-variate attention. This structure is aligned with the driver archive because prediction depends on conditional relationships among solar flux, geomagnetic indices, sunspot number, rotation count, and harmonic phase coordinates. The predictive task depends on temporal coherence within each channel and conditional dependence among channels. Variate-token encoding allows the encoder to learn cross-channel covariance among latent temporal features, thereby supporting direct modeling of dependencies across solar, geomagnetic, and harmonic inputs [122].

The implementation rearranges the convolutionally reduced feature map into the format required for inverted embedding, applies training-set normalization inherited from the forecast-state construction, and embeds the resulting latent series into dimension d_{model} . The stacked encoder layers then apply cross-variate MHA and variate-wise feed-forward transformations. The full pipeline consists of convolutional compression, variate-token embedding, cross-channel dependency modeling, and projection to the forecast sequence [135, 139, 122].

The encoder retains temporal context because each variate token contains temporal features

extracted by the convolutional backbone. The encoder, therefore, examines cross-channel relationships among temporally informed representations and identifies coordinated, conditional, or divergent channel behavior under varying solar and geomagnetic forcing states. In this setting, attention is applied at the level of multivariate interaction after local temporal structure has already been organized by the ResNeXt front end [121, 135, 122].

4.4.4 Output Head and Recursive Rollout

The output stage projects the final encoder representation into the daily forecast horizon. Let the final variate-token representation be $\mathbf{Z}_q^{(N_{\text{enc}})}$. The output projection map Φ_{out} generates the predicted target tensor

$$\widehat{\mathbf{Y}}_q^{\text{out}} = \Phi_{\text{out}} \left(\mathbf{Z}_q^{(N_{\text{enc}})} \right), \quad \widehat{\mathbf{Y}}_q^{\text{out}} \in \mathbb{R}^{L_{\text{out}} \times N_{\text{tar}}} \quad (4.39)$$

The forecasted target channels correspond to the solar radio flux family and $a_p(t)$,

$$\widehat{\mathbf{Y}}_q^{\text{out}} = \left[\widehat{F}_{10.7}^{\text{obs}}, \widehat{F}_{10.7}^{\text{adj}}, \widehat{F}_{10.7}^{\text{URSI}}, \widehat{a}_p \right]_q \quad (4.40)$$

sampled over the $L_{\text{out}} = 365$ future daily epochs. This decoder maps variate-level representations into a multivariate daily forecast spanning the annual prediction interval [136, 122, 138].

Let $r_{\text{roll}} \in \{0, \dots, R_{\text{roll}} - 1\}$ denote the rollout-block index, and let $\mathcal{H}_{\text{SW}}^{(r_{\text{roll}})}$ denote the historical-plus-predicted driver history available at the beginning of rollout block r_{roll} . The forecast block is generated by

$$\widehat{\mathbf{D}}_{\text{tar}}^{(r_{\text{roll}})} = \mathcal{M}_{\text{RIT}} \left[\mathcal{H}_{\text{SW}}^{(r_{\text{roll}})} \right] \quad (4.41)$$

where $\widehat{\mathbf{D}}_{\text{tar}}^{(r_{\text{roll}})} \in \mathbb{R}^{L_{\text{out}} \times N_{\text{tar}}}$ is the predicted target-driver block. The rollout-state assembly map $\mathcal{A}_{\text{roll}}$ appends the predicted target block, refreshes deterministic harmonic features, and reconstructs contextual inputs required for the next model application:

$$\mathcal{H}_{\text{SW}}^{(r_{\text{roll}}+1)} = \mathcal{A}_{\text{roll}} \left(\mathcal{H}_{\text{SW}}^{(r_{\text{roll}})}, \widehat{\mathbf{D}}_{\text{tar}}^{(r_{\text{roll}})} \right) \quad (4.42)$$

After R_{roll} recursive blocks, the assembled future driver trajectory is denoted by

$$\widehat{\mathbf{d}}_{\text{SW}}^{(R_{\text{roll}})}(t) \quad (4.43)$$

which supplies future $F_{10.7}$ forcing, conditional a_p generation, and subsequent NRLMSIS-compatible model input.

Recursive rollout introduces error accumulation because predicted values become part of the

subsequent input history. The architecture addresses this issue through three stabilizing structures: convolutional regularization of local temporal patterns, explicit harmonic phase coordinates, and cross-channel coherence imposed through the variate-token encoder. The resulting model combines analytic phase scaffolding with a learned multivariate transition rule, providing a structured mechanism for multi-year driver generation beyond single-block prediction [140, 121, 135, 122, 137, 123, 125, 124].

4.5 Training Protocol and Forecast Validation

The forecasting model is trained and evaluated using a validation protocol designed for prospective application, with the objective of avoiding retrospective curve fitting. The primary objective is stable, low-error performance across repeated causal splits that approximate operational deployment. This analysis focuses on three technical requirements: strict chronological partitioning, stable optimization under the specified training regime, and retained forecast skill across validation windows that sample distinct portions of the historical solar archive [141, 142].

4.5.1 Temporal Split and Cross-Validation

The validation protocol uses a five-fold expanding-window time-series cross-validation design. Let

$$N_{\text{cv}} = 5 \tag{4.44}$$

denote the number of temporal folds, and let $r_{\text{cv}} \in \{1, \dots, N_{\text{cv}}\}$ denote the fold index. Each fold is defined by a chronological cutoff epoch $t_{\text{cut}}^{(r_{\text{cv}})}$. The supervised sequence-index sets for the corresponding training and validation segments are denoted by

$$\mathcal{Q}_{\text{train}}^{(r_{\text{cv}})} \quad \text{and} \quad \mathcal{Q}_{\text{val}}^{(r_{\text{cv}})} \tag{4.45}$$

The training set contains all admissible supervised sequences available before the fold cutoff, while the validation set contains the contiguous forecast interval that immediately follows. This construction preserves causal ordering and prevents future observations from entering the model-selection process [141, 142].

For each fold, the trained forecasting map is estimated using only

$$\{\mathbf{X}_q^{\text{in}}, \mathbf{Y}_q^{\text{out}}\}_{q \in \mathcal{Q}_{\text{train}}^{(r_{\text{cv}})}} \tag{4.46}$$

and evaluated on

$$\{\mathbf{X}_q^{\text{in}}, \mathbf{Y}_q^{\text{out}}\}_{q \in \mathcal{Q}_{\text{val}}^{(r_{\text{cv}})}} \quad (4.47)$$

Thus, each validation forecast reflects the operational condition in which the model is trained on a past archive and then applied to a subsequent future interval.

This splitting method is appropriate for time-ordered data whose statistical properties are influenced by solar-cycle phase, solar-rotation modulation, and geomagnetic disturbance structure. Random partitioning mixes quiescent and active intervals across training and validation subsets, which can inflate skill metrics and obscure sensitivity to transitions, volatility, and regime-dependent forcing. The expanding-window protocol tests whether a model trained through $t_{\text{cut}}^{(r_{\text{cv}})}$ can forecast the following interval using only information available at that epoch [141, 142].

The implementation maintains the one-year input window and one-year forecast horizon established previously, with $L_{\text{in}} = 365$ and $L_{\text{out}} = 365$. Each fold, therefore, evaluates blockwise forward prediction with consistent forecast-state construction, normalization, harmonic augmentation, and recursive rollout logic. The resulting validation scores serve as operational tests of model utility across temporally segregated segments of the solar archive [133, 130, 141].

4.5.2 Optimization Configuration

Let θ_{RIT} denote the trainable parameter vector of the ResNeXt-iTransformer forecasting map $\mathcal{M}_{\text{RIT}}$. The model is trained by minimizing the mean squared error (MSE) between predicted and observed target-driver sequences. For a supervised sequence q , forecast day index $\ell \in \{1, \dots, L_{\text{out}}\}$, and target channel $j \in \{1, \dots, N_{\text{tar}}\}$, the scalar prediction is denoted by $\hat{Y}_{q,\ell,j}^{\text{out}}$, and the corresponding observed target is denoted by $Y_{q,\ell,j}^{\text{out}}$. The fold-level training objective is

$$\mathcal{L}_{\text{MSE}}(\theta_{\text{RIT}}; \mathcal{Q}_{\text{train}}^{(r_{\text{cv}})}) = \frac{1}{|\mathcal{Q}_{\text{train}}^{(r_{\text{cv}})}| L_{\text{out}} N_{\text{tar}}} \sum_{q \in \mathcal{Q}_{\text{train}}^{(r_{\text{cv}})}} \sum_{\ell=1}^{L_{\text{out}}} \sum_{j=1}^{N_{\text{tar}}} \left(\hat{Y}_{q,\ell,j}^{\text{out}} - Y_{q,\ell,j}^{\text{out}} \right)^2 \quad (4.48)$$

The training procedure uses AdamW optimization, cosine-annealed learning rates, gradient clipping, dropout, automatic mixed precision (AMP), and early stopping [143, 144, 145, 146, 147, 148]. The optimizer and scheduler configuration are

$$\eta_0 = 10^{-4}, \quad \lambda_{\text{wd}} = 10^{-2}, \quad \eta_{\text{min}} = 10^{-8}, \quad N_{\text{ep}} = 20 \quad (4.49)$$

with batch size $B_{\text{batch}} = 128$, dropout probability $p_{\text{drop}} = 0.25$, and gradient-norm clipping threshold $G_{\text{max}} = 1.0$. The optimizer update rule is denoted by $\mathcal{O}_{\text{AdamW}}$, while the cosine-annealing

learning-rate schedule is denoted by $\mathcal{S}_\eta$. The clipping operation constrains the total gradient norm according to

$$\|\nabla_{\theta_{\text{RIT}}} \mathcal{L}_{\text{MSE}}\|_2 \leq G_{\text{max}} \quad (4.50)$$

after clipping is applied.




Figure 4.8: Training Loss

The loss history in Figure 4.8 indicates stable convergence. Training MSE decreases from approximately 0.61 at epoch one to 0.22 by epoch 20, with the most pronounced reduction occurring during the first five epochs, followed by a slower monotonic decline. This profile indicates that dominant solar-cycle, annual, and intermediate-scale patterns are captured early, while later epochs refine residual structure in the target-driver sequence. The absence of oscillatory growth or late divergence indicates that the selected learning rate, weight decay, dropout, and gradient-clipping configuration is well calibrated for this dataset.

The convergence profile is consistent with the architectural decomposition. The convolutional layers extract local and mesoscale features from the normalized, harmonically augmented input state, reducing the representational load placed on the attention layers during early training. Subsequent epochs refine cross-channel dependencies and residual temporal structure through the inverted Transformer encoder. The loss curve therefore supports the separation between multiscale temporal feature extraction and cross-variate dependency modeling [121, 135, 122].

4.5.3 Fold-Level Forecast

Forecast skill remains consistent across the five temporal folds. For validation fold r_{cv} , the mean absolute error (MAE) is defined as

$$\text{MAE}^{(r_{cv})} = \frac{1}{|Q_{\text{val}}^{(r_{cv})}| L_{\text{out}} N_{\text{tar}}} \sum_{q \in Q_{\text{val}}^{(r_{cv})}} \sum_{\ell=1}^{L_{\text{out}}} \sum_{j=1}^{N_{\text{tar}}} |\hat{Y}_{q,\ell,j}^{\text{out}} - Y_{q,\ell,j}^{\text{out}}| \quad (4.51)$$

The corresponding root mean squared error (RMSE) is

$$\text{RMSE}^{(r_{cv})} = \left[\frac{1}{|Q_{\text{val}}^{(r_{cv})}| L_{\text{out}} N_{\text{tar}}} \sum_{q \in Q_{\text{val}}^{(r_{cv})}} \sum_{\ell=1}^{L_{\text{out}}} \sum_{j=1}^{N_{\text{tar}}} \left(\hat{Y}_{q,\ell,j}^{\text{out}} - Y_{q,\ell,j}^{\text{out}} \right)^2 \right]^{1/2} \quad (4.52)$$

The retained validation diagnostic collection is denoted by

$$\mathcal{R}_{\text{val}} = \left\{ \text{MAE}^{(r_{cv})}, \text{RMSE}^{(r_{cv})} \right\}_{r_{cv}=1}^{N_{cv}} \quad (4.53)$$

MAE and RMSE provide complementary absolute-error summaries: MAE reports the mean magnitude of residuals, while RMSE weights larger residual excursions more strongly through the squared-error operation [149].

Fold diagnostics indicate MAE values between 0.64 and 0.78, and RMSE values between 0.94 and 1.09. These values exhibit limited fold-to-fold variation, with no validation segment showing catastrophic performance loss. The model therefore maintains stable forecast skill across temporally distinct validation intervals.




Figure 4.9: Cross-Validation MAE by Fold

Figure 4.9 shows the fold-wise MAE distribution, indicating bounded variation across the expanding-window validation splits and confirming that absolute forecast error remains stable across temporally distinct evaluation intervals.




Figure 4.10: Cross-Validation RMSE by Fold

The corresponding RMSE behavior in Figure 4.10 shows the same fold-level stability while emphasizing the contribution of larger residual excursions. The absence of an anomalously large

fold indicates that the model does not experience validation-window collapse under the causal cross-validation protocol.

Lower errors are associated with validation windows that exhibit smoother forcing histories, while higher errors occur in windows characterized by irregular structure and cycle-phase transitions. This behavior reflects the state dependence of forecast skill, with increased difficulty corresponding to less regular and less compressible solar and geomagnetic intervals [133, 130, 123, 125, 124].

This validation protocol is sufficient for the current objective because the model demonstrates reproducible convergence and bounded error under causal temporal splits. Broader performance characterization would require baseline comparisons, architectural ablations, probabilistic calibration, and explicit uncertainty quantification [149, 140, 142]. For the downstream thermospheric forcing pipeline, the central requirement is that the model supports reproducible training, causal validation, and forward application without fold-specific instability or regime-local failure.

4.6 Residual Structure

Fold-averaged error metrics summarize aggregate accuracy without resolving the structure of forecast error. Forcing models intended for thermospheric specification require residual diagnostics because downstream density error depends on whether the unresolved component is diffuse, biased, temporally correlated, or spectrally organized. This analysis evaluates the marginal distribution, temporal evolution across the prediction window, and retained periodic content. Collectively, these diagnostics determine whether residual uncertainty can be treated as weak stochastic variability or should be propagated as a structured error component in subsequent forcing [149, 150, 151, 152].

Let the forecast-residual tensor for supervised sequence q be defined using the prediction-minus-observation convention:

$$\mathbf{E}_q^{\text{out}} = \widehat{\mathbf{Y}}_q^{\text{out}} - \mathbf{Y}_q^{\text{out}} \quad (4.54)$$

The scalar residual for sequence q , forecast-lead index ℓ , and target channel j is therefore

$$E_{q,\ell,j}^{\text{out}} = \widehat{Y}_{q,\ell,j}^{\text{out}} - Y_{q,\ell,j}^{\text{out}} \quad (4.55)$$

Under this convention, negative residuals indicate underprediction and positive residuals indicate overprediction. The retained residual-diagnostic collection is denoted by $\mathcal{R}_{\text{res}}$, including marginal, horizon-dependent, and spectral summaries.

4.6.1 Residual Distribution

The residual histogram displays a pronounced central peak near zero, with substantially lower probability assigned to large deviations. Most forecast errors remain within a moderate-deviation regime, indicating that errors arise primarily from localized mismatches in difficult intervals rather than persistent large-amplitude failures. This behavior is advantageous for multi-year forcing models because the forecast captures the slowly varying solar background and typical departures, leaving the unresolved component concentrated in the most irregular archival intervals [149].

The residual distribution exhibits asymmetry, with a longer negative tail. Given the prediction-minus-observation sign convention, the largest negative residuals correspond to underestimation of active behavior or incomplete recovery of irregular excursion amplitudes. This asymmetry is consistent with the intrinsic difficulty of forecasting active intervals shaped by rotational recurrence, evolving active regions, and transient enhancements that perturb the low-frequency cycle structure [112, 133, 130]. Smooth solar minima and broad cycle envelopes are represented with lower residual dispersion.




Figure 4.11: Histogram of Prediction Residuals

Figure 4.11 summarizes the marginal residual distribution and shows a dominant central concentration near zero with asymmetric tail behavior. The pronounced central concentration near zero indicates that the hybrid model reconstructs the dominant forcing morphology throughout most of the archive. Residual error is concentrated in the distribution tails, which suggests effective learning of the principal solar modulation and a limited set of activity-dependent mismatches

remaining unresolved.

4.6.2 Horizon Residual Behavior

Time-resolved residual diagnostics demonstrate that forecast bias varies with lead time across the annual prediction window. The lead-dependent mean residual is written as

$$\bar{E}_{\ell,j}^{\text{out}} = \frac{1}{N_{\text{res}}} \sum_{q \in Q_{\text{res}}} E_{q,\ell,j}^{\text{out}} \quad (4.56)$$

where Q_{res} is the residual-evaluation sequence set and $N_{\text{res}} = |Q_{\text{res}}|$. When target channels are aggregated, the channel-averaged mean residual is denoted by

$$\bar{E}_{\ell}^{\text{out}} = \frac{1}{N_{\text{tar}}} \sum_{j=1}^{N_{\text{tar}}} \bar{E}_{\ell,j}^{\text{out}} \quad (4.57)$$

The mean residual begins negative, intensifies in the early horizon, and then gradually increases toward near-zero or slightly positive values at longer lead times. This pattern indicates systematic lead-dependent bias with bounded temporal evolution.




Figure 4.12: Mean Forecast Residual over the Prediction Horizon

Figure 4.12 makes the lead-dependent residual structure explicit by showing a smooth transition from early-horizon underprediction toward near-zero or slightly positive mean residual at later lead times. The structure of the residual has direct implications for long-horizon recursive and blockwise forecasting. Recursive forecasts can exhibit rapid error amplification when forecasted values feed

subsequent inputs and can exhibit stable bias drift when residuals remain bounded and vary smoothly with forecast lead [140]. The diagnostics presented here correspond to the bias-drift regime, as residuals evolve in a controlled and non-divergent manner across the annual horizon.

The reversal of residual sign across the forecast horizon indicates that error allocation varies dynamically with lead time. Initial underprediction is most prominent and then decreases as the horizon extends. This pattern does not imply uniformly superior late-horizon prediction. It demonstrates a smooth, bounded, and diagnosable bias structure. This organization supports lead-time-specific uncertainty modeling for downstream density forcing, which is preferable to a single undifferentiated residual envelope.

4.6.3 Spectral Content

Spectral diagnostics provide a physically informative perspective on the forecasting model. The Lomb–Scargle (LS) periodogram of a driver time series is denoted by $P_{\text{LS}}^{\text{drv}}(f)$, where f is temporal frequency. The corresponding residual periodogram is denoted by $P_{\text{LS}}^{\text{res}}(f)$. The LS periodogram is a least-squares spectral estimator originally developed for unevenly sampled data and widely used for detecting periodic content in astronomical and geophysical time series [151, 150, 152].

The power spectrum of the solar radio flux driver is dominated by low-frequency components, with LS analysis identifying a prominent long-period peak associated with the solar-cycle envelope [150, 133]. This spectral signature indicates that variance is primarily governed by slow Schwabe-scale modulation, with higher-frequency features superimposed. The forecast preserves this low-frequency organization, maintaining the thermospheric relevance of the predicted driver sequence [112, 35, 36].

The concentration of low-frequency power in the driver spectrum does not fully characterize the forecast error. Residual periodograms display persistent periodicity, especially near the solar-rotation band B_{rot} and at intermediate periods. The residual LS spectrum therefore reveals organized power at specific frequencies, indicating that forecast errors retain physically interpretable temporal structure after the dominant cycle is removed [150, 152, 112, 133].

Residual periodicity has direct consequences for downstream uncertainty modeling. A residual field with negligible spectral organization could be represented as weak perturbations around a deterministic forcing trajectory. The diagnostics indicate an unresolved component with periodic content associated with solar modulation, rotational recurrence, and intermediate-scale structure. Future thermospheric forcing therefore benefits from uncertainty representations that preserve this spectral organization, either through residual-informed stochastic models or ensemble forcing trajectories [35, 36].

The model captures the dominant low-frequency solar-cycle structure required for density specification over multi-month and multi-year periods, while the residuals isolate remaining organized forcing components. This separation between primary learned structure and residual periodic content provides a physically coherent foundation for uncertainty-aware NRLMSIS forcing and improves upon purely deterministic atmospheric-driver construction.




Figure 4.13: Lomb–Scargle Diagnostics of 10.7 cm Solar Radio Flux

Figure 4.13 shows the Lomb–Scargle power spectrum of the solar radio flux driver, with dominant low-frequency power associated with the solar-cycle envelope. This confirms that the primary driver variance is organized by long-period solar modulation relevant to thermospheric forcing.




Figure 4.14: Lomb–Scargle Diagnostics of 10.7 cm Solar Radio Flux Residuals

The residual spectrum in Figure 4.14 shows that forecast error retains organized periodic content at solar-rotation and intermediate timescales after the dominant cycle-scale structure has been modeled. This residual organization motivates downstream uncertainty representations that preserve temporal correlation and spectral structure.

4.7 Multi-Year Forecasts of Solar and Geomagnetic Drivers

The performance of the forecasting model depends on its ability to generate physically coherent driver histories over extended multi-year intervals. The principal criteria are preservation of the low-frequency solar-cycle envelope, internal consistency among the radio-flux proxies, and generation of a statistically plausible geomagnetic forcing process suitable for empirical thermospheric modeling [112, 133, 130, 123]. This analysis evaluates recursive forecasts for the observed, adjusted, and URSI-D 10.7 cm solar radio flux channels, followed by the conditional stochastic projection of a_p .

Let the future daily forecast epoch set be denoted by $\mathcal{T}_{\text{fcst}}$, and let the multi-year forecast horizon be

$$H_{\text{fcst}}^{\text{MY}} = 22 \text{ yr} \quad (4.58)$$

The recursively assembled forecast driver vector over this interval is written as

$$\widehat{\mathbf{d}}_{\text{SW}}^{\text{MY}}(t), \quad t \in \mathcal{T}_{\text{fcst}} \quad (4.59)$$

where the solar radio flux subset is

$$\widehat{\mathbf{S}}_{\text{F10.7}} = \left\{ \widehat{F}_{10.7}^{\text{obs}}(t), \widehat{F}_{10.7}^{\text{adj}}(t), \widehat{F}_{10.7}^{\text{URSI}}(t) \right\}_{t \in \mathcal{T}_{\text{fcst}}} \quad (4.60)$$

These forecasted driver histories form the radiative component of the future thermospheric forcing pipeline, while the geomagnetic component is supplied through the conditional stochastic a_p model.

4.7.1 Forecast of Observed Solar Radio Flux

The recursive observed-flux forecast preserves the historical cycle morphology over the 22-year horizon. The forecast maintains alternating minima and maxima, including a post-historical low-activity interval followed by two major cycle ascents and decays. This behavior is essential for forcing generation because thermospheric density responds to gradual solar-activity modulation across monthly, annual, and solar-cycle timescales [112, 133, 130, 123, 125, 124].




Figure 4.15: Observed 10.7 cm Solar Radio Flux Forecast

Figure 4.15 places the observed-channel forecast against the historical record, showing that the recursive projection retains the broad solar-cycle envelope while avoiding direct replication of isolated historical transient spikes. The projected observed solar radio flux amplitude distribution remains physically plausible relative to the historical archive. Minima approach quiet-Sun levels, while forecast maxima reach elevated forcing states without reproducing isolated extreme transient outliers. This behavior is appropriate for long-horizon forcing construction because the forecast maintains the cycle envelope and characteristic variability, while intermittent transient departures remain assigned to residual uncertainty and scenario generation.

The observed solar radio flux continuation preserves direct measurement-space interpretation and permits transparent comparison between historical and forecast records. Its use in atmospheric coupling is secondary, however, because the observed channel retains Sun-Earth distance effects that are removed in the adjusted channel [114].

4.7.2 Forecast of Adjusted Solar Radio Flux

The adjusted-flux forecast functions as the primary radiative forcing channel for downstream modeling. The adjustment corrects the observed measurement to a standard Sun-Earth distance, producing a solar-activity proxy more directly suited to empirical upper-atmosphere forcing [114, 34, 35]. Within this pipeline, $\widehat{F}_{10.7}^{\text{adj}}(t)$ also informs the conditional a_p model, thereby coupling the future radiative state to the geomagnetic forcing realization.




Figure 4.16: Adjusted 10.7 cm Solar Radio Flux Forecast

Figure 4.16 shows the adjusted-channel continuation used as the principal radiative forcing input for downstream NRLMSIS-compatible density generation. The forecasted adjusted series retains the broad cycle organization of the observed solar radio flux continuation in a form better suited to forcing applications. The initial descent into low activity is followed by two major cycle ascents, with maxima remaining below isolated historical spikes while remaining distinct from quiet-Sun baselines. The forecast remains dynamically structured throughout the multi-year interval, avoiding flat climatology and preserving interpretable amplitude variation across minima, ascent phases, maxima, and decays [123, 125, 124, 140].

This behavior is critical because thermospheric calculations depend on accurate cycle phase and forcing level more strongly than exact reproduction of transient solar radio flux enhancements. The adjusted channel provides the preferred radiative input for conditional geomagnetic modeling and subsequent NRLMSIS-compatible forcing [114, 35, 81].

4.7.3 Forecast of URSI-D Solar Radio Flux

The URSI-D continuation functions as a consistency check within the solar radio flux family. As a standardized transformation of the adjusted flux, its forecast should preserve the underlying solar-cycle progression while maintaining the expected scale and morphology of the URSI-D representation. The projected series satisfies these requirements: minima, rising phases, maxima, and decays remain synchronized with the cycle structure of the observed and adjusted forecasts,

and the absolute amplitudes align with URSI-D scaling [114].




Figure 4.17: URSI-D 10.7 cm Solar Radio Flux Forecast

Figure 4.17 demonstrates that the URSI-D forecast remains phase-aligned with the observed and adjusted solar radio flux forecasts while preserving its channel-specific scaling. This agreement tests whether the multichannel model captures a shared low-frequency solar state. Divergence in timing or morphology among the three related channels would indicate poor recovery of the common solar driver. The observed alignment demonstrates that the model preserves a unified cycle-scale structure across distinct channel definitions [123, 125, 124].

The URSI-D series, therefore, serves as structural validation for multichannel forecasting. The model captures the common forcing geometry within the radio-flux family while preserving channel-specific scaling. This behavior is necessary for coherent future forcing because atmospheric-driver construction depends on physically consistent multivariate forecasts.

4.7.4 Conditional Stochastic Forecast of a_p

The geomagnetic component of the forcing pipeline is designed to remain statistically and climatologically consistent with the forecasted solar state. The objective is a stochastic forcing sequence anchored to the evolving solar background, with daily variability and disturbed intervals represented as realizations around a conditional mean. In the current implementation, the conditional mean of a_p is parameterized using the forecasted adjusted flux $\hat{F}_{10.7}^{\text{adj}}(t)$, its 81-day centered mean $\bar{F}_{10.7}^{81c}(t)$,

and solar-cycle harmonic channels [34, 35, 81].

The conditional mean is represented schematically as

$$\mu_{a_p|SW}(t) = \mathcal{M}_{a_p} \left(\hat{F}_{10.7}^{\text{adj}}(t), \bar{F}_{10.7}^{81c}(t), \mathbf{s}_{\text{cyc}}(t) \right) \quad (4.61)$$

where $\mathcal{M}_{a_p}$ denotes the conditional geomagnetic-driver model. The residual component is modeled using an autoregressive (AR) process,

$$\epsilon_{a_p}(t) = \phi_{a_p} \epsilon_{a_p}(t - \Delta t_d) + \sigma_{a_p} \zeta_{a_p}(t) \quad (4.62)$$

where ϕ_{a_p} is the autoregressive coefficient, σ_{a_p} is the innovation scale, and $\zeta_{a_p}(t)$ is a stochastic innovation [153]. The sampled geomagnetic realization is then written as

$$\hat{a}_p(t) = \mu_{a_p|SW}(t) + \epsilon_{a_p}(t) \quad (4.63)$$

with clipping or admissibility enforcement applied during implementation to preserve physically valid geomagnetic-index ranges [115, 116, 117, 127].




Figure 4.18: Conditional Stochastic a_p Forecast

Figure 4.18 shows the conditional stochastic a_p forecast, including the historical geomagnetic activity record, the solar-state-conditioned mean trajectory, and a sampled realization about that mean. The figure demonstrates that the future geomagnetic forcing sequence retains both low-

frequency modulation tied to the projected solar state and daily stochastic variability required for disturbance-aware thermospheric forcing.

This methodology separates geomagnetic forcing into low-frequency conditional structure and stochastic residual variability. The conditional mean imparts solar-cycle modulation to the forecasted a_p , ensuring sensitivity to the evolving solar forcing state. The autoregressive residual introduces daily variability and disturbed intervals around this mean. The resulting forecast bridges deterministic solar-state conditioning and stochastic geomagnetic scenario generation [115, 116, 117, 127, 153].

This distinction is fundamental for downstream atmospheric modeling. NRLMSIS-compatible forcing requires a geomagnetic driver that remains statistically consistent with the solar cycle and retains realistic variability across extended future intervals [35, 81]. The a_p model fulfills these requirements: the mean forecast establishes the cycle-scale geomagnetic envelope, while sampled realizations represent the natural variability of geomagnetic disturbances. This driver is therefore appropriate for uncertainty-aware density forcing and extended scenario generation.

The updated forecast demonstrates this behavior. The conditional mean follows both quiescent and active solar phases, while sampled realizations span the expected disturbance range. The resulting geomagnetic driver can be integrated directly into the atmospheric-forcing pipeline as a stochastic history that remains coherent with the projected solar state.

4.8 NRLMSIS 2.0 Coupling

Forecasted solar and geomagnetic drivers become operationally valuable when integrated with an atmospheric model that produces density fields at the spatial and temporal resolutions required by the propagator. This integration is performed using the official Fortran implementation of Naval Research Laboratory Mass Spectrometer and Incoherent Scatter Radar model 2.0 (NRLMSIS 2.0), distributed by the Community Coordinated Modeling Center (CCMC) and driven by forecast-conditioned solar and geomagnetic parameters [81, 35]. The coupling procedure generates a time-indexed sequence of forecast-conditioned density fields, which can be reused across orbit simulations, stochastic-driver scenarios, shell-level sensitivity studies, and prospective uncertainty analyses.

4.8.1 Mapping Forecast Drivers

NRLMSIS 2.0 is an empirical neutral-atmosphere model that provides temperature, constituent number densities, and total mass density as functions of geophysical state, including location,

day of year, time of day, solar activity, and geomagnetic activity [35, 81]. The forecasted driver histories are interfaced with the model after conversion into the MSIS input convention. The principal coupling requirements are variable translation, temporal synchronization, location-aware grid evaluation, and consistent forcing maintenance throughout the forecast horizon.

Let the MSIS input record at epoch t and geodetic state $(\varphi_g, \lambda_g, h_g)$ be denoted by

$$\mathbf{u}_{\text{MSIS}}(t, \varphi_g, \lambda_g, h_g) \quad (4.64)$$

This record contains the calendar date, day of year, Universal Time (UT) seconds, geodetic latitude, geodetic longitude, geodetic altitude, local apparent solar time, an averaged solar-flux input, a daily solar-flux input, and a geomagnetic-history vector. In the implementation, the code-level fields `sfluxavg`, `sflux`, and `ap(1:7)` are mapped from forecast-conditioned solar and geomagnetic histories [34, 35, 81].

The solar flux inputs are represented by

$$\bar{F}_{10.7}^{\text{MSIS}}(t) \quad \text{and} \quad F_{10.7}^{\text{MSIS}}(t) \quad (4.65)$$

where $\bar{F}_{10.7}^{\text{MSIS}}(t)$ is implemented from the centered 81-day adjusted-flux record and $F_{10.7}^{\text{MSIS}}(t)$ is implemented from the previous-day adjusted-flux value. This construction preserves the distinction between the slowly varying solar background and the daily solar forcing level [34, 35, 114]. The geomagnetic input is represented by the seven-component history vector

$$\mathbf{a}_{p,7}(t) = \begin{bmatrix} A_p^{\text{day}}(t) & a_p^{0\text{h}}(t) & a_p^{-3\text{h}}(t) & a_p^{-6\text{h}}(t) & a_p^{-9\text{h}}(t) & \bar{a}_p^{-12:-33\text{h}}(t) & \bar{a}_p^{-36:-57\text{h}}(t) \end{bmatrix}^T \quad (4.66)$$

This vector contains the daily Ap term, the current 3-hour geomagnetic index, three recent 3-hour lag terms, and two lagged multi-bin averages. The stochastic geomagnetic forecast $\hat{a}_p(t)$ is therefore converted into an MSIS-compatible history vector before density generation [115, 116, 117, 127, 35].

The interface clarifies the operational role of the stochastic geomagnetic model. The atmospheric model requires large-scale solar forcing and geomagnetic-disturbance histories, while the forecast system is designed to produce statistically consistent disturbance scenarios instead of deterministic storm-event reconstructions. The conditional a_p framework provides a disturbance-aware input that can be mapped into $\mathbf{a}_{p,7}(t)$, thereby supporting multi-year empirical thermosphere specification.

The implementation uses the official Fortran NRLMSIS 2.0 code distributed by CCMC,

maintaining the forcing description at the model-interface level [81]. This approach provides detailed control over the spatial grid, output cadence, binary storage, and integration with the orbit-simulation workflow. The resulting system operates as an atmospheric precomputation pipeline that separates density-field generation from the satellite propagation loop.

The preprocessing sequence compiles a driver file from forecast-conditioned solar and geomagnetic records. This file, together with the model parameter set, is supplied to the Fortran grid generator, which produces raw density grids. The forecast-conditioned density epoch set is written as

$$\mathcal{T}_\rho^{\text{fcst}} = \{t_\ell^\rho\}_{\ell=1}^{N_\rho}, \quad \Delta t_\rho = 3 \text{ h} \quad (4.67)$$

The grid uses World Geodetic System 1984 (WGS84)-consistent geodetic latitude and altitude [82]. The node sets are

$$\mathcal{L}_\varphi = \{-90^\circ, -85^\circ, \dots, 90^\circ\}, \quad \mathcal{L}_\lambda = \{0^\circ, 5^\circ, \dots, 355^\circ\} \quad (4.68)$$

and

$$\mathcal{L}_h = \{115, 120, \dots, 600\} \text{ km} \quad (4.69)$$

Thus,

$$\Delta\varphi_g = 5^\circ, \quad \Delta\lambda_g = 5^\circ, \quad \Delta h_g = 5 \text{ km} \quad (4.70)$$

with altitude bounds

$$h_{g,\min}^\rho = 115 \text{ km}, \quad h_{g,\max}^\rho = 600 \text{ km} \quad (4.71)$$

For atmospheric scenario r_{scn} , the forecast-conditioned density snapshot is

$$\mathcal{G}_{\rho,\text{fc}}^{(\ell, r_{\text{scn}})} = \mathcal{M}_{\text{NRLMSIS}} \left[\mathbf{u}_{\text{MSIS}} \left(t_\ell^\rho, \varphi_g, \lambda_g, h_g \right) \right] \quad (4.72)$$

evaluated for all grid nodes in $\mathcal{L}_\varphi \times \mathcal{L}_\lambda \times \mathcal{L}_h$. The output is total mass density ρ in kilograms per cubic meter, stored as float32 binary fields [35, 81]. Each temporal slice is archived as a separate binary file with the form

`rho_YYYYMMDD_utHH_f32.bin`

This file-based architecture decouples atmospheric generation from orbit propagation. A forecast-conditioned thermospheric state can be precomputed for a specified driver scenario and reused across multiple satellites, orbital shells, and sensitivity analyses without redundant NRLMSIS calls inside the propagator loop.

4.8.2 Density Products for Future Satellite Propagation

Atmospheric products provided to the propagator are structured as time-indexed density grids, which provide substantially greater physical specificity than scalar daily density summaries. This multidimensional parameterization is required for large LEO satellite constellations because atmospheric drag depends on geodetic latitude, geodetic longitude, altitude, universal time, local solar time, and forecasted space-weather conditions [34, 35, 12, 36, 1]. Maintaining spatial structure in the NRLMSIS coupling output supports trajectory integration, shell-specific sensitivity analysis, and inverse-problem formulations that depend on physically consistent density histories.

For scenario r_{scn} , the time-ordered density product is denoted by

$$\mathcal{G}_{\rho, \text{seq}}^{(r_{\text{scn}})} = \left\{ \mathcal{G}_{\rho, \text{fc}}^{(\ell, r_{\text{scn}})} \right\}_{t_{\ell}^{\rho} \in \mathcal{T}_{\rho}^{\text{fcst}}} \quad (4.73)$$

When multiple deterministic or stochastic driver scenarios are retained, the atmospheric ensemble is indexed by

$$r_{\text{scn}} \in \{1, \dots, N_{\text{scn}}\} \quad (4.74)$$

The binary archive associated with each scenario is denoted by $\mathcal{B}_{\rho}^{(r_{\text{scn}})}$, while derived diagnostic products are collected in $\mathcal{D}_{\rho}^{\text{diag}}$.

Three primary product classes support downstream applications. The first class is the raw three-dimensional density field at each epoch, which provides direct input for drag computation. The second class is the temporal sequence of these fields, representing the atmosphere as a dynamic environmental history. The third class consists of reduced diagnostics derived from the primary fields, including altitude-time histories, latitude-time cross-sections, shell-specific climatologies, and global reference curves. Raw fields are used for simulation, while diagnostics support interpretation, calibration, and comparative analysis [35, 81, 36].

This framework is well-suited to satellite propagation because operational satellites are distributed across multiple shells and experience different altitude, latitude, longitude, and local-time histories. A global mean density would obscure spatial gradients and local-time effects, thereby masking drag variability throughout the satellite constellation. Gridded atmospheric fields allow the propagator to interpolate ρ_{dyn} at each satellite's geodetic position and epoch, supporting direct evaluation of atmospheric drag along each trajectory. Incorporating forecast-driven solar and geomagnetic inputs therefore produces a temporally dynamic and spatially resolved density model [34, 35, 12, 1].

The coupling architecture also enables scenario analysis. Processing multiple stochastic geomagnetic driver realizations through the same NRLMSIS grid generator produces alternative

density histories under common large-scale solar conditions. The resulting ensemble of atmospheric states supports propagation uncertainty quantification, ballistic-coefficient estimation, and fleet-level sensitivity studies within a forecast-conditioned thermospheric environment [35, 36].

Chapter 5

Physical Parameter Estimation via Bayesian Inference

5.1 TLE Decay Observables

This analysis formulates a constrained inverse problem in late-life LEO dynamics. The objective is to infer an orbit-effective ballistic coefficient, denoted BC_s^{eff} , for decaying Starlink satellites using public GP/TLE catalog histories [154, 155, 156, 12, 36]. The inferred parameter is subsequently propagated through a forecast-conditioned force model to support physically consistent decay and re-entry simulations [35, 81, 157]. Under the existing convention $BC = m/(C_D A_D)$, this parameter consolidates unresolved aerodynamic and environmental influences, including attitude variation, projected-area history, gas-surface interaction, and thermospheric model discrepancy, since these factors remain unresolved in archival catalog data [158, 38, 159, 36]. The formulation estimates an interval-effective drag-response parameter that maintains predictive fidelity for drag-dominated propagation while remaining identifiable from the observed decay record.

5.1.1 Late-Life Satellite Decay as an Inverse Problem

Late-life satellite trajectories provide a robust basis for inferring drag parameters. Significant orbital-energy loss during this phase generates a clear secular trend in the catalog-derived orbital record, and the uniformity of satellite class and operational regime facilitates cross-satellite analysis [154, 155, 156]. For satellite s , the retained late-life decay arc is denoted by $\mathcal{D}_s^{\text{dec}}$. Records within this arc are indexed by k_{dec} , a local decay-record index introduced to preserve the existing use of k as the coverage-multiplicity symbol. Each retained record contains a catalog epoch, a GP/TLE

mean-element state, and a catalog mean motion from which the decay observable is constructed.

The inverse problem is constrained by public observability. The accessible observations are time-ordered catalog states, while attitude, projected area, accommodation behavior, drag coefficient, and local density error remain latent or only partially represented [16, 17, 160, 12]. The objective is to identify the effective aerodynamic response that maximizes consistency between observed decay and a specified forward dynamical model.

The available data constrain an aggregate drag response more strongly than separate fixed values of drag coefficient and geometric area. In practical terms, the record constrains the parameter that reproduces the observed semi-major-axis evolution over the retained analysis interval [154, 155, 156, 159]. The resulting BC_s^{eff} is therefore interpreted as a simulation-calibrated effective ballistic coefficient. This interpretation accounts for the degeneracy inherent in decay signatures, since comparable orbital decay can arise from different combinations of atmospheric density, area history, and drag coefficient [158, 38, 159, 36]. Forward propagation requires the aggregate parameter governing $\mathbf{a}_{\text{drag}}$ over the analysis window, which makes BC_s^{eff} the identifiable quantity for the present public-data formulation.

5.1.2 Semi-Major Axis from TLE Mean Motion

The primary observable is the semi-major axis derived from catalog mean motion. Let $n_{s,k_{\text{dec}}}^{\text{rev}}$ denote the raw GP/TLE mean motion for satellite s at decay-record index k_{dec} , reported in revolutions per day. Let $n_{s,k_{\text{dec}}}^{\text{cat}}$ denote the same quantity after conversion to radians per second. The observed semi-major-axis proxy is computed from the Keplerian mean-motion relation

$$a_{s,k_{\text{dec}}}^{\text{obs}} = \left(\frac{\mu_{\oplus}}{\left(n_{s,k_{\text{dec}}}^{\text{cat}} \right)^2} \right)^{1/3}, \quad n_{s,k_{\text{dec}}}^{\text{cat}} = \frac{2\pi}{86400} n_{s,k_{\text{dec}}}^{\text{rev}} \quad (5.1)$$

where $\mu_{\oplus}$ is the Earth standard gravitational parameter. The unit conversion is applied before the Keplerian relation is evaluated [2, 1, 17, 88, 20].

The semi-major axis is selected as the decay-sensitive scalar most reliably recoverable from GP/TLE products. Atmospheric drag primarily dissipates orbital mechanical energy, and this loss appears in the catalog as a secular reduction of the orbit. The derived scalar $a_{s,k_{\text{dec}}}^{\text{obs}}$ provides a stable and interpretable signal for the inverse problem, with stronger relevance to drag-dominated decay than isolated inspection of angular elements [154, 155, 156]. This construction respects the model-conditioned mean-element character of GP/TLE records and uses the semi-major-axis

sequence as a catalog-consistent proxy for drag-induced orbital-energy loss [16, 17, 160, 88].

5.1.3 Anchored Residuals and the Inference-Prediction Split

The observation series is expressed using anchored semi-major-axis residuals over the retained decay arc $\mathcal{D}_s^{\text{dec}}$. For satellite s , each residual is formed relative to the first retained semi-major-axis value,

$$\Delta a_{s,k_{\text{dec}}}^{\text{obs}} = a_{s,k_{\text{dec}}}^{\text{obs}} - a_{s,0}^{\text{obs}} \quad (5.2)$$

This transformation directs the inference toward relative orbital evolution across the decay arc. Anchoring the series at the initial retained epoch reduces sensitivity to catalog-propagation offsets and isolates the decay dynamics most relevant for inverse inference. The likelihood construction therefore emphasizes relative orbital lowering, reducing dependence on precise initial alignment with catalog elements. This approach is appropriate for public GP/TLE data, where the objective is to recover the decay trajectory indicated by the mean-element record [17, 160, 154].

Anchoring isolates the temporal evolution of the semi-major axis, which is strongly influenced by atmospheric drag during late-life decay. The resulting residual series provides a stable scalar observation model for estimating BC_s^{eff} , while preserving the secular trend required for posterior predictive simulation.

The retained time series is partitioned into an inference interval and a held-out prediction interval. A satellite-specific split epoch $t_{\text{split},s}$ separates the two subsets. Samples satisfying $t_{s,k_{\text{dec}}}^{\text{cat}} \leq t_{\text{split},s}$ define the inference subset $\mathcal{D}_s^{\text{inf}}$, and samples satisfying $t_{s,k_{\text{dec}}}^{\text{cat}} > t_{\text{split},s}$ define the held-out prediction subset $\mathcal{D}_s^{\text{pred}}$. The inference subset informs estimation of the effective drag response, while the prediction subset enables posterior evaluation beyond the calibration window. This partitioning supports assessment of both parameter plausibility and forward-simulation skill.

5.1.4 Retained Satellite Decay Arcs

The dataset includes six late-life Starlink decay arcs selected to provide tractable, nontrivial inference cases. The retained index set $\mathcal{I}_{\text{dec}}$ contains satellites sat1010, sat1011, sat1013, sat1014, sat1019, and sat1030. These arcs exhibit various lengths, observation counts, and terminal-decay profiles, exposing the estimation framework to variable late-life behavior within a common satellite class. This selection supports individual inference and cross-satellite comparison while avoiding the computational and interpretive burden of a full-fleet analysis.




Figure 5.1: Retained Late-Life Semi-major Axis Histories

Figure 5.1 shows the retained semi-major-axis histories for the six late-life decay arcs, illustrating the distinct observation spans, terminal descent rates, and curvature regimes used to condition the satellite-specific inverse problems.

The heterogeneity of the retained arcs motivates satellite-specific estimation of BC_s^{eff} . Differences in observation span, altitude range, and decay morphology indicate that a single pooled parameter would obscure satellite-level aerodynamic and environmental differences. Subsequent cross-satellite analysis is therefore performed using posterior summaries, uncertainty intervals, and predictive performance.

Table 5.1: Starlink Decay Arcs for Ballistic Coefficient Inference

Satellite ID	Observations	Epoch Span (UTC)	Altitude Range (km)
sat1010	55	2025-10-25 22:00 to 2025-11-15 09:04	153.9–288.0
sat1011	50	2025-11-19 22:00 to 2025-12-06 00:42	152.2–277.1
sat1013	44	2025-08-16 12:07 to 2025-09-01 02:56	143.1–268.2
sat1014	60	2025-06-21 14:00 to 2025-07-15 17:18	153.3–308.3
sat1019	20	2025-12-27 22:00 to 2025-12-31 16:00	385.9–388.9
sat1030	49	2025-08-30 22:00 to 2025-09-19 09:13	147.2–297.1

5.2 Forecast-Conditioned Forward Model for Satellite Drag Simulation

The forward model maps an orbit-effective drag parameter into simulated Starlink decay histories by directly propagating the translational state. The purpose is to construct a reproducible and physically interpretable simulator suitable for quantitative comparison with TLE-derived semi-major-axis records [2, 1, 12]. The retained force terms capture the dominant contributors to late-life LEO decay over the short analysis intervals considered here, while the atmospheric pathway incorporates forecast-conditioned thermospheric density into the likelihood evaluation [35, 81, 157].

5.2.1 Equations of Motion and Retained Perturbations

Propagation is performed in Cowell form in an Earth-centered inertial frame. Let $\mathbf{r}(t)$ and $\mathbf{v}(t)$ denote the satellite position and velocity. The translational dynamics are written as

$$\dot{\mathbf{r}} = \mathbf{v}, \quad (5.3)$$

$$\dot{\mathbf{v}} = -\frac{\mu_{\oplus}}{r^3}\mathbf{r} + \mathbf{a}_{J_2} + \mathbf{a}_{3b}^{(\odot)} + \mathbf{a}_{3b}^{(\zeta)} + \mathbf{a}_{\text{SRP}} + \mathbf{a}_{\text{drag}} \quad (5.4)$$

where $\mu_{\oplus}$ is the Earth standard gravitational parameter and $r = \|\mathbf{r}\|$. The retained perturbation set is intentionally compact while preserving the dominant physical mechanisms for drag-sensitive late-life propagation. It includes the leading zonal component of Earth's nonspherical gravity, direct-minus-indirect third-body forcing from the Sun and Moon, SRP with eclipse handling, and atmospheric drag evaluated using a forecast-conditioned thermospheric density field [2, 1, 37, 35].

The Earth gravity field is represented at the degree-two axisymmetric level, preserving first-order secular precession and short-period geometric corrections while allowing drag to dominate dissipative evolution [2, 1, 37, 161]. In the Earth-fixed frame, the retained disturbing potential is written as

$$U_{20}(\mathbf{r}, t) = \frac{\mu_{\oplus}}{r} \left(\frac{R_{\oplus}}{r} \right)^2 \bar{P}_{20}(\sin \phi) \bar{C}_{20} \quad (5.5)$$

The corresponding inertial acceleration $\mathbf{a}_{J_2}$ is obtained by transforming the inertial position into the Earth-fixed frame, evaluating the gradient contribution associated with U_{20} , and rotating the resulting acceleration back into the inertial frame. This representation preserves the essential secular geometry for LEO without expanding the force model beyond the accuracy needed for short, drag-dominated decay arcs [2, 1, 37].

For a retained third body $b \in \{\odot, \mathbb{C}\}$, with geocentric position $\mathbf{r}_b$ and gravitational parameter μ_b , the third-body acceleration is written in direct-minus-indirect form as

$$\mathbf{a}_{3b}^{(b)} = \mu_b \left(\frac{\mathbf{r}_b - \mathbf{r}}{\|\mathbf{r}_b - \mathbf{r}\|^3} - \frac{\mathbf{r}_b}{\|\mathbf{r}_b\|^3} \right) \quad (5.6)$$

Although these perturbations are secondary over short, drag-dominated arcs, their inclusion improves geometric consistency during repeated propagation and reduces model-form error during likelihood evaluation [2, 1].

SRP is retained in cannonball form with the existing eclipse factor χ_{ecl} . Let $r_\odot$ denote the Sun-satellite distance and let $\hat{\mathbf{s}}$ denote the Sun-line unit vector in the acceleration direction. The SRP acceleration is modeled as

$$\mathbf{a}_{\text{SRP}} = \chi_{\text{ecl}} P_\odot C_R \left(\frac{A_{\text{SRP}}}{m} \right) \left(\frac{\text{AU}}{r_\odot} \right)^2 \hat{\mathbf{s}} \quad (5.7)$$

SRP has a small effect on late-life semi-major-axis decay for Starlink-class LEO objects relative to atmospheric drag. Its inclusion nevertheless prevents the inferred drag parameter from compensating for a known non-gravitational perturbation [2, 1, 12].

5.2.2 NRLMSIS Density Realization

Atmospheric forcing is modeled using the forecast-conditioned thermospheric workflow developed previously. The forward model evaluates a time-dependent NRLMSIS 2.0 density realization conditioned on solar and geomagnetic drivers [35, 81]. Drag is therefore evaluated using a driver-aware thermosphere that evolves with epoch, geodetic location, altitude, and forecasted space-weather state.

The satellite-sampled density along the propagated trajectory is written as

$$\rho_{\text{fc},s}(t) = \widehat{\rho}_{\text{fc}}(t, \varphi_g(t), \lambda_g(t), h_g(t)) \quad (5.8)$$

where $\widehat{\rho}_{\text{fc}}$ is the forecast-conditioned thermospheric density field, φ_g is geodetic latitude, λ_g is geodetic longitude, and h_g is geodetic altitude. The drag acceleration is then expressed using the existing ballistic-coefficient convention $BC = m/(C_D A_D)$:

$$\mathbf{a}_{\text{drag}} = -\frac{1}{2} \frac{\rho_{\text{fc},s}(t)}{BC_s^{\text{eff}}} \|\mathbf{v}_{\text{rel}}\| \mathbf{v}_{\text{rel}} \quad (5.9)$$

with

$$\mathbf{v}_{\text{rel}} = \mathbf{v} - \boldsymbol{\omega}_{\oplus} \times \mathbf{r} \quad (5.10)$$

where $\boldsymbol{\omega}_{\oplus}$ is the Earth rotation vector and BC_s^{eff} is the orbit-effective ballistic coefficient inferred for satellite s . This formulation keeps the inferred parameter aligned with the physical drag model while allowing thermospheric variability to enter through the forecast-conditioned density realization. The resulting estimate reflects the drag response required by the specified atmospheric pathway, instead of forcing a static atmosphere to absorb space-weather-driven density variability [12, 158, 38, 159, 35, 157].

5.2.3 Offline Density Grids and Runtime Interpolation

Evaluating the full thermosphere model during each force computation is computationally prohibitive in repeated-likelihood scenarios, especially when propagating multiple Markov Chain Monte Carlo chains across multiple satellites. To address this constraint, the implementation separates the atmospheric model evaluation from runtime force evaluation. Forecast-driven NRLMSIS 2.0 outputs are precomputed offline on fixed geodetic grids and stored as time-ordered density fields. These grids are indexed, optionally compressed, and accessed during propagation through an interpolation mechanism [35, 81, 76].

The grid uses fixed latitude, longitude, and altitude nodes, supplemented by multiple Universal Time snapshots each day. During propagation, inertial satellite states are converted to geodetic coordinates, and local density is obtained through spatial and temporal interpolation within the grid sequence. For atmospheric scenario r_{scn} , this runtime operation is represented compactly as

$$\rho_{\text{grid},s}^{(r_{\text{scn}})}(t) = \mathcal{I}_{\rho} \left[\mathcal{G}_{\rho,\text{seq}}^{(r_{\text{scn}})} \right] (t, \varphi_g(t), \lambda_g(t), h_g(t)) \quad (5.11)$$

where $\mathcal{I}_{\rho}$ denotes the runtime interpolation operator, $\mathcal{G}_{\rho,\text{seq}}^{(r_{\text{scn}})}$ is the time-ordered forecast-conditioned density-grid sequence, and $(\varphi_g, \lambda_g, h_g)$ are the geodetic coordinates of the propagated satellite state. This expression formalizes the query performed by the forward model without reintroducing the full empirical atmosphere evaluation inside the integrator.

This method preserves the information content of the driver-conditioned thermosphere while transforming repeated atmospheric queries into a structured interpolation task. It serves as a computational surrogate for the density lookup process, altering only the runtime querying mechanism within the likelihood evaluation while preserving the underlying forecast-conditioned density field. The resulting implementation enables tractable Markov Chain Monte Carlo loops without replacing

the physical thermosphere pathway used to generate the density grids [35, 81, 76].

5.2.4 Terminal Event and Propagation Horizon

The forward model focuses on late-stage, drag-dominated propagation rather than unconstrained long-horizon integration. Propagation terminates at an altitude-defined event, marking the transition beyond the regime of meaningful orbital analysis for the intended application. Let

$$h(t) = \|\mathbf{r}(t)\| - R_{\oplus} \quad (5.12)$$

then the terminal condition is defined as

$$g_{\text{re}}(t, \mathbf{y}) = h(t) - h_{\text{stop}} \quad (5.13)$$

with a terminal negative crossing. The threshold altitude h_{stop} defines the lower limit of relevant orbital propagation for the Starlink decay scenario. In implementation, this corresponds to an event function with terminal behavior and downward-crossing detection, consistent with event-driven numerical integration [74, 162, 76]. This event is combined with a satellite-specific calendar cutoff $t_{\text{cal},s}^{\text{max}}$, constraining propagation to the finite interval $\mathcal{T}_s^{\text{prop}}$ used for posterior-predictive simulation.

The propagation horizon is therefore determined by two constraints: dynamical termination at the threshold altitude and chronological termination at the prescribed analysis cutoff. This dual stopping criterion supports the inferential objective of reproducing and extending observed decay histories within the operationally significant late-life regime, without modeling coupled aerothermodynamic entry below the altitude where the retained force formulation remains valid [2, 12, 157].

5.3 Effective Ballistic-Coefficient Parameterization

The inverse problem under consideration does not produce a uniquely determined aerodynamic state. The TLE decay record constrains the drag response that most accurately reproduces observed orbital-energy loss within the selected arc [154, 155, 156]. This response is efficiently represented by an effective ballistic coefficient. The choice of parameterization is dictated by observability: the available archive does not provide direct measurements of projected-area variability, attitude dynamics, gas-surface interaction parameters, or density-model discrepancy as independent variables. Forward simulations are primarily sensitive to the combined influence of these factors through the

drag law [158, 38, 159, 36]. The inferred parameter is therefore an orbit-relevant quantity defined over the modeling window and selected to optimize predictive accuracy in satellite decay simulations, with reconstruction of latent physical mechanisms left outside the identifiable scope of the catalog record.

5.3.1 Effective Drag Representation

The retained aerodynamic state is expressed through the orbit-effective drag-coupling parameter

$$B_s^{\text{eff}} = \frac{C_{D,s} A_{D,s}^{\text{eff}}}{m_s} \quad (5.14)$$

and the corresponding effective ballistic coefficient

$$BC_s^{\text{eff}} = \frac{1}{B_s^{\text{eff}}} = \frac{m_s}{C_{D,s} A_{D,s}^{\text{eff}}} \quad (5.15)$$

Here, $C_{D,s}$ represents the drag coefficient, $A_{D,s}^{\text{eff}}$ is the effective drag area exposed to the flow, and m_s is the satellite mass. The symbol B_s^{eff} is used for reciprocal ballistic-coefficient coupling because β is reserved elsewhere in this thesis for the visibility-cap half-angle. Drag acceleration is determined by the aggregate aerodynamic coupling, not by individually identifiable values of $C_{D,s}$, $A_{D,s}^{\text{eff}}$, and m_s [158, 38, 12, 159]. This parameterization aligns with the resolution permitted by the available public data. The simulation requires a scalar quantity that translates modeled thermospheric density into observed semi-major-axis decay; B_s^{eff} and BC_s^{eff} provide this role with minimal parameter dimensionality.

This representation provides a physically interpretable inferential objective. A larger B_s^{eff} signifies stronger drag coupling and more rapid orbital-energy dissipation under identical density and relative-velocity conditions. A larger BC_s^{eff} indicates weaker drag coupling and slower decay. The inferred parameter serves as the scalar determinant of the satellite response to atmospheric forcing over the analyzed arc.

5.3.2 Random-Tumbling Effective-Area Surrogate

The principal unresolved geometric component in the drag law is the effective area exposed to flow. Since time-resolved attitude data are absent from the TLE record, the model substitutes the instantaneous projected area with a time-averaged surrogate suitable for unresolved attitude variation. Let $A_{\text{face},s}$ denote the planform area of a flat-plate analog for satellite s . Under isotropic

orientation averaging, the mean projected area is related to surface area through Cauchy’s surface-area formula [163]. Restricting this approximation to a two-sided flat plate and omitting edge contributions gives

$$A_{D,s}^{\text{eff}} = \frac{A_{\text{face},s}}{2} \quad (5.16)$$

The retained drag-coupling parameter can then be written as

$$B_s^{\text{eff}} = \frac{C_{D,s}}{m_s} \left(\frac{1}{2} s_{\text{att},s} A_{\text{face},s} \right) \quad (5.17)$$

where $s_{\text{att},s}$ is a scalar correction factor that accounts for systematic deviations from ideal isotropic tumbling [154, 163].

This surrogate is intentionally low dimensional, where the role is to encapsulate unresolved attitude variability within a geometrically meaningful effective parameter without requiring an unobservable attitude history. The approximation preserves the dominant influence of projected-area variability on drag acceleration while remaining matched to the resolution supported by the decay data. The resulting parameterization is a controlled approximation for ballistic-coefficient inference under public catalog observability.

5.3.3 Identifiability Under TLE-Derived Decay

Decay histories constrain the aggregate drag response, which limits unique attribution among thermospheric density, projected geometry, and gas-surface interaction effects. This distinction is central to the interpretation of recovered ballistic coefficients. Variations in thermospheric density, effective projected area, and free-molecular drag coefficient all modify the observed decay rate [158, 38, 159, 36]. The TLE-derived semi-major-axis sequence is primarily sensitive to the aggregate aerodynamic loading that appears in the drag law, with weaker separability among $C_{D,s}$, $A_{D,s}^{\text{eff}}$, m_s , and the density realization [154, 155, 156, 12]. The recovered posterior is therefore a conditional estimate of effective aerodynamic coupling under the specified forward model, forecast-conditioned density pathway, and catalog-derived observation process.

This limitation improves the interpretive clarity of the inferred parameter for simulation. The workflow does not attempt to recover laboratory-grade drag coefficients, time-resolved attitude histories, or independent density corrections from cataloged data alone, because these quantities are inaccessible from the retained public observation layer. Instead, the inference targets the parameter combination that allows the forward model to reproduce the observed decay arc with posterior consistency. The absence of physical separability is therefore handled through the effective

ballistic-coefficient parameterization.

5.3.4 Effective Parameterization

The ballistic-coefficient framework is sufficient for simulating Starlink late-life decay and re-entry uncertainty without requiring comprehensive aerodynamic system identification. Once the posterior for $A_{\text{face},s}$, B_s^{eff} , or BC_s^{eff} is inferred from the TLE decay history, posterior samples are propagated through the forward model to generate semi-major-axis trajectories, terminal descent times, and re-entry windows. These quantities directly depend on the effective drag response defined by the chosen parameterization [154, 155, 156].

This approach balances physical structure with observational realism. It preserves the essential dependencies of drag on density, relative velocity, and aerodynamic loading while avoiding latent variables that are unconstrained by the catalog record. For decaying Starlink objects, the inferred BC_s^{eff} provides the appropriate scalar parameter for posterior predictive propagation and uncertainty quantification across the late-life interval.

5.4 Bayesian Inference and MCMC Implementation

The statistical model links the retained TLE decay observable to the forecast-conditioned forward propagator through a one-parameter posterior distribution that characterizes the effective aerodynamic state. The primary inferential objective is to recover a usable orbit-effective ballistic coefficient for satellite constellation simulation. The posterior distribution is constructed using anchored semi-major-axis residuals, a bounded prior for the one-sided face-area surrogate, and a random-walk Metropolis algorithm evaluated through the forward model [164, 165, 166].

5.4.1 Likelihood on Anchored Semi-Major-Axis Residuals

For satellite s , the retained TLE-derived observable is the anchored semi-major-axis residual series over the inference subset $\mathcal{D}_s^{\text{inf}}$. Using the existing decay-record index k_{dec} , the observed anchored residual is

$$\Delta a_{s,k_{\text{dec}}}^{\text{obs}} = a_{s,k_{\text{dec}}}^{\text{obs}} - a_{s,0}^{\text{obs}} \quad (5.18)$$

The latent one-parameter state is written as

$$\vartheta_s = \log A_{\text{face},s} \quad (5.19)$$

where $A_{\text{face},s}$ is the one-sided face-area surrogate introduced in the effective-area parameterization. This choice keeps the inferred parameter positive after exponentiation and places the inference directly on the unresolved geometric component that enters B_s^{eff} and BC_s^{eff} [164].

For a candidate value ϑ_s , the forward model is evaluated at the retained catalog epochs to obtain the modeled anchored series

$$\Delta a_{s,k_{\text{dec}}}^{\text{mod}}(\vartheta_s) = a_{s,k_{\text{dec}}}^{\text{mod}}(\vartheta_s) - a_{s,0}^{\text{mod}}(\vartheta_s) \quad (5.20)$$

and the residual mismatch is written as

$$r_{s,k_{\text{dec}}}(\vartheta_s) = \Delta a_{s,k_{\text{dec}}}^{\text{obs}} - \Delta a_{s,k_{\text{dec}}}^{\text{mod}}(\vartheta_s) \quad (5.21)$$

The retained likelihood assumes conditionally independent Gaussian residuals,

$$r_{s,k_{\text{dec}}}(\vartheta_s) \sim \mathcal{N}(0, \sigma_{a,s}^2) \quad (5.22)$$

which yields the log-likelihood

$$\log p(\Delta \mathbf{a}_s^{\text{obs}} \mid \vartheta_s, \sigma_{a,s}) = -\frac{1}{2} \sum_{k_{\text{dec}} \in \mathcal{D}_s^{\text{inf}}} \left(\frac{r_{s,k_{\text{dec}}}(\vartheta_s)}{\sigma_{a,s}} \right)^2 - N_s^{\text{inf}} \log \sigma_{a,s} - \frac{N_s^{\text{inf}}}{2} \log(2\pi) \quad (5.23)$$

Here, $\Delta \mathbf{a}_s^{\text{obs}}$ denotes the retained vector of anchored observed residuals and N_s^{inf} is the number of samples in the inference interval [164].

The observational scale $\sigma_{a,s}$ is estimated separately for each satellite from a detrended semi-major-axis record using a robust median-absolute-deviation procedure, with regularization by a lower bound $\sigma_{a,\text{min}}$ to suppress excessive likelihood curvature in short or anomalously smooth arcs [167]. A least-squares linear trend is subtracted prior to scale estimation, outliers are clipped in the detrended residuals, and a semi-major-axis scale floor of 0.25 km is imposed. This likelihood formulation stabilizes repeated posterior evaluation and preserves the principal decay signal in the anchored TLE series, without requiring strict Gaussianity or temporal independence as literal data-generating assumptions.

5.4.2 Prior on the Inferred Aerodynamic State

The latent parameter is $\vartheta_s = \log A_{\text{face},s}$, where $A_{\text{face},s}$ represents the one-sided face area of the satellite flat-plate surrogate described previously. The prior distribution is Gaussian in log space

and truncated to retain physically admissible support:

$$\vartheta_s \sim \mathcal{N}\left(\mu_{\log A,s}, \sigma_{\log A}^2\right), \quad A_{\min} \leq \exp(\vartheta_s) \leq A_{\max} \quad (5.24)$$

The prior mean $\mu_{\log A,s}$ corresponds to the nominal satellite face-area surrogate derived from the adopted solar-array planform and the largest chassis face. The log-space prior width is set to $\sigma_{\log A} = 0.75$, and support is restricted to $0.5 \leq A_{\text{face},s} \leq 200 \text{ m}^2$. The corresponding log prior is

$$\log p(\vartheta_s) = -\frac{1}{2} \left(\frac{\vartheta_s - \mu_{\log A,s}}{\sigma_{\log A}} \right)^2 - \log \sigma_{\log A} - \frac{1}{2} \log(2\pi) \quad (5.25)$$

with value $-\infty$ outside the retained support.

Combining the likelihood and prior gives the unnormalized satellite-specific posterior

$$\log \pi_s(\vartheta_s \mid \Delta \mathbf{a}_s^{\text{obs}}, \sigma_{a,s}) = \log p(\Delta \mathbf{a}_s^{\text{obs}} \mid \vartheta_s, \sigma_{a,s}) + \log p(\vartheta_s) + C \quad (5.26)$$

where C is independent of ϑ_s and is not required for random-walk Metropolis acceptance ratios [164, 166].

This prior accommodates substantial uncertainty in effective geometry while excluding non-physical or operationally irrelevant parameter regions. The regularizing effect stabilizes inference in the one-parameter nonlinear inverse problem and prevents posterior concentration on implausible geometric assumptions.

5.4.3 Random-Walk Metropolis and Warmup Adaptation

Sampling uses a random-walk Metropolis kernel defined in logarithmic area coordinates, with the satellite-specific state represented by $\vartheta_s = \log A_{\text{face},s}$. For chain c_{ch} and MCMC iteration m_{mc} , proposal states are generated according to

$$\vartheta'_s = \vartheta_{s,m_{\text{mc}}}^{(c_{\text{ch}})} + \sigma_{\text{prop},s,m_{\text{mc}}} \varepsilon_{m_{\text{mc}}}, \quad \varepsilon_{m_{\text{mc}}} \sim \mathcal{N}(0, 1) \quad (5.27)$$

where

$$\sigma_{\text{prop},s,m_{\text{mc}}} = \exp(\ell_{\text{prop},s,m_{\text{mc}}}) \quad (5.28)$$

is the proposal standard deviation in log-area space. The Metropolis-Hastings acceptance probability is

$$\alpha_{\text{MH}}(\vartheta_s, \vartheta'_s) = \min \left\{ 1, \exp \left[\log \pi_s \left(\vartheta'_s \mid \Delta \mathbf{a}_s^{\text{obs}}, \sigma_{a,s} \right) - \log \pi_s \left(\vartheta_s \mid \Delta \mathbf{a}_s^{\text{obs}}, \sigma_{a,s} \right) \right] \right\} \quad (5.29)$$

Here, $\log \pi_s(\cdot)$ denotes the unnormalized satellite-specific log posterior. Each posterior evaluation requires direct numerical propagation at the retained observation epochs, so the dominant computational cost is repeated solution of the forward initial-value problem [165, 166, 74, 162, 76].

For each satellite, $N_{\text{ch}} = 4$ independent chains are initialized. Each chain contains $N_{\text{iter}} = 1000$ total iterations, with $N_{\text{warm}} = 500$ warmup-adaptation iterations and $N_{\text{post}} = 500$ post-warmup retained samples. Proposal adaptation is limited to the warmup phase and is performed using a Robbins-Monro stochastic approximation for the logarithmic proposal scale:

$$\ell_{\text{prop},s,m_{\text{mc}}+1} = \ell_{\text{prop},s,m_{\text{mc}}} + \gamma_{m_{\text{mc}}}^{\text{RM}} (I_{\text{acc},s,m_{\text{mc}}} - p_{\text{acc}}^{\star}) \quad (5.30)$$

where $I_{\text{acc},s,m_{\text{mc}}} \in \{0, 1\}$ is the acceptance indicator, $p_{\text{acc}}^{\star} = 0.40$ is the target acceptance probability, and

$$\gamma_{m_{\text{mc}}}^{\text{RM}} = \min \left(0.05, \frac{1}{\sqrt{m_{\text{mc}} + 1}} \right) \quad (5.31)$$

After warmup, adaptation is discontinued and the proposal scale remains fixed for post-warmup sampling [168, 169]. Posterior summaries are computed from the retained post-warmup draws. The processed archive enforces a minimum observation spacing of $\Delta t_{\text{obs},\text{min}} = 2$ days and a maximum retained observation count of $N_{\text{obs},s}^{\text{max}} = 60$ per satellite, reducing computational overhead while preserving the dominant decay trend in the likelihood evaluations.

5.4.4 Diagnostic Criteria and Archive Limitations

Interpretation of posterior summaries depends on rigorous assessment of Monte Carlo quality. The primary diagnostics are rank-normalized split- $\widehat{R}$, bulk and tail effective sample sizes, and Monte Carlo standard error [170, 171]. These diagnostics are appropriate for this low-dimensional posterior because independent chain initializations, stable between-chain agreement, and adequate effective sample sizes are required before posterior summaries can be used in downstream satellite constellation simulations.

Table 5.2: Posterior Summaries and Monte Carlo Diagnostics

Satellite	BC Median (kg m^{-2})	95% HDI (kg m^{-2})	$\widehat{R}$	ESS_{bulk}	ESS_{tail}	MCSE (kg m^{-2})	n_{obs}
sat1010	36.033	[36.001, 36.110]	1.0116	287.38	426.88	0.0020	55
sat1011	38.351	[38.326, 38.405]	1.0075	270.54	155.04	0.0024	50
sat1013	28.894	[28.863, 28.969]	1.0107	282.53	353.50	0.0020	44
sat1014	21.642	[21.618, 21.710]	1.0110	168.99	162.43	0.0034	60
sat1019	30.975	[30.612, 31.360]	1.0031	472.00	419.42	0.0089	20
sat1030	29.794	[29.763, 29.948]	1.0088	231.57	244.92	0.0039	49

For the six satellites analyzed, the processed archive produces split- $\widehat{R}$ values near unity, bulk effective sample sizes spanning 169 to 472, and comparable tail effective sample sizes. The Monte Carlo standard error values for BC_s^{eff} are small relative to the posterior scale, supporting the use of posterior medians and highest-density intervals in subsequent propagation [172, 171]. The primary limitations stem from archive incompleteness: exported samples lack chain identifiers and log-posterior values, and stored traces are partial. This necessitates reconstruction of synthetic equal-length chains from post-warmup samples and prevents comprehensive acceptance-trace diagnostics. The diagnostic summary supports posterior reporting, although it remains weaker than a complete stepwise chain log containing chain index, proposal state, acceptance indicator, posterior density, and proposal scale at every iteration.

5.5 Per-Satellite and Constellation-Level Ballistic Coefficients

This section presents the recovered Starlink ballistic-coefficient structure under the retained forecast-conditioned density model. The analysis produces satellite-specific posterior distributions in place of a fleet-averaged coefficient. The posterior locations, interval widths, and pairwise separations quantify whether the retained decay arcs support common or distinct effective drag regimes. This differentiation directly affects simulation fidelity. Although a pooled coefficient can provide a compact fleet-level summary, posterior propagation for an individual satellite should rely on the corresponding satellite-specific distribution unless the cross-satellite posterior structure demonstrates genuine interchangeability [164, 172, 171].

5.5.1 Per-Satellite Posterior Summaries

The recovered posterior medians, denoted by $\widetilde{BC}_s$, display a clear ordering among the retained satellites. `sat1011` and `sat1010` possess the highest coefficients, with posterior medians near 38.35 kg m^{-2} and 36.03 kg m^{-2} , respectively. `sat1019`, `sat1030`, and `sat1013` yield intermediate values, with medians near 30.98 , 29.79 , and 28.89 kg m^{-2} . `sat1014` has the lowest median at 21.64 kg m^{-2} . This hierarchy indicates meaningful variation in the effective drag response within the retained model. Higher BC_s^{eff} values correspond to weaker effective drag coupling for a given density realization, whereas lower values indicate stronger drag coupling and more rapid orbital-energy dissipation [158, 12].

Table 5.3: Per-Satellite Posterior Summaries

Satellite	BC_s^{eff} Median (kg m^{-2})	95% HDI (kg m^{-2})
<code>sat1010</code>	36.033	[36.001, 36.110]
<code>sat1011</code>	38.351	[38.326, 38.405]
<code>sat1013</code>	28.894	[28.863, 28.969]
<code>sat1014</code>	21.642	[21.618, 21.710]
<code>sat1019</code>	30.975	[30.612, 31.360]
<code>sat1030</code>	29.794	[29.763, 29.948]

The posterior intervals further reinforce the differentiation among satellites. In five of six cases, the 95% highest-density intervals are narrow relative to the separations between posterior centers, indicating that the retained data and forward model support distinct effective drag regimes [172, 171]. The broadest interval occurs for `sat1019`, which corresponds to the shortest retained arc and provides limited information on longer-term decay curvature. This structure reflects the geometry of the processed decay archive: longer and more informative decay arcs enable more precise effective drag identification, while shorter arcs produce greater posterior uncertainty under the same observation model.




Figure 5.2: Starlink Ballistic Coefficients Forest

Figure 5.2 provides the corresponding forest summary of the per-satellite BC_s^{eff} posteriors, showing the posterior medians, interval widths, and separation among the retained Starlink decay cases.

5.5.2 Conditional Posterior Concentration

Narrow highest-density intervals require careful interpretation. These intervals represent conditional concentration within the retained inferential structure, which combines the anchored semi-major-axis likelihood, the forecast-conditioned forward model, the selected density realization pathway, and the effective-area surrogate. The sharpness of the posteriors follows from the model’s focus on a dominant drag parameter and the substantial information content supplied by the retained decay histories within this framework [164, 172].

The reported posterior widths quantify uncertainty in the retained effective drag response under the specified modeling assumptions, with broader physical factors represented only through the selected model structure. Alternative density realizations, time-resolved attitude models, or more flexible observation models would generally alter posterior locations and interval widths. The inferred BC_s^{eff} values are therefore sharply identified within the current conditional framework, and this concentration supports posterior-predictive simulation under the same assumptions used during inference [164, 154].

5.5.3 Cross-Satellite Heterogeneity

The recovered ballistic coefficients exhibit a range exceeding a factor of 1.7 across the retained sample, spanning from approximately 21.64 kg m^{-2} for sat1014 to 38.35 kg m^{-2} for sat1011. This variation exceeds the scale attributable to Monte Carlo error and indicates a structured pattern inconsistent with numerical artifact. The observed dispersion therefore marks substantive cross-satellite heterogeneity in effective drag response. Potential contributors include differences in late-life attitude dynamics, projected-area evolution, unmodeled configuration changes, and discrepancies between the retained atmospheric density realization and the thermospheric conditions encountered along each trajectory [38, 159, 12, 157]. Although the present data do not resolve these sources independently, aggregate drag response varies materially across the retained sample.

This heterogeneity is operationally significant for satellite orbital simulation. A pooled summary may be useful for descriptive fleet-level characterization, but propagation of individual decaying satellites under uncertainty should retain object-specific posterior structure. The posterior dispersion across satellites supports the satellite-level inference strategy adopted here and weakens the justification for imposing a single global ballistic coefficient in subsequent re-entry simulations [154, 155, 156].

5.5.4 Constellation-Level Posterior Structure

Per-satellite posteriors serve as the primary simulation inputs, while satellite constellation-level summaries describe the aggregate drag structure of the retained sample. Equal-weight pooling yields a median ballistic coefficient $\widetilde{BC}_{\text{eq}}$ of approximately 30.20 kg m^{-2} , while ESS-weighted pooling yields $\widetilde{BC}_{\text{ESS}} \approx 30.90 \text{ kg m}^{-2}$. These medians lie near the central three-object group and provide a concise characterization of the six-satellite ensemble. The distinction between object-specific posterior propagation and pooled descriptive summaries is consistent with hierarchical Bayesian practice, in which pooling is useful only when the between-object structure supports information sharing [164].




Figure 5.3: Pairwise Wasserstein-1 Distances between Per-Satellite Posterior Distributions

Figure 5.3 summarizes the pairwise separation among the per-satellite posterior distributions and makes the low- BC , intermediate- BC , and high- BC regimes visible at the distributional level.

Pairwise posterior similarity is quantified by the first Wasserstein distance. Let $P_{BC,s}$ and $P_{BC,s'}$ denote the posterior marginal distributions of BC_s^{eff} and $BC_{s'}^{\text{eff}}$, and let $F_{BC,s}^{-1}$ and $F_{BC,s'}^{-1}$ denote their quantile functions. The pairwise distance is

$$W_{1,s,s'}^{BC} = \int_0^1 \left| F_{BC,s}^{-1}(u) - F_{BC,s'}^{-1}(u) \right| du \quad (5.32)$$

In this sample, the resulting distance matrix delineates a structured posterior landscape [173, 174, 175]. **sat1014** occupies the low- BC regime, **sat1010** and **sat1011** anchor the high- BC regime, and **sat1013**, **sat1019**, and **sat1030** form an intermediate group with smaller mutual distances. This structure identifies where partial pooling or hierarchical modeling may be appropriate and where such models could obscure significant object-level variation [164].

5.6 Posterior-Predictive Satellite Simulation and Re-Entry Uncertainty

The retained inference framework generates a family of decay trajectories by propagating the posterior distribution of the effective ballistic coefficient through the forecast-conditioned forward model. For each satellite s , the posterior over BC_s^{eff} supplies samples that define physically plausible drag responses under the selected thermospheric realization. Propagating these samples produces corresponding late-life orbital decay histories. The resulting posterior-predictive ensemble captures dynamical uncertainty induced by parameter uncertainty, allowing orbital lowering and terminal descent timing to be interpreted as downstream consequences of the inferred drag response [164, 176, 35, 81, 157].

5.6.1 Propagation of Posterior Draws

Let $\{\vartheta_s^{(q_{\text{pp}})}\}_{q_{\text{pp}}=1}^{N_{\text{pp},s}}$ denote the retained posterior samples of the inferred aerodynamic state for satellite s , where q_{pp} indexes posterior-predictive draws. Each draw induces a corresponding ballistic-coefficient sample $BC_s^{(q_{\text{pp}})}$ and reciprocal drag-coupling sample $B_s^{(q_{\text{pp}})} = 1/BC_s^{(q_{\text{pp}})}$ through the retained effective-area parameterization. For each posterior-predictive draw, the forecast-conditioned forward model integrates from the initial epoch over $\mathcal{T}_s^{\text{prop}}$, producing a modeled state history

$$\left(\mathbf{r}_s^{(q_{\text{pp}})}(t), \mathbf{v}_s^{(q_{\text{pp}})}(t) \right), \quad t \in \mathcal{T}_s^{\text{prop}} \quad (5.33)$$

and the corresponding semi-major-axis history

$$a_s^{(q_{\text{pp}})}(t) \quad (5.34)$$

The posterior-predictive semi-major-axis ensemble is therefore written as

$$\mathcal{A}_s^{\text{pp}} = \left\{ a_s^{(q_{\text{pp}})}(t) : q_{\text{pp}} = 1, \dots, N_{\text{pp},s}, t \in \mathcal{T}_s^{\text{prop}} \right\} \quad (5.35)$$

This ensemble is generated using a common force model and a consistent forecast-conditioned density realization, with drag coupling modulated by the posterior uncertainty in BC_s^{eff} [164, 2, 12].

This propagation step establishes the connection between inference and simulation. The posterior distribution over BC_s^{eff} defines a distribution over plausible drag responses, which is mapped into state variables relevant to late-life Starlink decay. The posterior-predictive ensemble is

the simulation output used for subsequent uncertainty quantification, with the parameter posterior serving as the generative input to the forward model [176, 164].

5.6.2 Semi-Major Axis Trajectory Fan

The posterior-predictive trajectories display a fan-like structure in semi-major-axis space. During the initial propagation interval, the ensemble remains tightly clustered, indicating that posterior spread in BC_s^{eff} has limited dynamical influence over the earliest portion of the propagated arc. As orbital decay progresses and the drag regime intensifies, differences in the sampled drag response accumulate through repeated aerodynamic energy dissipation, leading to broader ensemble spread and diverging trajectories. This fan structure demonstrates the direct dynamical impact of posterior uncertainty within a nonlinear dissipative model [177, 12, 157].




Figure 5.4: Posterior-Predictive Semi-Major Axis Trajectories

Figure 5.4 shows the posterior-predictive semi-major-axis fan produced by propagating retained BC_s^{eff} samples through the forecast-conditioned forward model. The ensemble remains compact early in the arc and then spreads rapidly as aerodynamic dissipation strengthens near terminal descent. This behavior is consistent with the cumulative nature of atmospheric drag, which depends on thermospheric density, relative velocity, and effective aerodynamic loading [158, 177, 12]. Initially, small differences in $BC_s^{(q_{pp})}$ produce limited state separation, but these differences become increasingly consequential as the orbit reaches denser thermospheric layers. The trajectory

fan therefore shows that the inferred BC_s^{eff} posterior encodes simulation-relevant information. A narrow parameter posterior can still produce meaningful predictive dispersion in state space because orbital sensitivity to drag-response uncertainty increases near terminal descent.

5.6.3 Re-Entry Timing Uncertainty

The distribution of terminal descent times is obtained by mapping the posterior-predictive trajectory ensemble onto the altitude threshold. For posterior-predictive draw q_{pp} , the operational re-entry time for satellite s is defined as

$$T_{\text{re},s}^{(q_{\text{pp}})} = \inf \left\{ t \in \mathcal{T}_s^{\text{prop}} : h_s^{(q_{\text{pp}})}(t) \leq h_{\text{stop}} \right\} \quad (5.36)$$

where

$$h_s^{(q_{\text{pp}})}(t) = \left\| \mathbf{r}_s^{(q_{\text{pp}})}(t) \right\| - R_{\oplus} \quad (5.37)$$

is the propagated altitude history, and h_{stop} is the terminal altitude threshold introduced previously. The posterior-predictive re-entry-time distribution is represented by the sample set

$$\mathcal{T}_{\text{re},s}^{\text{pp}} = \left\{ T_{\text{re},s}^{(q_{\text{pp}})} : q_{\text{pp}} = 1, \dots, N_{\text{pp},s} \right\} \quad (5.38)$$

The re-entry window $\mathcal{W}_{\text{re},s}$ is then obtained by summarizing $\mathcal{T}_{\text{re},s}^{\text{pp}}$ through posterior quantiles or interval bounds, depending on the reporting convention used for the simulation product [164, 172].

This construction makes the ballistic-coefficient posterior operationally explicit. Variations in $BC_s^{(q_{\text{pp}})}$ alter drag coupling, modify the semi-major-axis decay history, and shift threshold-crossing times. Predictive re-entry-time uncertainty therefore arises as the cumulative dynamical outcome of uncertainty in the effective drag parameter. Deterministic late-life simulations cannot represent this dispersion, because nonlinear drag accumulation can produce substantial terminal descent-time spread even when the inferred BC_s^{eff} posterior is comparatively concentrated [177, 12, 157, 178].

Chapter 6

Physics-Constrained System Identification of Low-Thrust Dynamics

6.1 Mean-Element Observation Model

The orbit-level inverse problem is formulated using public catalog histories, which consist of time-stamped GP records, including TLE records and equivalent GP products, together with externally assigned maneuver-phase labels [16, 17, 160, 88, 20]. The objective is to model the evolution of catalog mean-element summaries through a reduced, phase-conditioned dynamical model parameterized by effective drag and actuation variables. Both the data product and the forward model are maintained within the mean-element domain, so the estimated unknowns retain the same observational semantics as the public catalog layer [16, 17, 160]. This construction defines the mean-element observation model (MEOM) used for the low-thrust system-identification problem.

6.1.1 Catalog Mean-Element Observation Model

Let $s \in \mathcal{S}$ denote a satellite in the catalog population, and let t_k denote the k th catalog epoch associated with that object. The catalog mean-element record for satellite s at epoch t_k is written as

$$\mathbf{y}_{s,k}^{\text{cat}} = (\bar{n}_{s,k}, e_{s,k}, i_{s,k}, \Omega_{s,k}, \omega_{s,k}, M_{s,k}) \quad (6.1)$$

where $\bar{n}_{s,k}$ denotes catalog mean motion, $e_{s,k}$ is eccentricity, $i_{s,k}$ is inclination, $\Omega_{s,k}$ is the RAAN, $\omega_{s,k}$ is the argument of perigee, and $M_{s,k}$ is the mean anomaly. The overbar on $\bar{n}_{s,k}$ emphasizes catalog provenance and distinguishes the catalog-reported mean motion from any modeled or fitted mean-motion variable introduced later in the system-identification formulation. These quantities

are interpreted as catalog-inferred mean elements optimized for compatibility with SGP4 analytical orbit models, with direct interpretation constrained by the GP observation layer rather than by numerical osculating-state propagation [16, 17, 160, 88, 20].

In this inverse problem, the raw catalog fields are mapped to a reduced observation vector that emphasizes secular and slowly varying structure relevant to low-thrust and drag-dominated evolution. The principal derived observable is a catalog-consistent semi-major axis proxy,

$$a_{s,k}^{\text{cat}} = \left(\frac{\mu}{\bar{n}_{s,k}^2} \right)^{1/3} \quad (6.2)$$

where μ is the Earth gravitational parameter and $\bar{n}_{s,k}$ is expressed in radians per second [2, 1, 17, 160]. This quantity is treated as a catalog-consistent semi-major axis proxy, with interpretive fidelity tied to the public mean-element archive and its associated propagation conventions.

A second key observable is the catalog mean-longitude proxy,

$$\lambda_{s,k} = \text{wrap}_{2\pi} (\Omega_{s,k} + \omega_{s,k} + M_{s,k}) \quad (6.3)$$

where $\text{wrap}_{2\pi}(\cdot)$ maps angular quantities onto a 2π -periodic domain. The proxy $\lambda_{s,k}$, together with $\Omega_{s,k}$, provides an along-track-sensitive summary compatible with the catalog mean-element representation [2, 1]. Accordingly, the reduced catalog observation vector is defined as

$$\mathbf{z}_{s,k} = \begin{bmatrix} a_{s,k}^{\text{cat}} \\ \Omega_{s,k} \\ \lambda_{s,k} \end{bmatrix} \quad (6.4)$$

Finite-interval differences and rates are constructed from $\mathbf{z}_{s,k}$ over labeled maneuver arcs, with angular components unwrapped before differencing when required.

The basic observational unit is a phase-labeled trajectory segment or arc. Let r denote an arc index, let $p_r \in \mathcal{P}$ denote the maneuver-phase label assigned to that arc, and let $\mathcal{P}$ denote the admissible phase-label set. For arc r , with observation times $t_{r,0}, \dots, t_{r,N_r-1}$, define the arc-local index set

$$\mathcal{K}_r = \{0, \dots, N_r - 1\} \quad (6.5)$$

The phase-labeled catalog trajectory segment is then written as

$$\mathcal{Y}_r = \left\{ \left(t_{r,k}, a_{r,k}^{\text{cat}}, \Omega_{r,k}, \lambda_{r,k} \right) \right\}_{k \in \mathcal{K}_r} \quad (6.6)$$

This representation preserves the temporal ordering required to distinguish insertion, operational maintenance, and disposal phases, while remaining consistent with the information content of the catalog record.

For adjacent samples within arc r , the finite interval and reduced observation increment are defined as

$$\Delta t_{r,k} = t_{r,k+1} - t_{r,k}, \quad k = 0, \dots, N_r - 2 \quad (6.7)$$

and

$$\Delta \mathbf{z}_{r,k} = \mathbf{z}_{r,k+1} - \mathbf{z}_{r,k} \quad (6.8)$$

These increments provide the observational basis for subsequent phase-conditioned system identification, where the modeled secular response of a^{cat} , Ω , and λ is related to effective drag and low-thrust actuation parameters.

6.1.2 Phase-Conditioned Effective Parameterization

Unknown quantities are modeled as effective parameters indexed by mission phase, rather than as directly recoverable internal thruster states. The admissible set of mission phases is defined as

$$\mathcal{P} = \{\text{ins}, \text{ops}, \text{disp}\} \quad (6.9)$$

where ins denotes insertion or orbit raising, ops denotes operational shell maintenance, and disp denotes disposal lowering. These labels define the phase variable $p_r \in \mathcal{P}$ assigned to each trajectory arc $\mathcal{Y}_r$. The latent effective parameter vector is partitioned into a global block and phase-conditioned blocks,

$$\boldsymbol{\theta} = (\boldsymbol{\theta}_g, \boldsymbol{\theta}_{\text{ins}}, \boldsymbol{\theta}_{\text{ops}}, \boldsymbol{\theta}_{\text{disp}}) \quad (6.10)$$

The global block $\boldsymbol{\theta}_g$ contains quantities shared across all phases, while $\boldsymbol{\theta}_p$ contains phase-specific effective quantities for $p \in \mathcal{P}$.

At the reduced level considered here, the global block contains parameters common to all mission phases, including an effective satellite mass scale and an effective specific impulse. Each phase block governs the net mean-element evolution for its corresponding maneuver regime. In the compact posterior formulation developed later, this structure yields

$$\boldsymbol{\theta} = \left[m_{\text{eff}}, I_{\text{sp}}, T_{\text{ins}}, D_{\text{ins}}, T_{\text{ops}}, D_{\text{ops}}, T_{\text{disp}}, D_{\text{disp}} \right]^{\top} \quad (6.11)$$

where m_{eff} denotes the effective mass scale, I_{sp} represents the effective specific impulse, T_p is

the phase-level effective actuation scale, and D_p is the phase-level effective drag scale for each phase $p \in \mathcal{P}$. The quantities T_p and D_p are not interpreted as directly measured thrust or drag coefficients. These are aggregate parameters that summarize the catalog-observable contribution of actuation-like and drag-like effects to the reduced mean-element evolution [154, 155, 156, 83, 179].




Figure 6.1: Mission Phases for Gen 1 Starlink TLE History

Figure 6.1 shows the phase-labeled Gen 1 Starlink catalog history used to define insertion or orbit raising, operational shell maintenance, and disposal lowering intervals. The figure illustrates that the inverse problem is conditioned on a temporally structured maneuver archive rather than on an undifferentiated sequence of catalog states.

This parameterization balances operational utility with the limitations of public catalog data. Public mean-element archives do not provide thrust vector histories, duty cycles, attitude histories, cross-sectional geometry histories, electrical currents, discharge voltages, mass-flow rates, or plasma-state variables [17, 88, 20, 83]. The reduced parameters are therefore defined at an aggregation level directly supported by the data, capturing effective combinations that induce catalog-compatible changes in semi-major axis and angular mean-element summaries over labeled intervals. This level of granularity is sufficient to distinguish insertion, shell maintenance, and disposal regimes in a dynamically consistent manner without assigning physical specificity unsupported by the observation layer.

The phase-conditioned forward model is expressed abstractly as

$$\widehat{\mathcal{Y}}_r = \mathcal{M}\left(\boldsymbol{\theta}, p_r, \mathbf{x}_{r,0}^{\text{mean}}, \{t_{r,k}\}_{k \in \mathcal{K}_r}\right) \quad (6.12)$$

where $\widehat{\mathcal{Y}}_r$ denotes the predicted catalog-compatible trajectory segment for arc r , $\mathbf{x}_{r,0}^{\text{mean}}$ denotes the initial mean-state estimate for that arc, and $\mathcal{M}$ is a reduced phase-conditioned dynamical map. The map $\mathcal{M}$ returns predicted catalog-compatible summaries rather than full osculating trajectories. Consequently, the inverse problem is posed on phase-conditioned mean-element evolution, which is the appropriate observational level for fleet-scale behavioral modeling with public catalog records [16, 17, 160].

6.1.3 Inference Scope and Effective Interpretation

The inferential objective is to obtain a low-dimensional representation of drag-like and actuation-like effects that explains the observed mean-element trajectories. This formulation estimates effective phase-dependent interactions observable in the catalog and suitable for forward propagation in satellite behavior models, while leaving internal Hall-thruster plasma states, commanded thrust histories, and the full osculating force environment outside the identifiable parameter space [17, 83, 179, 180].

Formally, the observed arc is modeled as the output of a reduced dynamical system composed with a catalog-compatible observation map,

$$\mathcal{Y}_r = \mathcal{H}\left(\mathcal{M}\left(\boldsymbol{\theta}, p_r, \mathbf{x}_{r,0}^{\text{mean}}, \{t_{r,k}\}_{k \in \mathcal{K}_r}\right)\right) + \boldsymbol{\varepsilon}_r \quad (6.13)$$

where $\mathcal{H}$ extracts the mean-element summaries represented in the public archive, and $\boldsymbol{\varepsilon}_r$ accounts for catalog error, phase-label mismatch, model discrepancy, and unmodeled forcing. Identifiability is governed by the composite operator $\mathcal{H} \circ \mathcal{M}$. The recovered parameters are therefore interpreted as effective quantities that summarize the observable low-thrust and drag behavior for each phase [17, 160, 154, 155, 156].

This interpretation enables analysis at the satellite constellation scale. Stage A produces a deterministic, physics-constrained trajectory fit. Stage B constructs an approximate posterior over reduced parameters using simulation-based inference (SBI), yielding

$$\pi(\boldsymbol{\theta} \mid \mathcal{Y}) \quad (6.14)$$

where $\mathcal{Y}$ denotes the selected collection of catalog-compatible trajectory summaries used for in-

ference [181, 182, 183]. The resulting posterior is a distribution over catalog-compatible effective parameters conditioned on the selected summaries and reduced dynamics. Although the full internal satellite state remains unresolved, the posterior quantifies the phase-conditioned drag and actuation combinations consistent with observed mean-element histories. This information is directly applicable to regime-specific propagation, behavior classification, and satellite constellation-scale dynamical modeling.

6.2 Data Construction, Segmentation, and Observable Structure

The observational pipeline utilizes historical TLE records for the Starlink Gen 1 satellite constellation in conjunction with an independent archive of maneuver-phase interval labels. The raw catalog stream provides epoch-indexed mean-element vectors for each tracked object, while the label archive partitions this history into operational regimes defined by distinct dynamical signatures. Data construction requires precise alignment of two heterogeneous archives: the public catalog orbital-fit archive and the curated phase-label archive. Let these archives be denoted by $\mathcal{D}_{\text{cat}}$ and $\mathcal{D}_{\text{lab}}$, respectively. The resulting aligned dataset is transformed into learning objects appropriate for reduced trajectory modeling and simulator-based posterior inference [16, 17, 88, 20, 181].

The constructed dataset consists of $N_{\text{seg}} = 18,952$ phase-labeled two-point segments spanning $N_{\text{sat}} = 500$ satellites and three maneuver phases. For trajectory fitting, phase-labeled intervals are further converted into $N_{\text{arc}} = 12,209$ multi-epoch arcs with explicit temporal ordering. Segment and arc representations address distinct analytical requirements. The segment view supports summary statistics and Stage B posterior inference, while the arc view maintains within-interval temporal structure required for Stage A trajectory matching [184, 181].

6.2.1 TLE Ingestion, Phase Intervals, and Trajectory-Arc Construction

Each TLE record is parsed into a time-stamped mean-element vector that includes catalog mean motion, eccentricity, inclination, RAAN, argument of perigee, and mean anomaly [16, 17, 160, 88, 20]. These parameters are mapped to the reduced catalog observation vector $\mathbf{z}_{s,k}$, with particular emphasis on the semi-major axis proxy $a_{s,k}^{\text{cat}}$, the nodal angle $\Omega_{s,k}$, and the mean-longitude proxy,

$$\lambda_{s,k} = \text{wrap}_{2\pi} (\Omega_{s,k} + \omega_{s,k} + M_{s,k}) \quad (6.15)$$

This mapping produces a unified, catalog-consistent state representation for all satellites before phase conditioning [2, 1, 185].

Parsed TLE histories are subjected to limited smoothing to suppress isolated outliers while preserving secular and slowly varying trends that are critical for inverse modeling. The resulting time series are then sorted by satellite identifier and epoch to prepare for phase-interval label alignment [186, 167].

The phase-label archive defines intervals $[t_{\text{start},j}, t_{\text{end},j}]$ with associated phase states, including insertion, orbit raising, operational shell maintenance, and disposal lowering. These phase states are mapped onto the admissible phase-label set, where for each labeled interval, the closest matching TLE epochs within an alignment tolerance τ_{align} are matched to the interval boundaries. This procedure forms a two-point segment from the initial and final catalog states [187, 17, 160].

Let $j \in \mathcal{J}$ denote a segment index. Segment j for satellite s and phase $p_j \in \mathcal{P}$ is defined as

$$\mathcal{G}_j = (s, p_j, t_{j,0}, t_{j,1}, \mathbf{z}_{j,0}, \mathbf{z}_{j,1}) \quad (6.16)$$

where $\mathcal{G}_j$ denotes the phase-labeled two-point segment object, $t_{j,0}$ and $t_{j,1}$ are the aligned initial and final catalog epochs, and $\mathbf{z}_{j,0}$ and $\mathbf{z}_{j,1}$ denote the corresponding reduced mean-element vectors. The notation $\mathcal{G}_j$ is used for segment objects so that $\mathcal{S}$ remains reserved for the satellite index set.

The primary segment-level increments are defined as

$$\Delta a_j^{\text{cat}} = a_{j,1}^{\text{cat}} - a_{j,0}^{\text{cat}}, \quad \Delta \Omega_j = \text{wrap}_{2\pi}(\Omega_{j,1} - \Omega_{j,0}), \quad \Delta \lambda_j = \text{wrap}_{2\pi}(\lambda_{j,1} - \lambda_{j,0}) \quad (6.17)$$

with the corresponding elapsed duration

$$\Delta t_j = t_{j,1} - t_{j,0} \quad (6.18)$$

When rate-normalized summaries are required, the finite-interval rates are

$$\dot{a}_j^{\text{cat}} = \frac{\Delta a_j^{\text{cat}}}{\Delta t_j}, \quad \dot{\Omega}_j = \frac{\Delta \Omega_j}{\Delta t_j}, \quad \dot{\lambda}_j = \frac{\Delta \lambda_j}{\Delta t_j} \quad (6.19)$$

The angular increments are evaluated on the periodic domain before rate normalization, consistent with circular-variable treatment of angular data [185]. The compact segment-level increment vector

is written as

$$\mathbf{q}_j = \begin{bmatrix} \Delta a_j^{\text{cat}} \\ \Delta \Omega_j \\ \Delta \lambda_j \\ \Delta t_j \end{bmatrix} \quad (6.20)$$

and the corresponding rate vector is written as

$$\mathbf{u}_j = \begin{bmatrix} \dot{a}_j^{\text{cat}} \\ \dot{\Omega}_j \\ \dot{\lambda}_j \end{bmatrix} \quad (6.21)$$

These increments, durations, and rate summaries constitute the compact observational units used for segment-based diagnostics and simulator-driven posterior inference. The collection of such segment summaries is denoted by $\mathcal{D}_{\text{seg}}$ [184, 181].

Stage A analysis requires a more granular structure than a two-point segment, since its primary objective is to capture semi-major axis evolution within a labeled interval. Each admissible phase interval is therefore converted into a trajectory arc comprising all TLE observations within the interval, filtered by observation count and minimum duration. Let $\mathcal{R}$ denote the admissible arc index set. Arc $r \in \mathcal{R}$ is constructed as

$$\mathcal{Y}_r = \left\{ \left(t_{r,k}, a_{r,k}^{\text{cat}}, \Omega_{r,k}, \lambda_{r,k} \right) \right\}_{k \in \mathcal{K}_r}, \quad \mathcal{K}_r = \{0, \dots, N_r - 1\}, \quad N_r \geq N_{\min} \quad (6.22)$$

where $N_{\min} = 5$ in the present construction. When a duration filter is imposed, an admissible arc also satisfies

$$t_{r,N_r-1} - t_{r,0} \geq \Delta t_{\min} \quad (6.23)$$

The arc retains the phase label $p_r \in \mathcal{P}$, maintains internal temporal ordering, facilitates the definition of path-loss terms on $a^{\text{cat}}(t)$, and enables multiple-shooting continuity constraints for extended intervals. The collection of admissible trajectory arcs is denoted by $\mathcal{D}_{\text{arc}}$.

Segment and arc representations function as complementary constructs. The segment view reduces each labeled interval to finite increments and rates, while the arc view encodes the within-interval trajectory geometry required for Stage A fitting. This dual representation allows the same catalog and phase-label archives to support both compact posterior summaries and temporally resolved trajectory matching [184, 181].

6.2.2 Phase Imbalance and Segment-Duration Structure

The labeled archive demonstrates substantial phase imbalance. Let the phase-specific segment subset be defined as

$$\mathcal{J}_p = \{j \in \mathcal{J} : p_j = p\}, \quad p \in \mathcal{P} \quad (6.24)$$

with corresponding phase-specific segment count and empirical phase fraction

$$N_{\text{seg},p} = |\mathcal{J}_p|, \quad \rho_{\text{seg},p} = \frac{N_{\text{seg},p}}{N_{\text{seg}}} \quad (6.25)$$

Operational-shell segments are the most prevalent, disposal-lowering segments constitute a notable secondary category, and insertion or orbit-raise segments are comparatively rare. This distribution reflects the lifecycle structure of the satellite constellation, since satellites primarily occupy long-duration shell-maintenance regimes while orbit raising and disposal lowering occur over shorter transitional intervals.

This imbalance has direct implications for modeling. Aggregate error metrics are weighted toward operational-shell behavior, and global fit statistics are most strongly influenced by the regime with the greatest empirical support. Because shell maintenance dominates long-term satellite constellation evolution, this statistical concentration increases the operational relevance of the fitted model for persistent fleet behavior. Insertion maneuvers are comparatively rare and dynamically abrupt, thus contributing disproportionately to the tail of the error distribution despite limited representation in the labeled archive.




Figure 6.2: Segment Counts by Maneuver Phase

Figure 6.2 quantifies this phase imbalance, showing that operational-shell segments dominate

the labeled archive, disposal-lowering segments form the principal secondary class, and insertion or orbit-raise segments represent the sparsest regime.

The segment-duration sample set is defined as

$$\mathcal{D}_{\Delta t} = \{\Delta t_j : j \in \mathcal{J}\} \quad (6.26)$$

Empirically, this distribution is strongly right-skewed. Most labeled intervals are short or moderate in duration, while a sparse tail extends to longer intervals. The reduced model must therefore accommodate localized dynamics within compact intervals and accumulated error, label impurity, and perturbation effects over extended segments.

The observed semi-major axis increment sample set is defined as

$$\mathcal{D}_{\Delta a} = \{\Delta a_j^{\text{cat}} : j \in \mathcal{J}\} \quad (6.27)$$

The distribution of Δa_j^{cat} is tightly centered near zero, with a dense core and infrequent large deviations. This pattern reflects the dominance of shell-maintenance intervals, where net changes in catalog-consistent semi-major axis are small compared with the shifts associated with orbit raising or disposal lowering. The archive, therefore, contains a majority of near-equilibrium intervals and a smaller number of maneuver-rich, dynamically challenging segments, supporting high-fidelity learning of dominant shell-maintenance behavior while retaining evidence of transitional regimes.




Figure 6.3: Distribution of Segment Duration and Observed Semi-Major Axis Increment

Figure 6.3 provides the empirical basis for this interpretation by showing the right-skewed segment-duration distribution together with the sharply concentrated semi-major-axis increment distribution.

6.2.3 Observable Channels for Stage A and Stage B

Stage A is formulated as a trajectory-matching problem on phase-labeled arcs. Its dominant observable is the catalog-consistent semi-major axis path

$$\mathbf{a}_r^{\text{cat}} = \begin{bmatrix} a_{r,0}^{\text{cat}} & a_{r,1}^{\text{cat}} & \dots & a_{r,N_r-1}^{\text{cat}} \end{bmatrix}^\top \quad (6.28)$$

evaluated at all available internal epochs within arc r . Secondary endpoint diagnostics are formed from the wrapped nodal and mean-longitude differences,

$$\Delta\Omega_r = \text{wrap}_{2\pi}(\Omega_{r,N_r-1} - \Omega_{r,0}), \quad \Delta\lambda_r = \text{wrap}_{2\pi}(\lambda_{r,N_r-1} - \lambda_{r,0}) \quad (6.29)$$

The corresponding angular endpoint diagnostic vector is

$$\boldsymbol{\gamma}_r = \begin{bmatrix} \Delta\Omega_r \\ \Delta\lambda_r \end{bmatrix} \quad (6.30)$$

The Stage A objective is therefore trajectory-rich in the semi-major axis channel and endpoint-based in the angular channels. This asymmetry is intentional: the implemented reduced model is designed primarily to match the path geometry of $a^{\text{cat}}(t)$, with angular closure used as a secondary diagnostic on nodal and along-track consistency [1, 2].

Stage B operates on a reduced summary representation built from segment increments and duration-normalized rates. For each segment j , define the six-dimensional Stage B summary vector

$$\boldsymbol{\eta}_j = \begin{bmatrix} \Delta a_j^{\text{cat}} \\ \Delta\Omega_j \\ \Delta\lambda_j \\ \dot{a}_j^{\text{cat}} \\ \dot{\Omega}_j \\ \dot{\lambda}_j \end{bmatrix} \quad (6.31)$$

where

$$\dot{a}_j^{\text{cat}} = \frac{\Delta a_j^{\text{cat}}}{\Delta t_j}, \quad \dot{\Omega}_j = \frac{\Delta\Omega_j}{\Delta t_j}, \quad \dot{\lambda}_j = \frac{\Delta\lambda_j}{\Delta t_j} \quad (6.32)$$

The Stage B summary dataset is then written as

$$\mathcal{X}_B = \{(\boldsymbol{\eta}_j, p_j) : j \in \mathcal{J}\} \quad (6.33)$$

This summary dataset defines the core observation map used by the simulator-based posterior stage [184, 181, 182, 183].

The two-stage design is deliberately complementary. Stage A uses the richer trajectory object $\mathcal{Y}_r$ to learn the dominant semi-major axis behavior of each maneuver phase, while Stage B uses the compressed summary $\boldsymbol{\eta}_j$ to quantify uncertainty over the corresponding effective parameters at satellite constellation scale. The compression discards within-arc path information, especially in angular channels and intermediate semi-major axis curvature, but it enables computationally practical phase-conditioned posterior inference over a large labeled archive [184, 181].

6.3 Reduced Dynamical and Propulsion Model

The forward model incorporates a reduced mean-element trajectory map, a phase-conditioned actuation law, a krypton Hall thruster beam closure, and a surrogate chemistry regularizer. This architecture applies deliberately uneven physical fidelity across model layers. The trajectory layer is designed to reproduce catalog-scale semi-major axis evolution under low-thrust and drag perturbations. The propulsion layer constrains the latent parameter space to physically plausible domains while leaving internal discharge physics outside the resolved state. The model serves three primary functions: orbit-level state evolution, effective actuation parameterization, and feasibility regularization over latent propulsion variables [2, 1, 158, 83, 179, 188, 180].

6.3.1 Mean-Element Trajectory and Phase-Conditioned Net Acceleration

For a labeled arc r in phase p_r , the model assigns a signed net tangential acceleration that drives secular semi-major axis evolution. Since a is reserved for semi-major axis, acceleration quantities are denoted using α . Let T_r denote the arc-level effective thrust magnitude, with $T_r = T_{p_r}$ when thrust is conditioned only on mission phase. Let d_r denote the effective duty factor after power-cap scaling, m_{eff} the effective satellite mass scale, and $\alpha_{D,r}$ the drag deceleration magnitude. The implemented net tangential acceleration is

$$\alpha_{\text{net},r} = \alpha_{T,r} - \alpha_{D,r} \quad (6.34)$$

with

$$\alpha_{T,r} = \chi_r d_r \frac{T_r}{1000 m_{\text{eff}}} \quad (6.35)$$

where T_r is represented in newtons, m_{eff} is represented in kilograms, and the factor 1000 converts the resulting acceleration from m s^{-2} to km s^{-2} for consistency with semi-major axis expressed in kilometers. The coefficient χ_r is a dimensionless signed actuation-direction coefficient [2, 1].

The actuation-direction coefficient depends on the maneuver phase. Let

$$\sigma_{p_r} \in \{-1, 0, +1\} \quad (6.36)$$

denote the phase-sign coefficient for arc r , where the zero-sign case corresponds to operational shell maintenance. For insertion and disposal, χ_r is a signed quantity built from the phase sign and a constrained direction-strength parameter. For operational-shell arcs, χ_r is replaced by a drag-compensation coefficient that permits near-balanced shell maintenance. The implemented form is

$$\chi_r = \begin{cases} f_{\text{shell},r}, & |\sigma_{p_r}| < 0.5 \\ \sigma_{p_r} \text{ clip}_{[0.25,1.0]}(\text{softplus}(\zeta_r) + 0.25), & |\sigma_{p_r}| \geq 0.5 \end{cases} \quad (6.37)$$

where $f_{\text{shell},r}$ is the shell drag-compensation fraction, ζ_r is a learned direction-strength latent variable, and

$$\text{softplus}(x) = \log(1 + \exp x) \quad (6.38)$$

is a smooth positive transform. This formulation treats operational-shell behavior as a drag-balance control regime and treats insertion and disposal arcs as signed low-thrust maneuver regimes.

The duty factor is constrained by an electrical power cap. If $P_{\text{nom},r}$ denotes the nominal electrical power implied by T_r , I_{sp} , and the total effective thruster efficiency $\eta_{\text{tot},r}$, and if $P_{\text{cap},r}$ denotes the phase-level power limit, the implementation uses

$$d_r = \text{clip}_{[10^{-4}, d_{\text{max}}]} \left[d_{0,r} \min \left(\frac{P_{\text{cap},r}}{\max(P_{\text{nom},r}, 1)}, 1 \right) \right] \quad (6.39)$$

where $d_{0,r}$ is the raw phase-conditioned duty factor and d_{max} is the thermal or operational duty ceiling. The nominal power relation used internally is

$$P_{\text{nom},r} = \frac{T_r g_0 I_{\text{sp}}}{2 \eta_{\text{tot},r}} \quad (6.40)$$

where g_0 is standard gravitational acceleration. This relation follows from the usual beam-power scaling in which exhaust power is proportional to thrust multiplied by effective exhaust velocity $g_0 I_{\text{sp}}$, with total efficiency mapping input electrical power to directed jet power [83, 179]. The actuation level is reduced whenever the parameterization predicts nominal power beyond the imposed

phase-level cap.

The default reduced propagation law uses a Gauss-variation-of-parameters-style semi-major axis update with constant net tangential acceleration over an arc segment. If $a_{r,0}^{\text{cat}}$ is the initial catalog-consistent semi-major axis and

$$n_{r,0} = \sqrt{\frac{\mu}{\left(a_{r,0}^{\text{cat}}\right)^3}} \quad (6.41)$$

is the corresponding initial mean motion, then the trajectory-level prediction is

$$\widehat{a}_r(t) = a_{r,0}^{\text{cat}} + \frac{2 \alpha_{\text{net},r}}{n_{r,0}} (t - t_{r,0}) \quad (6.42)$$

This reduced dynamical model forms the basis of the trajectory-matching code, providing a computationally efficient, differentiable, and observation-aligned propagation scheme [2, 1].

A higher-fidelity option is also implemented for nonlinear propagation with altitude-dependent drag. In that mode, the model integrates

$$\frac{d\widehat{a}_r}{dt} = \frac{2}{n(\widehat{a}_r)} (\alpha_{T,r}(\widehat{a}_r, t) - \alpha_{D,r}(\widehat{a}_r, t)), \quad n(\widehat{a}_r) = \sqrt{\frac{\mu}{\widehat{a}_r^3}} \quad (6.43)$$

A classical fourth-order Runge-Kutta (RK4) integrator operates on a uniform internal time grid, and the resulting modeled states are interpolated to the TLE observation epochs. The drag term is recalculated at each instantaneous altitude using a USSA76 table-based density model, which steepens the decay response as altitude decreases. This nonlinear mode captures altitude-dependent drag and the variation of mean motion with semi-major axis, both of which are suppressed in the linear constant-acceleration approximation [80, 74, 177].

The orbit-averaged drag deceleration used in the nonlinear path is written as

$$\alpha_D(\widehat{a}_r) = \frac{1}{2} \beta_{\text{ref}} \rho(h) v^2(\widehat{a}_r), \quad h = \widehat{a}_r - R_{\oplus} \quad (6.44)$$

where $\rho(h)$ is the tabulated atmospheric density, $R_{\oplus}$ is the reference Earth radius, and

$$v(\widehat{a}_r) = \sqrt{\frac{\mu}{\widehat{a}_r}} \quad (6.45)$$

is the near-circular orbital speed approximation. The reference inverse ballistic coefficient is

$$\beta_{\text{ref}} = \frac{C_D A_{\text{eff}}}{m_{\text{eff}}} = \frac{1}{BC_{\text{ref}}}, \quad BC_{\text{ref}} = \frac{m_{\text{eff}}}{C_D A_{\text{eff}}} \quad (6.46)$$

which preserves the ballistic-coefficient convention introduced earlier. In the learned model, this baseline drag is multiplied by phase-level and satellite-level drag scales,

$$\alpha_{D,r}(\widehat{a}_r) = \kappa_{p_r}^D \kappa_{s(r)}^D \frac{1}{2} \beta_{\text{ref}} \rho(h) v^2(\widehat{a}_r) \quad (6.47)$$

where $\kappa_{p_r}^D$ is the phase-level multiplicative drag scale, $\kappa_{s(r)}^D$ is the satellite-level multiplicative drag scale, and $s(r)$ denotes the satellite associated with arc r . This parameterization functions as an effective drag law that absorbs catalog mismatch, geometry uncertainty, and residual force-model error into identifiable reduced parameters [158, 12, 38, 159].

The angular channels are propagated only at arc endpoints and only through second-zonal-harmonic secular rates. Let $\widehat{a}_{r,f}$ be the predicted final semi-major axis, let

$$\widehat{a}_{r,m} = \frac{a_{r,0}^{\text{cat}} + \widehat{a}_{r,f}}{2} \quad (6.48)$$

be the midpoint approximation, and let

$$\Delta t_r = t_{r,N_r-1} - t_{r,0} \quad (6.49)$$

be the total arc duration. The nodal endpoint is computed as

$$\widehat{\Omega}_{r,f} = \Omega_{r,0} + \dot{\Omega}_{J_2}(\widehat{a}_{r,m}, e_{r,0}, i_{r,0}) \Delta t_r \quad (6.50)$$

while the mean-longitude endpoint uses a trapezoidal average of the secular rate,

$$\widehat{\lambda}_{r,f} = \lambda_{r,0} + \frac{1}{2} \left[\dot{\lambda}_{J_2}(a_{r,0}^{\text{cat}}, e_{r,0}, i_{r,0}) + \dot{\lambda}_{J_2}(\widehat{a}_{r,m}, e_{r,0}, i_{r,0}) \right] \Delta t_r \quad (6.51)$$

These endpoint channels provide diagnostics derived from the reduced semi-major axis trajectory. Fit quality is optimized primarily in the semi-major axis evolution, with angular endpoint channels used to assess weaker nodal and along-track consistency [2, 1, 3].

6.3.2 Krypton Beam Relations and Effective Propulsion Variables

The propulsion module implements a low-order krypton Hall thruster (HT) closure, formalized through canonical one-dimensional beam relations. Let V_b denote beam voltage, I_b denote beam current, γ denote a current-utilization factor, and η_m denote mass-utilization efficiency. The reduced krypton beam-law map is denoted by $\mathcal{B}_{\text{Kr}}$, with thrust and specific impulse approximated as

$$T_{\text{beam}} = C_T^{\text{Kr}} \gamma I_b \sqrt{V_b}, \quad I_{\text{sp,beam}} = C_I^{\text{Kr}} \gamma \eta_m \sqrt{V_b} \quad (6.52)$$

Here, T_{beam} denotes the beam-law thrust magnitude, $I_{\text{sp,beam}}$ denotes the beam-law specific impulse, and C_T^{Kr} and C_I^{Kr} are krypton-specific coefficients determined by the representative krypton ion mass m_{Kr} , the elementary charge q_e , and the one-dimensional beam approximation. For a singly charged krypton-ion surrogate, these coefficients may be written in the idealized form

$$C_T^{\text{Kr}} = \sqrt{\frac{2m_{\text{Kr}}}{q_e}}, \quad C_I^{\text{Kr}} = \frac{1}{g_0} \sqrt{\frac{2q_e}{m_{\text{Kr}}}} \quad (6.53)$$

with utilization factors absorbing departures from the ideal beam law. The effective exhaust velocity follows as

$$v_e = g_0 I_{\text{sp,beam}} \quad (6.54)$$

where v_e is the effective exhaust velocity and g_0 is standard gravitational acceleration [83, 179, 188].

The discharge or electrical input-power proxy is expressed as

$$P_d \approx \frac{V_b I_b}{\eta_b \eta_v} \quad (6.55)$$

where P_d is the discharge-power proxy, η_b denotes a beam-current or current-utilization efficiency factor, and η_v denotes voltage utilization. The total effective thruster efficiency is factorized as

$$\eta_{\text{tot}} = \eta_b \eta_v \eta_m \eta_o \quad (6.56)$$

where η_o is an aggregate residual efficiency factor representing divergence, facility, and unresolved loss mechanisms [83, 179, 188, 180].

The closure addresses the ill-posedness of the reduced low-thrust inverse problem by constraining the latent actuation variables. Beam relations define a physically interpretable manifold within the latent parameter space, preventing the trajectory fit from selecting arbitrarily high-thrust,

low-power, and high- I_{sp} combinations. The propulsion block, therefore, serves as a structural prior over the low-thrust parameterization.

Latent variables in the trajectory model represent effective propulsion quantities, summarizing actuation authority consistent with observed mean-element evolution and beam-law constraints. These variables do not yield a unique reconstruction of the internal thruster operating state. Multiple combinations of voltage utilization, current utilization, mass utilization, divergence, satellite attitude, and duty-cycle history can produce indistinguishable semi-major axis signatures at the catalog scale. The reduced closure restricts the admissible set of propulsion parameters, while internal-state non-identifiability remains intrinsic to the public catalog inverse problem [83, 179, 17].

For krypton Hall thrusters, thrust, specific impulse, and efficiency depend on operating regime, magnetic topology, propellant processing, and facility interactions [188, 180]. The model omits detailed discharge physics and instead uses krypton beam relations to regularize the effective low-thrust parameter space. This regularization is sufficient to stabilize phase-ordered behavior estimates across insertion, operational-shell maintenance, and disposal regimes.

6.3.3 Surrogate Chemistry and Feasibility Regularization

The chemistry module functions as a reduced-order surrogate designed to penalize latent parameter sets incompatible with plausible krypton discharge regimes. An electron-temperature proxy is defined as

$$T_e \approx T_{e,0} + c_V V_b + c_\xi \xi_{\text{ion}} \quad (6.57)$$

where T_e is expressed in electron-volts (eV), $T_{e,0}$ is a baseline electron-temperature offset, c_V and c_ξ are fixed surrogate coefficients, and ξ_{ion} is a dimensionless ionization-state latent variable. From this proxy, a Voronov-type krypton ionization-rate-coefficient surrogate is computed as

$$K_{\text{iv}}^{\text{Kr}} \approx 5 \times 10^{-15} T_e^{1.25} \exp\left(-\frac{9.7}{T_e}\right) \quad (6.58)$$

where $K_{\text{iv}}^{\text{Kr}}$ is expressed in $\text{m}^3 \text{s}^{-1}$ [189, 190]. The symbol $K_{\text{iv}}^{\text{Kr}}$ is used because the quantity is a rate coefficient rather than a geometric cross section.

Neutral density is estimated using a pressure-temperature relation,

$$n_n = \frac{p}{k_B T_n} \quad (6.59)$$

where n_n is the neutral number-density proxy, p is the neutral pressure proxy, k_B is Boltzmann's constant, and T_n is the neutral-temperature proxy. The electron number-density proxy is modeled as

$$n_e \approx \xi_{\text{ion}} n_n \quad (6.60)$$

and an effective ionization mean free path is formed as

$$\ell_i = \frac{v_n}{n_e K_{\text{iv}}^{\text{Kr}}} \quad (6.61)$$

where v_n is a neutral thermal-speed proxy. The mean free path is denoted by ℓ_i so that λ remains reserved for the catalog mean-longitude proxy. The final scalar feasibility diagnostic is the ionization-length ratio,

$$r_i = \frac{\ell_i}{L_{\text{ref}}} \quad (6.62)$$

where L_{ref} is a fixed reference channel length.

Quantities in this module function as feasibility diagnostics. The implementation labels the resulting term as a `chemistry_penalty`, denoted here by $\mathcal{R}_{\text{chem}}$, to preserve diagnostic interpretation. This methodological stance is justified by the observation layer: the public catalog does not provide direct measurements of discharge current, plasma potential structure, ionization-front location, electron energy distribution, wall interaction, or neutral density inside the discharge channel. A soft regularizer is therefore appropriate, allowing the latent parameter search to remain within a physically interpretable krypton-discharge regime [83, 188, 180, 189, 190].

At the optimization layer, the reduced trajectory model, beam closure, and chemistry surrogate are incorporated through distinct penalty channels. Since $\mathcal{J}$ is already reserved for the segment index set, the Stage A deterministic trajectory-fitting objective is denoted by $\mathcal{L}_A$. In compact notation,

$$\mathcal{L}_A = \mathcal{L}_{\text{obs}} + w_{\text{coll}} \mathcal{L}_{\text{coll}} + w_{\text{prior}} \mathcal{L}_{\text{prior}} + w_{\text{hall}} \mathcal{L}_{\text{hall}} + w_{\text{chem}} \mathcal{L}_{\text{chem}} + w_{\text{feas}} \mathcal{L}_{\text{feas}} \quad (6.63)$$

Here, $\mathcal{L}_{\text{obs}}$ denotes the observational trajectory-matching loss, $\mathcal{L}_{\text{coll}}$ denotes the collocation or multiple-shooting continuity penalty, $\mathcal{L}_{\text{prior}}$ denotes prior regularization, $\mathcal{L}_{\text{hall}}$ denotes the Hall-thruster beam-closure penalty, $\mathcal{L}_{\text{chem}}$ denotes the chemistry-surrogate penalty, and $\mathcal{L}_{\text{feas}}$ denotes the remaining feasibility penalty. The nonnegative weights w_{coll} , w_{prior} , w_{hall} , w_{chem} , and w_{feas} control the relative influence of each penalty channel.

The trajectory-matching Stage A path simplifies the objective by focusing on the semi-major axis path, endpoint consistency, and multiple-shooting continuity. RAAN and mean-longitude end-

point terms are incorporated through curriculum scheduling and phase weighting. Propulsion and chemistry components function chiefly as structural regularizers on the latent parameter manifold, with the semi-major axis path serving as the principal empirical anchor.

6.4 Trajectory Estimation and Learned Phase Structure

6.4.1 Composite Trajectory Objective and Weighting Structure

For arc r , let

$$\widehat{\mathbf{a}}_r = \begin{bmatrix} \widehat{a}_{r,0} & \cdots & \widehat{a}_{r,N_r-1} \end{bmatrix}^\top \quad (6.64)$$

denote the predicted semi-major axis path at the catalog observation epochs, where

$$\widehat{a}_{r,k} = \widehat{a}_r(t_{r,k}) \quad (6.65)$$

The observed catalog-consistent semi-major axis path is retained from the Stage A observable definition,

$$\mathbf{a}_r^{\text{cat}} = \begin{bmatrix} a_{r,0}^{\text{cat}} & \cdots & a_{r,N_r-1}^{\text{cat}} \end{bmatrix}^\top \quad (6.66)$$

The Stage A deterministic trajectory-fitting objective is denoted by $\mathcal{L}_A$, preserving $\mathcal{J}$ for the segment index set. The objective is primarily influenced by the semi-major axis path term, the semi-major axis endpoint term, and a multiple-shooting continuity term. Additional regularization is provided by prior, Hall, chemistry, and feasibility penalties. This structure is consistent with nonlinear least-squares trajectory estimation and direct optimal-control transcription, where observation residuals, endpoint residuals, continuity defects, and regularization terms are balanced within a composite objective [191, 192, 193]. In compact form,

$$\begin{aligned} \mathcal{L}_A = & w_{\text{path}} \mathcal{L}_{\text{path}} + w_{a_f} \mathcal{L}_{a_f} + w_{\Omega} \mathcal{L}_{\Omega} + w_{\lambda} \mathcal{L}_{\lambda} + w_{\text{cont}} \mathcal{L}_{\text{cont}} \\ & + w_{\text{prior}} \mathcal{L}_{\text{prior}} + w_{\text{hall}} \mathcal{L}_{\text{hall}} + w_{\text{chem}} \mathcal{L}_{\text{chem}} + w_{\text{feas}} \mathcal{L}_{\text{feas}} \end{aligned} \quad (6.67)$$

Here, w_{path} , w_{a_f} , w_{Ω} , w_{λ} , w_{cont} , w_{prior} , w_{hall} , w_{chem} , and w_{feas} are nonnegative objective weights [193, 192].

The path term includes all internal observation epochs within each arc,

$$\mathcal{L}_{\text{path}} = \sum_{r \in \mathcal{R}} \varpi_r \frac{1}{N_r} \sum_{k \in \mathcal{K}_r} \left(\widehat{a}_{r,k} - a_{r,k}^{\text{cat}} \right)^2 \quad (6.68)$$

where ϖ_r is a nonnegative arc weight. The endpoint semi-major axis term penalizes the final-state mismatch,

$$\mathcal{L}_{a_f} = \sum_{r \in \mathcal{R}} \varpi_r \left(\widehat{a}_{r,N_r-1} - a_{r,N_r-1}^{\text{cat}} \right)^2 \quad (6.69)$$

The angular endpoint diagnostic losses are formed on wrapped angular residuals, so squared errors are evaluated after mapping each residual to the periodic angular domain [185].

$$\mathcal{L}_{\Omega} = \sum_{r \in \mathcal{R}} \varpi_r \left[\text{wrap}_{2\pi} \left(\widehat{\Omega}_{r,f} - \Omega_{r,N_r-1} \right) \right]^2 \quad (6.70)$$

and

$$\mathcal{L}_{\lambda} = \sum_{r \in \mathcal{R}} \varpi_r \left[\text{wrap}_{2\pi} \left(\widehat{\lambda}_{r,f} - \lambda_{r,N_r-1} \right) \right]^2 \quad (6.71)$$

For long arcs partitioned into subarcs, continuity is enforced through a multiple-shooting penalty. Let B_r denote the number of subarcs used for arc r , and let b index adjacent subarc boundaries. The continuity penalty is written as

$$\mathcal{L}_{\text{cont}} = \sum_{r \in \mathcal{R}} \sum_{b=0}^{B_r-2} \left(\widehat{a}_{r,b}^{\text{end}} - \widehat{a}_{r,b+1}^{\text{start}} \right)^2 \quad (6.72)$$

where $\widehat{a}_{r,b}^{\text{end}}$ is the predicted terminal semi-major axis of subarc b , and $\widehat{a}_{r,b+1}^{\text{start}}$ is the starting semi-major axis assigned to the following subarc. This penalty enforces consistency across internal shooting boundaries while retaining local trajectory flexibility for extended intervals [194, 192].

Within the resolved trajectory-mode configuration, the weighting scheme is intentionally asymmetric. The semi-major axis path term is assigned $w_{\text{path}} = 5.0$, the endpoint term is set to $w_{a_f} = 1.0$, and the continuity term is maintained at $w_{\text{cont}} = 0.1$. This weighting prioritizes the measurement channel most directly tied to the reduced dynamical state while retaining continuity regularization for long intervals [191, 192].

Angular endpoint penalties are assigned

$$w_{\Omega} = 0, \quad w_{\lambda} = 0 \quad (6.73)$$

so RAAN and mean-longitude endpoint channels serve as diagnostics in this configuration. This formulation prioritizes semi-major axis fidelity, assigns angular closure to diagnostic evaluation, and determines the interpretation of subsequent validation results.




Figure 6.4: Stage A Log-Loss Decomposition

Figure 6.4 provides the log-scale decomposition of the Stage A loss history, making visible the dominant early contraction of the objective and the lower-amplitude variations that remain after the principal semi-major-axis structure has been fitted.

6.4.2 Optimization Behavior and Convergence




Figure 6.5: Stage A Trajectory Training

Figure 6.5 summarizes the Stage A optimization behavior by separating the trajectory-loss components from the semi-major-axis path RMSE. The figure shows the rapid early reduction in the trajectory objective and the subsequent plateau associated with convergence of the reduced phase-conditioned model.

Training diagnostics indicate rapid initial loss reduction followed by a sustained plateau. This pattern is characteristic of a reduced model that captures the dominant low-dimensional trajectory structure early and then confines further refinement to a narrower parameter region. The path root-mean-square error (RMSE) decreases sharply during the initial optimization iterations, after which the total objective and semi-major axis terms show only marginal improvement. In later iterations, the traces remain stable, indicating convergence within a localized parameter basin [193].

For arc r , the semi-major axis path RMSE is defined as

$$\text{RMSE}_{\text{path},r} = \left[\frac{1}{N_r} \sum_{k \in \mathcal{K}_r} \left(\hat{a}_{r,k} - a_{r,k}^{\text{cat}} \right)^2 \right]^{1/2} \quad (6.74)$$

The Stage A validation distribution yields an arc-level mean path RMSE of approximately $\overline{\text{RMSE}}_{\text{path}} \approx 3.26$ km with a median value of approximately $\widetilde{\text{RMSE}}_{\text{path}} \approx 0.36$ km. The segment-level endpoint validation yields $\text{RMSE}_{a_f} \approx 9.25$ km across 18,952 phase-labeled segments. The arc-level path RMSE evaluates within-arc trajectory geometry, while the segment-level endpoint RMSE evaluates final-state agreement over the full segment archive. Jointly, these diagnostics indicate that the reduced model reproduces the principal semi-major axis dynamics observed in the archive, with larger residuals concentrated in maneuver-intensive or label-impure arcs [191].

The limited change in angular endpoint diagnostics during training reflects the prioritization established by the objective function. The resolved configuration applies zero weight to endpoint RAAN and mean-longitude residuals, so the optimizer is directed toward semi-major axis path fidelity. This convergence pattern is consistent with the reduced-order formulation and produces a solution optimized for downstream behavior modeling in the dominant semi-major axis channel [193, 192].

6.4.3 Learned Global and Phase-Level Parameter Evolution

Learned parameter traces demonstrate convergence to a stable phase hierarchy. Let u denote the optimization-iteration index for Stage A. The global traces $m_{\text{eff}}^{(u)}$ and $I_{\text{sp}}^{(u)}$ remain close to their initial values, while most optimization flexibility is allocated to phase-conditioned variables and satellite-level scale factors. The final global estimates are approximately $m_{\text{eff}}^{\star} \approx 252$ kg and $I_{\text{sp}}^{\star} \approx 1.5 \times 10^3$ s.

These quantities are interpreted as effective trajectory-fit parameters regularized by the propulsion model, not as independently recovered spacecraft telemetry [83, 179, 188].




Figure 6.6: Epoch-by-Epoch Evolution of the Phase-Conditioned Parameters

Figure 6.6 shows the epoch-by-epoch evolution of thrust, duty, drag, and shell-drag parameters for the three maneuver phases. The traces indicate that the fitted solution organizes phase behavior primarily through effective actuation magnitude and duty cycle, while the drag-scale channels remain comparatively constrained by the imposed reduced model structure.

Phase-level parameter evolution provides the principal dynamical interpretation. The insertion or orbit-raise phase exhibits the highest effective actuation, disposal lowering is intermediate, and the operational-shell phase converges to a substantially reduced effective thrust and duty cycle. The final learned thrust magnitudes maintain the hierarchy

$$T_{\text{ins}}^{\star} \approx 5.61 \times 10^{-2} \text{ N}, \quad T_{\text{disp}}^{\star} \approx 3.36 \times 10^{-2} \text{ N}, \quad T_{\text{ops}}^{\star} \approx 3.36 \times 10^{-3} \text{ N}$$

with corresponding effective duties

$$d_{\text{ins}}^* \approx 0.45, \quad d_{\text{disp}}^* \approx 0.31, \quad d_{\text{ops}}^* \approx 0.026$$

Learned drag scales show minimal differentiation in this configuration, since shell-drag compensation remains fixed at imposed phase-specific values. The principal learned distinction is therefore expressed through actuation magnitude and duty cycle [83, 179, 188].




Figure 6.7: Final Effective Phase Parameters

Figure 6.7 summarizes the final effective thrust and duty hierarchy across the three maneuver phases. The insertion or orbit-raise phase occupies the highest actuation regime, disposal lowering remains intermediate, and operational-shell maintenance converges to the lowest effective thrust and duty level.

This phase ordering has a direct dynamical interpretation. The insertion regime, characterized by frequent orbit-raising maneuvers, requires the largest effective actuation scale to reproduce positive semi-major axis excursions. Disposal lowering remains active but is less dominant in the fitted hierarchy. The operational-shell phase functions as a low-control regime in which maintenance behavior is expressed through small effective actuation and drag compensation. The model therefore learns a reduced phase structure that is operationally plausible and suitable for forward modeling across distinct satellite behavior categories [2, 1, 83].

6.5 Trajectories, Residual Structure, and Failure Modes

This section evaluates the primary strengths and limitations of the Stage A fit. Semi-major axis trajectory alignment demonstrates high fidelity within the operational-shell regime and retains diagnostic value across all phases. Residual errors are concentrated within a sparse subset of challenging arcs, indicating structured failure modes concentrated in specific regimes and observable channels. The analysis is framed as a residual-based validation of a reduced orbit-determination and trajectory-estimation model, with sensitivity diagnostics used to test robustness to controlled perturbations [191, 195].

6.5.1 Aggregate Fit Quality and Phase-Resolved Error Hierarchy

Arc-level validation indicates that semi-major axis path fitting is the most robust result of Stage A. The use of residual distributions, endpoint errors, and phase-conditioned error summaries follows standard orbit-estimation practice, where fit quality is assessed through observation-minus-computation structure rather than only through final parameter values [191]. Let the empirical set of arc-level path errors be

$$\mathcal{E}_{\text{path}} = \{\text{RMSE}_{\text{path},r} : r \in \mathcal{R}\}. \quad (6.75)$$

Among $N_{\text{arc}} = 12,209$ labeled arcs, the aggregate mean path RMSE is $\overline{\text{RMSE}}_{\text{path}} = 3.26$ km with median $\widetilde{\text{RMSE}}_{\text{path}} = 0.360$ km and 90th percentile $Q_{0.90}(\mathcal{E}_{\text{path}}) = 8.54$ km. The pronounced skew, characterized by a low median and elevated upper quantile, confirms that a small subset of arcs contributes disproportionately to the error tail, while the majority of arcs are accurately fit.

Phase stratification remains consistent. Let $N_{\text{arc},p}$ denote the number of admissible arcs assigned to phase p , and let $\overline{\text{RMSE}}_{\text{path},p}$ and $\widetilde{\text{RMSE}}_{\text{path},p}$ denote the corresponding phase-conditioned mean and median path errors. The operational-shell regime contains $N_{\text{arc,ops}} = 7,466$ arcs and achieves $\overline{\text{RMSE}}_{\text{path,ops}} = 0.623$ km and $\widetilde{\text{RMSE}}_{\text{path,ops}} = 0.134$ km. Disposal lowering contains $N_{\text{arc,disp}} = 3,641$ arcs and yields $\overline{\text{RMSE}}_{\text{path,disp}} = 6.62$ km and $\widetilde{\text{RMSE}}_{\text{path,disp}} = 2.52$ km. Insertion or orbit raising contains $N_{\text{arc,ins}} = 1,102$ arcs and produces $\overline{\text{RMSE}}_{\text{path,ins}} = 10.08$ km and $\widetilde{\text{RMSE}}_{\text{path,ins}} = 7.39$ km. This hierarchy reflects the underlying data composition, where operational-shell dominance and limited semi-major axis excursions support superior fit quality.




Figure 6.8: Distribution of Trajectory Path RMSE by Phase

Figure 6.8 visualizes this phase-resolved error hierarchy, showing a sharply concentrated operational-shell distribution and broader disposal-lowering and insertion/orbit-raise tails associated with maneuver-rich and label-sensitive intervals. Across $N_{\text{seg}} = 18,952$ segments, the overall endpoint semi-major axis RMSE is $\text{RMSE}_{a_f} = 9.25$ km. By phase, the endpoint semi-major axis errors are $\text{RMSE}_{a_f, \text{ops}} = 0.367$ km, $\text{RMSE}_{a_f, \text{disp}} = 15.72$ km, and $\text{RMSE}_{a_f, \text{ins}} = 18.28$ km. The reduced model performs best in the operational shell, which is physically dominant, and maintains consistent phase ordering across regimes with increased maneuver complexity.

Angular endpoint statistics are of secondary importance in this Stage A configuration. Since RAAN and mean longitude are circular variables, wrapped angular errors are interpreted on the periodic domain [185]. At the arc level, the aggregate RAAN mean absolute error (MAE) is $\text{MAE}_{\Omega} = 1.59$ rad and the mean-longitude MAE is $\text{MAE}_{\lambda} = 1.55$ rad. The weak phase differentiation in these angular statistics reflects the Stage A emphasis on semi-major axis path fidelity, with angular channels serving as diagnostics.

6.5.2 Observed-Predicted Structure and Residual Geometry

The observed-versus-predicted structure reveals significant channel asymmetry. In the semi-major axis endpoint panel, predicted values align closely with the identity relation, confirming accurate representation of the dominant endpoint behavior. In the wrapped $\Delta\Omega$ and $\Delta\lambda$ panels, predicted values form diffuse and branch-limited structures with weaker global correspondence. This comparison separates linear residual behavior in the semi-major-axis channel from circular residual behavior in angular channels [191, 185].

Figure 6.9 confirms that weaker angular performance follows from the modeling and weighting

structure. The estimator is designed to reproduce semi-major axis evolution, while angular channels are retained as secondary diagnostics. The endpoint scatter structure motivates the residual definitions by separating the semi-major axis channel from the angular endpoint channels.




Figure 6.9: Observed versus Predicted Endpoint Structure

For segment j , define the endpoint residuals

$$\epsilon_{a_1,j} = \widehat{a}_{j,1} - a_{j,1}^{\text{cat}}, \quad (6.76)$$

$$\epsilon_{\Omega,j} = \text{wrap}_{2\pi} \left(\widehat{\Omega}_{j,1} - \Omega_{j,1} \right), \quad (6.77)$$

and

$$\epsilon_{\lambda,j} = \text{wrap}_{2\pi} \left(\widehat{\lambda}_{j,1} - \lambda_{j,1} \right) \quad (6.78)$$

The corresponding empirical residual sets are

$$\mathcal{E}_{a_1} = \{\epsilon_{a_1,j} : j \in \mathcal{J}\}, \quad \mathcal{E}_{\Omega} = \{\epsilon_{\Omega,j} : j \in \mathcal{J}\}, \quad \mathcal{E}_{\lambda} = \{\epsilon_{\lambda,j} : j \in \mathcal{J}\} \quad (6.79)$$

Semi-major axis residuals are tightly centered with a sparse heavy tail, indicating strong fit quality interrupted by a limited number of challenging segments. Wrapped angular residuals are broader and less structured, reflecting weaker agreement between the reduced angular endpoint model and catalog observations. Collectively, scatter and residual analyses demonstrate that Stage A isolates the dominant semi-major axis signal, with mismatches concentrated in specific channels and cases [191, 185].

6.5.3 Robustness to Timing and Model Perturbation

Timing-sensitivity analysis provides a targeted robustness assessment beyond aggregate fit metrics. The diagnostic follows the general sensitivity-analysis principle of perturbing selected inputs and

evaluating the induced change in model outputs [195]. In this diagnostic, segment durations are shifted over a symmetric window,

$$\delta_t \in \mathcal{D}_{\text{shift}}, \quad \mathcal{D}_{\text{shift}} \subset [-\Delta t_{\text{shift}}, \Delta t_{\text{shift}}], \quad \Delta t_{\text{shift}} = 6 \text{ hr} \quad (6.80)$$

The fitted Stage A model is re-evaluated under each uniform timing perturbation without retraining. For observable channel

$$c \in \{a_f, \Omega, \lambda\} \quad (6.81)$$

the perturbed validation error is denoted by

$$\text{RMSE}_c(\delta_t) \quad (6.82)$$

The resulting error curves display pronounced channel dependence: semi-major axis endpoint RMSE remains nearly invariant, RAAN error shifts minimally, and mean-longitude RMSE varies substantially with a non-monotonic profile.

The fitted semi-major axis solution demonstrates robustness to moderate timing perturbations, supporting the conclusion that dominant path fidelity is not dependent on exact boundary alignment. The wrapped mean-longitude channel exhibits stronger timing sensitivity, consistent with the previously observed weaker global $\Delta\lambda$ closure. Mean longitude is therefore the least robust observable in the Stage A validation, while semi-major axis remains the most stable and propagation-relevant metric.




Figure 6.10: Sensitivity of Validation RMSE to Uniform Timing Shifts

Figure 6.10 quantifies this robustness pattern by showing near-invariance in semi-major axis RMSE across the imposed timing-shift window, modest sensitivity in RAAN, and substantially

stronger variation in mean-longitude closure.

Interpretation of post hoc force-model perturbations requires caution. Force-model ablation is an input-perturbation diagnostic, so its interpretation depends on the selected observable channels and on the dynamical role of each omitted perturbation [195, 2, 1]. Let $\mathcal{A}_{\text{abl}}$ denote the set of force-model ablation cases, and let $F \in \mathcal{A}_{\text{abl}}$ denote a particular term-toggling configuration. Semi-major axis ablation plots can suggest weak sensitivity to individual force-model terms, but that interpretation is incomplete when angular channels are excluded. Angular endpoints respond differently to reduced force-model structure, especially when J_2 terms are omitted. The primary robustness discussion is therefore based on timing sensitivity and channel asymmetry, while force-model ablations are interpreted through the full validation metric table.

The combined diagnostics confirm that Stage A provides robust, high-fidelity performance in the semi-major axis channel, particularly within operational-shell regimes. Mean longitude remains the weakest channel, and the largest residuals originate from sparse, structured arcs. These characteristics make the fitted phase-dependent model useful as a behavioral propagation layer before posterior analysis.

6.6 Posterior Inference and Identifiability

Stage B reformulates the deterministic Stage A fit as a reduced Bayesian inverse problem operating on simulator-derived summary statistics. This approach restricts the posterior to effective parameters consistent with the selected summary features, prior support, and simulator configuration [164, 181]. Stage A uses phase-labeled semi-major axis paths across multi-epoch arcs, whereas Stage B uses a condensed segment-level summary vector, thereby introducing an intentional information bottleneck [184]. The resulting posterior represents distributions over reduced parameters compatible with the observed summaries, which precludes unique reconstruction of the complete satellite state, commanded thrust history, or internal propulsion state [196, 197].

6.6.1 Reduced Posterior Parameterization and SNPE Setup

The Stage B parameter vector is intentionally compact. To distinguish this posterior parameterization from the broader latent vector θ , define

$$\theta_B = \left[m_{\text{eff}}, I_{\text{sp}}, T_{\text{ops}}, D_{\text{ops}}, T_{\text{disp}}, D_{\text{disp}}, T_{\text{ins}}, D_{\text{ins}} \right]^T \in \Theta_B \quad (6.83)$$

where m_{eff} denotes the effective satellite mass scale, I_{sp} represents the effective specific impulse, and T_p, D_p denote phase-conditioned effective thrust and drag parameters for $p \in \mathcal{P}$. The compact representation has dimension $d_\theta = \dim(\theta_B) = 8$ and preserves the principal phase asymmetries identified in Stage A [83, 179].

The observation map used for inference is similarly reduced. For each segment j , Stage B conditions on the existing six-dimensional summary vector

$$\boldsymbol{\eta}_j = \begin{bmatrix} \Delta a_j^{\text{cat}} \\ \Delta \Omega_j \\ \Delta \lambda_j \\ \dot{a}_j^{\text{cat}} \\ \dot{\Omega}_j \\ \dot{\lambda}_j \end{bmatrix}, \quad d_\eta = \dim(\boldsymbol{\eta}_j) = 6 \quad (6.84)$$

where

$$\dot{a}_j^{\text{cat}} = \frac{\Delta a_j^{\text{cat}}}{\Delta t_j}, \quad \dot{\Omega}_j = \frac{\Delta \Omega_j}{\Delta t_j}, \quad \dot{\lambda}_j = \frac{\Delta \lambda_j}{\Delta t_j} \quad (6.85)$$

The symbol $\boldsymbol{\eta}_j$ is retained from the Stage B summary construction so that $\mathbf{x}_{r,0}^{\text{mean}}$ remains reserved for initial mean-state estimates. The effective-mode implementation excludes context features that are invariant or weakly informative under the reduced posterior design. Specifically, the summary omits a_0, e_0, i_0 , elapsed time, and phase sign. As a result, the posterior is conditioned on compressed segment features rather than on complete arc geometries [184, 181].

Let $\mathcal{S}_B$ denote the Stage B simulator that maps reduced parameters to segment-level summaries,

$$\boldsymbol{\eta} = \mathcal{S}_B(\boldsymbol{\theta}_B) \quad (6.86)$$

with simulator stochasticity or sampling variability implicit when applicable. The reduced Bayesian target is

$$p_B(\boldsymbol{\theta}_B \mid \boldsymbol{\eta}_{\text{obs}}) \propto p_B(\boldsymbol{\eta}_{\text{obs}} \mid \boldsymbol{\theta}_B) \pi_B(\boldsymbol{\theta}_B) \quad (6.87)$$

where $\pi_B(\boldsymbol{\theta}_B)$ is the Stage B prior and $\boldsymbol{\eta}_{\text{obs}}$ is the observed summary used for conditioning. The likelihood term is implicit because the simulator can generate parameter-summary pairs without requiring closed-form evaluation of $p_B(\boldsymbol{\eta}_{\text{obs}} \mid \boldsymbol{\theta}_B)$ [164, 181].

The inference algorithm is Sequential Neural Posterior Estimation (SNPE), implemented with the `sbi` toolkit using a masked autoregressive flow (MAF) density estimator [181, 182, 198, 183].

The learned neural posterior approximation is written as

$$\widehat{p}_{\phi}(\boldsymbol{\theta}_B \mid \boldsymbol{\eta}) \quad (6.88)$$

where ϕ denotes the trainable density-estimator parameters. The simulator produces training pairs

$$\left\{ \left(\boldsymbol{\theta}_B^{(i)}, \boldsymbol{\eta}^{(i)} \right) \right\}_{i=1}^{N_{\text{sim}}}, \quad N_{\text{sim}} = 3000 \quad (6.89)$$

from which the neural density estimator learns the posterior approximation. The current configuration uses a MAF estimator, draws $N_{\text{post}} = 2000$ posterior samples, and evaluates adequacy using $N_{\text{pp},B} = 256$ Stage B posterior-predictive draws. The resulting posterior is an amortized approximation defined relative to the reduced summary operator, the prior, the simulator, and the simulation budget specified here [176, 181].

This inference configuration requires substantially less computational effort than repeated full-archive trajectory fitting while still quantifying uncertainty in effective parameters. Stage B converts the Stage A point-estimate structure into a phase-conditioned posterior, allowing uncertainty propagation in drag and actuation behavior from a distribution over reduced operating regimes [164, 181].

6.6.2 Posterior Marginals and Compensating Parameter Structure

Posterior marginals display substantial breadth while maintaining discernible structure. The effective mass posterior has a mean of $\bar{m}_{\text{eff},B} \approx 287$ kg and a standard deviation of $s_{m,B} \approx 57.6$ kg, and the effective specific-impulse posterior has a mean of $\bar{I}_{\text{sp},B} \approx 1990$ s and a standard deviation of $s_{I,B} \approx 729$ s. These variances indicate that multiple global operating points remain compatible with the reduced summary features, although posterior mass remains concentrated within physically plausible regions of the prior [164, 196].

The phase-level marginals show a similar pattern with a more direct dynamical interpretation. The operational-shell thrust posterior has mean $\bar{T}_{\text{ops},B} \approx 0.120$ N with wide support extending toward $Q_{0.95}(T_{\text{ops}}) \approx 0.23$ N. In contrast, the operational-shell drag marginal is more tightly concentrated around $\bar{D}_{\text{ops},B} \approx 2 \times 10^{-8}$ km s⁻². Both disposal-lowering thrust and insertion thrust display broad, right-skewed distributions, whereas corresponding drag channels remain comparatively concentrated. This asymmetry indicates that the posterior constrains drag-like combinations more strongly than thrust-like combinations while retaining meaningful phase separation [154, 155, 156, 197].




Figure 6.11: Posterior Correlation Structure for Reduced Stage B Parameters

Figure 6.11 displays the posterior correlation matrix for the reduced Stage B parameter vector, making explicit the compensating parameter directions that arise when phase-conditioned thrust and drag combinations produce similar segment-level summaries [196, 197].

The posterior dependence structure reveals compensation trends among parameter channels. Let $\mathbf{C}_{\theta,B}$ denote the posterior correlation matrix of θ_B , with entries

$$\rho_{uv}^B = \text{corr}(\theta_{B,u}, \theta_{B,v} \mid \boldsymbol{\eta}_{\text{obs}}) \quad (6.90)$$

Correlation heatmaps and low-dimensional projections show that operational-shell thrust and drag exhibit pronounced negative dependence, with similar compensation behavior in disposal lowering and structured coupling in insertion parameters. The posterior, therefore, organizes around families

of drag-actuation combinations that produce similar segment summaries. Broad marginals reflect the many-to-one structure of the inverse problem while preserving phase-specific behavioral content required for satellite constellation-scale modeling [196, 197].

This compensation structure follows from the adopted summary design. The six-dimensional segment summary removes distinctions required for unique identification of global mass, specific impulse, and phase-conditioned thrust histories. For satellite constellation-scale analysis, the posterior remains informative because it captures phase-dependent behavior and the uncertainty induced by the reduced observation operator [184, 181].

6.6.3 Thruster Performance

Table 6.1: Effective Phase-Resolved Thruster Performance

Metric	Insertion / Orbit Raise	Operational Shell	Disposal Lowering
Duty cycle	0.4458	0.0446	0.3430
Mean thrust [mN]	27.14	1.634	13.505
Mean I_{sp} [s]	1499.76	1499.76	1499.76
Mean efficiency	0.4752	0.4752	0.4752
Mean power [W]	273.12	16.44	135.90
Thrust-to-power [mN/kW]	99.37	99.37	99.37
Total impulse [kN s]	22241.65	378.44	9839.12
Total Δv [m/s]	89673.81	1502.42	39273.91
Total propellant [kg]	1512.25	25.73	668.98
Total energy [kWh]	62173.72	1057.87	27504.04
Median per-satellite impulse [N s]	46552.68	711.34	17474.30
Median per-satellite Δv [m/s]	187.46	2.82	69.49
Median per-satellite propellant [kg]	3.1652	0.0484	1.1881

The effective global parameters used in the performance reconstruction are $m_{\text{eff}} = 252.15$ kg, $m_{\text{dry}} = 149.98$ kg, $I_{sp} = 1499.76$ s, and $\eta_{\text{tot},0} = 0.4800$. Specific impulse is defined as $I_{sp} = v_e/g_0$, and total impulse is the time integral of thrust. The table reports reconstructed effective archive-level quantities derived from fitted Stage A phase parameters and labeled intervals [83, 179, 188].

Table 6.1 presents operational performance metrics derived from the fitted Stage A phase parameters and aggregated over the labeled archive. For phase $p \in \mathcal{P}_{\text{perf}}$, define the reconstructed mean

effective duty cycle, thrust, specific impulse, efficiency, and power as

$$\bar{d}_p^{\text{perf}}, \quad \bar{T}_p^{\text{perf}}, \quad \bar{I}_{\text{sp},p}^{\text{perf}}, \quad \bar{\eta}_{\text{tot},p}^{\text{perf}}, \quad \bar{P}_p^{\text{perf}} \quad (6.91)$$

The corresponding thrust-to-power ratio is written as

$$\Pi_{T/P,p} = \frac{\bar{T}_p^{\text{perf}}}{\bar{P}_p^{\text{perf}}} \quad (6.92)$$

with units reported as mNkW^{-1} . Aggregate phase-level impulse, velocity increment, propellant expenditure, and energy are denoted by

$$\mathcal{I}_p^{\text{tot}}, \quad \Delta V_p^{\text{tot}}, \quad m_{\text{prop},p}^{\text{tot}}, \quad E_p^{\text{tot}} \quad (6.93)$$

Median per-satellite quantities are denoted by

$$\widetilde{\mathcal{I}}_{\text{sat},p}, \quad \widetilde{\Delta V}_{\text{sat},p}, \quad \widetilde{m}_{\text{prop,sat},p} \quad (6.94)$$

The reconstructed performance hierarchy preserves the inferred phase sequencing. The insertion-orbit-raise phase dominates in effective thrust, duty cycle, impulse, and energy expenditure. The disposal-lowering phase remains dynamically active at lower reconstructed magnitudes. Operational-shell intervals are characterized by low mean thrust, limited duty cycle, and small cumulative Δv . The table establishes a direct connection between reduced-fit trajectories and fleet-scale behavior models, preserving the phase dependencies between drag compensation and low-thrust actuation identified in the preceding Stage A and Stage B analyses [83, 179, 188, 180].

6.6.4 Practical Identifiability and Low-Rank Structure

The primary diagnostic in this section is the local identifiability analysis [196, 197]. Let the Stage B sensitivity matrix be defined as

$$\mathbf{S}_B = \left. \frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\theta}_B} \right|_{\boldsymbol{\theta}_B = \boldsymbol{\theta}_B^\dagger} \quad (6.95)$$

where $\boldsymbol{\eta} \in \mathbb{R}^{d_\eta}$ is the Stage B summary vector, $\boldsymbol{\theta}_B \in \mathbb{R}^{d_\theta}$ is the compact posterior parameter vector, and $\boldsymbol{\theta}_B^\dagger$ denotes the selected operating point or posterior representative used for the local diagnostic.

The singular-value decomposition (SVD) of the sensitivity matrix is

$$\mathbf{S}_B = \mathbf{U}_B \mathbf{\Sigma}_B \mathbf{V}_B^\top \quad (6.96)$$

with singular-value vector

$$\mathbf{s}_B = (s_{1,B}, s_{2,B}, \dots, s_{\min(d_\eta, d_\theta), B}) \quad (6.97)$$

The reported numerical rank is

$$r_B = \text{rank}(\mathbf{S}_B) = 3 \quad (6.98)$$

for the eight-dimensional Stage B parameter vector, and the corresponding condition number is approximately

$$\kappa(\mathbf{S}_B) \approx 1.31 \times 10^{22} \quad (6.99)$$

This result indicates that the local inverse mapping from reduced summary features to the compact parameter vector is highly rank-deficient. The singular-value spectrum demonstrates this structure visually, with three dominant singular values followed by a sharp decline across several orders of magnitude into a near-null tail [197, 196].




Figure 6.12: Singular Values of the Local Stage B Sensitivity Operator

Figure 6.12 shows the low-rank singular-value spectrum of the local Stage B sensitivity operator, confirming that only a small subset of parameter combinations is locally resolved by the compressed summary representation [197].

The singular-value structure is consistent with the posterior marginals, pairwise compensation patterns, and parameter-sensitivity diagnostics [196, 197, 195]. Let $\mathcal{R}_B$ denote the locally resolved subspace associated with the dominant singular directions, and let $\mathcal{N}_B$ denote the near-null subspace associated with weakly constrained or compensating parameter directions. In the current diagnostic,

$\mathcal{R}_B$ is low-dimensional relative to the eight-dimensional parameterization. Weak identifiability is observed in effective mass, specific impulse, and selected thrust components. Direct sensitivity is concentrated primarily in drag-like directions, especially within the operational-shell regime. The posterior, therefore, resolves a limited set of dominant parameter combinations, while remaining parameter directions are weakly constrained or compensated through posterior dependence.

This interpretation maintains scientific rigor within the chapter’s observational constraints. The posterior analysis uses catalog-based mean-element summaries, and the simulator employs a reduced phase-conditioned model with surrogate chemistry regularization. Under these conditions, low-rank identifiability follows from the compressed inverse-problem formulation. Stage B provides a quantitative representation of compatible drag and actuation parameter combinations, their phase-dependent variability, and their uncertainty propagation in prediction. Posterior marginals, posterior predictive checks, and singular-value spectra therefore provide a coherent diagnostic basis for interpreting the reduced parameter estimates [176, 196, 197].

Chapter 7

Recommendations, Discussion, and Implications

Chapters 3 through 6 present four methodological layers, each developed as an independent analysis module while attaining greatest utility as part of an integrated satellite constellation-analysis workflow. The clustering methodology produces mesoscale representations of satellite constellation structure. Forecasting methods generate time-resolved environmental forcing for forward propagation. Bayesian inference provides satellite-specific uncertainty distributions for effective drag response. System identification derives reduced, phase-conditioned descriptions of actuation and secular orbital behavior. Chapter 2 establishes the high-fidelity propagation environment that permits these elements to interact dynamically. Together, these chapters constitute a unified framework for reconstructing, simulating, and optimizing large LEO satellite constellations using public orbital records.

Integration of the methodological layers allows public catalog histories to support analyses extending beyond shell counting and single-satellite trend inspection. The archive enables recovery of family structure, identification of distinct drag-response regimes, inference of reduced maneuver behavior, and generation of uncertainty-aware forward simulations under time-varying thermospheric forcing. The recovered quantities are effective and conditional, defined by their capacity to reproduce observable orbital dynamics within a transparent modeling framework. Their interpretability is achieved independently of confidential operator telemetry.

7.1 Integrated Findings Relative to Initial Hypotheses

A central premise of the thesis is that large satellite constellation analysis is most effective at an intermediate scale, positioned between individual-satellite inspection and fleet-scale aggregation. Clustering results support this premise. Positive controls demonstrate that the ROPV to UMAP to HDBSCAN pipeline recovers externally documented structure when such structure is present. For GPS, the methodology reconstructs a well-defined plane architecture. For OneWeb, the methodology remains effective despite increased geometric nonlinearity and reduced discrete separation. These benchmarks are essential because the Starlink analysis requires structural discovery using partial public priors and without direct family labels.

Analysis of Starlink indicates that global clustering alone is inadequate for highly stratified satellite constellations. Decomposition by shell reveals systematic topological variation across inclination regimes. The 53.0° shell forms a dense annular manifold with sector-like organization. The 53.2° shell is smoother and less discretized. The 70.0° shell is fragmented, exhibiting islands, filaments, and loops. The 97.6° shell displays the most discrete and persistent structure. These results demonstrate that modern satellite constellations exhibit substantial internal geometric heterogeneity, with inter-shell differences in mesoscale organization directly affecting downstream inference and optimization.

This conclusion assigns clustering a structural role within the broader workflow. The latent chart functions as a structural operator, transforming a complex fleet-scale problem into a family-conditioned representation in which local neighborhoods, shell-specific manifolds, and density-connected components are explicit. Once identified, this structure improves the formulation of downstream tasks. Ballistic-coefficient inference can be interpreted within dynamical families, system-identification residuals can be contextualized through phase-consistent neighbors, and optimization can be conducted over families or shell-conditioned representatives.

Forecasting analysis demonstrates the necessity of thermospheric models that evolve with solar-cycle phase and geomagnetic variability in future satellite constellation simulations. Validation using a five-fold expanding-window protocol yields MAE between 0.64 and 0.78 and RMSE between 0.94 and 1.09 under causal validation. These metrics support operationally relevant prospective skill instead of retrospective curve agreement.

Residual analysis in the forecasting chapter reveals structured, non-memoryless error, concentrated near zero with asymmetric tails and evolving bias. The spectral content retains power at physically meaningful periods. This finding is critical for downstream propagation because deterministic atmospheric forcing trajectories cannot capture post-fit structured uncertainty. The presence of residual solar rotation and intermediate-period content indicates that the atmosphere

entering the simulator is a forecast-conditioned stochastic field. This result defines the uncertainty geometry governing drag accumulation, shell occupancy, and long-horizon dispersion.

Bayesian inference shows that public decay histories sufficiently constrain satellite-specific effective drag responses to enable posterior-predictive simulation. Inferred ballistic-coefficient medians range from 21.64 kg/m^2 to 38.35 kg/m^2 , exceeding a factor of 1.7 across the sample. Posterior intervals are narrower for satellites with longer, cleaner decay arcs and broader for shorter arcs, reflecting information gain proportional to observable trajectory curvature.

This analysis indicates that satellite constellation drag state requires satellite-specific treatment, as significant heterogeneity persists even within narrow mission classes. Differences may arise from area evolution, attitude state, configuration history, density-model discrepancy, or unmodeled force contributions. Although the framework does not isolate a unique physical cause for each inferred response, the posterior over effective ballistic coefficient is sufficiently informative to generate posterior-predictive fans in semi-major axis and forecast distributions of re-entry time. This mapping from parameter uncertainty to propagated state uncertainty represents a principal operational outcome of the thesis.

System identification demonstrates that public mean-element histories contain sufficient information to resolve reduced actuation structure under phase-conditioned inference. Stage A reconstructs dominant semi-major axis behavior with an aggregate mean path RMSE of 3.26 km and a median path RMSE of 0.360 km. In the operational-shell phase, the mean path RMSE is 0.623 km, and the median path RMSE is 0.134 km. These results confirm that reduced dynamical models reproduce dominant orbital signals with fleet-relevant fidelity.

The posterior stage quantifies practical identifiability where the reduced sensitivity operator has numerical rank $r_B = 3$ for an eight-parameter posterior, with condition number $\kappa(\mathbf{S}_B) \approx 1.31 \times 10^{22}$. This spectrum establishes a sharply low-dimensional identifiable subspace within the selected summary design. The inverse mapping compresses the parameter space, constraining dominant combinations of drag-like and thrust-like behavior while leaving several orthogonal directions weakly resolved. This clarification refines the interpretation of recovered parameters. The phase-conditioned posterior provides a structured uncertainty model over effective operating regimes, preserving compensation geometry, phase separation, and predictive power while limiting attribution to hardware-specific telemetry.

A recurring pattern unifies the four technical methodologies where public archive information is concentrated in reduced objects. These include shell-conditioned families, forecast-conditioned forcing trajectories, posterior distributions over effective drag, and low-rank actuation parameters. When analysis is framed in these terms, the methodologies interconnect as elements of a coherent

estimation architecture, represented as

$$\mathcal{D}_{\text{cat}} \longrightarrow \mathcal{C}_{\text{fam}} \longrightarrow \widehat{\mathcal{F}}_{\text{atm}} \longrightarrow \left[p(BC_s | \mathcal{D}_s^{\text{decay}}), p_B(\theta_B | \eta_{\text{obs}}) \right] \longrightarrow \widehat{\mathcal{X}}_{[t_0, t_f]} \quad (7.1)$$

Here, $\mathcal{D}_{\text{cat}}$ represents the synchronized catalog archive, $\mathcal{C}_{\text{fam}}$ represents the shell-conditioned family structure, $\widehat{\mathcal{F}}_{\text{atm}}$ represents the forecast-conditioned atmospheric forcing process, BC_s is the satellite-specific effective ballistic coefficient for satellite s , $\mathcal{D}_s^{\text{decay}}$ is the satellite-specific decay arc used for drag inference, $p_B(\theta_B | \eta_{\text{obs}})$ is the Stage B posterior over reduced phase-conditioned parameters, and $\widehat{\mathcal{X}}_{[t_0, t_f]}$ denotes the resulting propagated satellite constellation-state ensemble over the interval $[t_0, t_f]$. The thesis thereby establishes a hierarchical estimation pipeline for satellite constellation behavior from public data, with each stage providing a reduced state, uncertainty model, or structural representation required by the next.

7.2 Importance and Applications for Satellite Constellations

The integrated framework provides scientific value by focusing on the intermediate, or mesoscale, level of satellite constellations. Traditional methodologies often concentrate on either individual satellites or entire fleets, thereby obscuring the operationally significant structures that emerge between these scales. Critical processes such as shell maintenance, deployment progression, family turnover, disposal routing, common-mode anomalies, and environmental response require analysis that incorporates family-conditioned dynamics. Single-satellite inspection lacks the population context required to identify shared structure, while fleet-average analysis suppresses the heterogeneity that governs operational behavior within specific shells and families.

The mesoscale regime constitutes the most effective domain for satellite constellation optimization. The simulator described in Chapter 2 models multi-satellite dynamics using high-fidelity force models and phase-dependent low thrust. Subsequent chapters reduce the decision space through clustering, family identification, medoid selection, and probabilistic forecasting. Bayesian inference produces drag priors tailored to specific objects or dynamically coherent families, while system identification constrains actuation magnitude, duty cycle, and phase behavior. This integrated approach enables optimization within a reduced, physically meaningful parameter space defined by satellite constellation family architecture.

A representative mathematical formulation is given by

$$\min_{\{\mathbf{u}_f^{\text{ctrl}}(t)\}_{f \in \mathcal{F}_{\text{fam}}}} \sum_{f \in \mathcal{F}_{\text{fam}}} w_f \mathbb{E}_{\widehat{\rho}, BC, \theta_B} \left[J_f(\mathbf{u}_f^{\text{ctrl}}; \widehat{\rho}, BC, \theta_B) \right] + w_{\text{sep}} R_{\text{sep}} + w_{\text{res}} R_{\text{res}} + w_{\text{pmd}} R_{\text{pmd}} \quad (7.2)$$

In this formulation, f indexes a recovered family, family medoid, or shell-conditioned representative group, and $\mathcal{F}_{\text{fam}}$ denotes the corresponding family index set. The term w_f encodes population size, mission priority, risk contribution, or operational importance, while J_f formalizes the family-level mission objective. The expectation is evaluated over forecast-conditioned atmospheric density, effective ballistic-coefficient uncertainty, and reduced phase-conditioned actuation uncertainty. The penalty terms R_{sep} , R_{res} , and R_{pmd} enforce geometric separation, shell-residence efficiency, and post-mission-disposal compliance, respectively. The nonnegative weights w_{sep} , w_{res} , and w_{pmd} determine the relative importance of these constraints. This formulation establishes an optimization paradigm that explicitly integrates uncertainty, family structure, and operational lifecycle requirements.

The mesoscale perspective supports a range of applications. In shell maintenance, propagating family-conditioned medoids under forecast density and reduced thrust priors produces control schedules that maintain altitude corridors, nodal spacing, and along-track separation with fewer decision variables than satellite-level approaches. In anomaly detection, deviations from a family-specific posterior envelope provide contextually meaningful signals because the reference baseline incorporates shared shell geometry, environmental forcing, and operational mode.

Re-entry and disposal forecasting are improved through posterior-predictive drag ensembles, which provide uncertainty distributions for threshold-crossing times. When combined with forecast-conditioned density trajectories, these ensembles enable probabilistic disposal assessments for internal planning, regulatory analysis, and mission-end risk characterization. For conjunction risk management, family-aware state compression, object-specific drag posteriors, and constrained actuation priors improve forward state prediction and maneuver-plausibility modeling, thereby supporting risk ranking and maneuver triage.

Reverse engineering satellite constellation behavior from public records becomes tractable through phase-conditioned reduced models and clustering-derived family structure. This approach enables mining for deployment patterns, shell-occupancy control, disposal routing, and actuation intensity across orbital populations. The resulting behavioral models remain approximate because public catalog records do not contain operator telemetry, maneuver commands, attitude histories, or proprietary ephemerides. Even under these limitations, the family-conditioned workflow yields a substantially more informative representation of satellite constellation behavior than analyses based only on disjoint orbital-element lists.

The integrated framework advances methodological practice in astrodynamics by incorporating machine learning as a support mechanism for physically grounded inference. Machine learning modules are introduced where data geometry, forcing complexity, or likelihood intractability requires flexible representation. UMAP and HDBSCAN recover geometric population structure; the

hybrid ResNeXt-iTransformer forecasts forcing variables for physics-based atmospheric modeling; Bayesian inference quantifies drag uncertainty; and simulation-based inference addresses actuation behavior under implicit likelihoods. Each component is coupled to a subsequent physical model, producing interpretable intermediate quantities such as families, forecast-conditioned densities, ballistic-coefficient posteriors, and phase-conditioned actuation parameters. The workflow, therefore, remains centered on scientific estimation, uncertainty propagation, and physically constrained simulation.

7.3 Limitations and Sources of Uncertainty

The structural features that provide utility to the framework also impose inherent limitations on its scope of inference. The modeled archive is constructed from public catalog products, which lack truth-state provenance, attitude telemetry, maneuver logs, onboard power histories, mass-flow records, and covariance information. Consequently, all analytical outcomes are conditioned by the information content of the public observation layer. Clustering is performed on synchronized catalog geometry, forecasting predicts external forcing and couples it to an empirical atmospheric model, ballistic-coefficient inference estimates effective drag response under the specified observation model and realized density field, and the propulsion posterior encodes phase-conditioned actuation within a reduced dynamical representation. The resulting outputs maintain scientific value as conditional estimates, but their interpretation must remain tied to the modeling hierarchy that produced them.

Temporal representation introduces an additional limitation since clustering produces synchronized structural snapshots of the satellite constellation, capturing family structure at specified epochs and supporting shell-conditioned organizational analysis. This representation does not constitute a dynamic state-space model for family evolution and does not explicitly estimate cluster birth, death, merger, splitting, reassignment, or transfer between operational regimes. Latent structures may arise from persistent architecture, occupancy phase, deployment chronology, local drag environment, incomplete plane populations, or catalog-sampling effects. Confidently separating these mechanisms would require a temporal clustering framework in which family membership evolves as a stochastic or deterministic state process.

Atmospheric model dependence further constrains the analysis. Forecast error and structured residuals propagate through the NRLMSIS-coupled density construction and enter the orbit simulator as forcing uncertainty. Forecast-conditioned density fields, therefore, transmit both mean environmental evolution and uncertainty geometry into downstream propagation. Because drag

acceleration depends nonlinearly on atmospheric density, relative velocity, altitude, and effective area, modest forcing discrepancies can accumulate into substantial long-horizon orbital dispersion. Explicit uncertainty propagation is therefore essential; reducing the thermosphere to a single deterministic reference trajectory would suppress a dominant source of forecast-state uncertainty.

Parameter semantics impose additional limitations. The effective ballistic coefficient is useful for orbit analysis, posterior-predictive simulation, and re-entry-window characterization, but it aggregates unresolved variability from atmospheric forcing, attitude evolution, projected-area changes, configuration history, and model discrepancy. Similarly, reduced propulsion parameters represent only the actuation signal recoverable from public mean-element histories. The parameters do not reconstruct internal thruster commands, detailed power-processing behavior, plume physics, or precise propellant mass flow. These effective parameters reflect the information geometry of the inverse problem and should be interpreted as catalog-consistent dynamical summaries.

Identifiability and model compression introduce additional constraints. The low-rank Stage B posterior is informative because it reveals which combinations of reduced actuation and drag-like parameters are supported by the selected summary operator. However, it also requires that optimization, forecasting, and control interpretation target identifiable parameter combinations. Controllers or decision rules based on weakly resolved posterior directions risk unwarranted confidence. Posterior widths quantify uncertainty within the selected simulator, summary statistics, priors, and neural posterior estimator, but do not exhaust all physical uncertainty outside the modeled parameterization.

Compounding uncertainty arises from the hierarchical integration of the framework components. Family assignments depend on feature representation, embedding hyperparameters, density-clustering choices, and shell decomposition. Forecasted forcing retains residual solar-rotation structure, intermediate-period variability, and evolving bias. Drag inference is conditional on the decay arc, density realization, and effective-area abstraction. Reduced actuation inference operates through a low-rank summary channel. Robust optimization, therefore, remains sensitive to multiple levels of model discrepancy. The framework provides value by localizing, propagating, and mathematically exposing these uncertainties, even when it cannot eliminate them.

These limitations delineate the principal uncertainty classes relevant to future satellite constellation analysis, denoted as

$$\mathcal{U}_{\text{tot}} = \{\mathcal{U}_{\text{meas}}, \mathcal{U}_{\text{model}}, \mathcal{U}_{\text{param}}, \mathcal{U}_{\text{struct}}, \mathcal{U}_{\text{comp}}\}$$

Here, $\mathcal{U}_{\text{meas}}$ denotes measurement and catalog uncertainty, including estimation error, synchronization error, mean-element interpretation, and missing covariance. The class $\mathcal{U}_{\text{model}}$ denotes model

discrepancy arising from atmospheric specification, reduced force models, surrogate maneuver schedules, and incomplete physical-state representation. The class $\mathcal{U}_{\text{param}}$ denotes parameter uncertainty encoded in ballistic-coefficient posteriors, phase-conditioned actuation posteriors, and posterior-predictive ensembles. The class $\mathcal{U}_{\text{struct}}$ denotes structural uncertainty induced by feature representation, shell decomposition, family assignment, temporal snapshotting, and summary-operator design. The class $\mathcal{U}_{\text{comp}}$ denotes computational approximation from reduced-order optimization, finite simulation budgets, numerical tolerances, emulator construction, and stochastic training. This taxonomy organizes the framework’s limitations into explicit sources of uncertainty, thereby supporting robust decision analysis for large LEO satellite constellations.

7.4 Strategic Recommendations

The preceding results support several recommendations grounded in the technical evidence developed across the thesis. These recommendations concern data products, ephemeris fidelity, family-level planning, disposal assessment, regulatory oversight, and methodological practice. Their common basis is that modern LEO satellite constellations exhibit structured mesoscale behavior that can be recovered, modeled, and propagated when public catalog records are interpreted through a hierarchical estimation framework.

The first recommendation concerns operational data products. Satellite constellation management benefits from structured archives that encode mission phase, maneuver state, and uncertainty information as primary data products. Phase segmentation, as developed in Chapter 6, enables meaningful modeling because dominant orbital signals depend strongly on operational state. A structured phase-metadata archive, denoted $\mathcal{D}_{\text{phase}}$, would record abstracted insertion, shell-maintenance, transfer, disposal, and collision-avoidance states without requiring disclosure of detailed command sequences. Standardized records of phase state, maneuver status, and uncertainty products would improve both internal modeling and external traffic coordination. In the present framework, independent phase labels make it possible to estimate reduced behavioral models with substantially less ambiguity. External disclosure of abstracted phase metadata would therefore improve conjunction assessment, maneuver-plausibility modeling, and public-data inference while preserving the distinction between operationally useful metadata and proprietary control logic.

The second recommendation emphasizes improved ephemeris fidelity. The gap between public catalog products and operator-maintained trajectory knowledge complicates many inverse problems, causing estimation effort to be spent separating observational artifacts from physical

behavior. Higher-fidelity ephemeris products, denoted $\mathcal{D}_{\text{eph}}$, should include consistent reference frames, covariance information Σ_{eph} , maneuver-status indicators χ_{man} , and clear provenance for the generating orbit-determination process. Sharing richer ephemeris products would reduce ambiguity in drag inference, low-thrust system identification, and conjunction risk assessment. Although the methodology developed in this thesis is designed to remain useful under public-data limitations, the results indicate that even incremental improvements in data richness can produce substantial analytical value. Covariance availability improves uncertainty propagation. Maneuver-state indicators reduce phase-label ambiguity. Consistent ephemeris formatting improves interoperability across operators, regulators, and civil space-traffic coordination systems.

The third recommendation proposes reframing satellite constellation planning around family representatives. A family-representative set $\mathcal{M}_{\text{fam}}$, containing medoids m_f for families $f \in \mathcal{F}_{\text{fam}}$, provides a practical basis for shell maintenance, drag-makeup analysis, disposal routing, and scenario evaluation. The representative objects preserve mesoscale geometry while reducing the dimensionality of the planning problem. Family-conditioned medoids enable efficient exploration of candidate control schedules, with full-fleet propagation reserved for selected policies after screening. Computational savings are valuable, but the stronger advantage is interpretability. A family-based control policy can be inspected at the level of shell structure, phase behavior, and environmental response. This representation is more auditable than a large collection of loosely coupled satellite-level decisions because each representative can be traced to a recoverable dynamical family and a physically meaningful region of the satellite constellation architecture.

The fourth recommendation concerns disposal and sustainability assessment. End-of-life behavior should be modeled probabilistically when uncertainty is dominated by drag, density variability, and phase-dependent actuation. A posterior distribution of threshold-crossing times provides more operational information than a single deterministic disposal prediction, especially for satellite constellations with heterogeneous ballistic coefficients and forecast-conditioned atmospheric environments. The methodology presented in Chapter 5 directly supports this probabilistic framework. Effective ballistic-coefficient posteriors generate posterior-predictive altitude and semi-major axis ensembles, which can be converted into re-entry windows, disposal-compliance probabilities, and uncertainty-aware lifecycle diagnostics. When combined with forecast-conditioned density fields from Chapter 4, these posterior-predictive products provide a defensible basis for internal planning, regulatory reporting, and mission-end risk analysis.

The fifth recommendation addresses regulatory oversight and civil space-traffic coordination. Oversight frameworks should incorporate family-aware indicators in addition to individual compliance snapshots. A family-aware regulatory indicator, denoted $\mathcal{I}_{\text{reg},f}$, may include shell-specific

residence-time distributions, family-conditioned disposal reliability, posterior drag-response distributions for late-life populations, or anomaly rates normalized by operational phase. These indicators better represent the dynamic structure of modern satellite constellations because systemic risk often emerges at the shell or family level. Civil space-traffic coordination can also benefit from a family-aware risk representation $\mathcal{R}_{\text{stm}}$, integrating state prediction, covariance, maneuver plausibility, phase metadata, and environmental uncertainty. Such a representation would support more informative risk ranking and maneuver triage than assessments based only on object-level proximity metrics.

The final recommendation is methodological and applies to both operators and public data analysts. Satellite constellation analysis is most effective when structured as a hierarchical estimation process, with geometric decomposition, forecast-conditioned environmental modeling, parameter inference, and behavioral system identification proceeding sequentially. Each analytic layer supplies essential information to the next. Clustering defines the family structure, forecasting supplies environmental forcing, Bayesian inference quantifies drag uncertainty, and system identification constrains phase-conditioned actuation. This hierarchy preserves structural information that is often lost when optimization begins directly from raw catalog histories. It also provides a transparent sequence of assumptions, intermediate products, and uncertainty representations. The integrated workflow developed in this thesis can therefore be interpreted as a decision-support architecture: public catalog data are transformed into family structure, environmental forcing, effective parameter posteriors, reduced actuation models, and uncertainty-aware propagated states. A policy or coordination layer Π_{policy} can then consume these products for planning, oversight, sustainability assessment, or space-traffic coordination.

Chapter 8

Conclusion

This work examined the potential of public orbital data, augmented by advanced machine learning, Bayesian inference, and high-fidelity propagation, to establish a principled analytical and optimization framework for large LEO satellite constellations. The findings indicate that public data can support such a framework, provided the inference problem is formulated at an appropriate level of abstraction, and each estimated quantity accurately reflects the information content of the archive.

This work examined the extent to which public orbital data, when augmented with manifold learning, Bayesian inference, forecast-conditioned thermospheric modeling, phase-conditioned system identification, and high-fidelity propagation, can support a principled analytical and optimization framework for large LEO satellite constellations. The results show that public GP and TLE archives contain substantial information about satellite constellation geometry, population structure, environmental response, and effective operational behavior. That information becomes scientifically useful when the inverse problem is posed at a level of abstraction consistent with the catalog observation layer. Each estimated quantity must therefore be interpreted as an effective, catalog-constrained representation of behavior, uncertainty, or forcing, not as direct telemetry or a reconstructed operator truth state.

8.1 Concluding Synthesis

Tractability in satellite constellation analysis is achieved by identifying and leveraging intermediate-scale structures. Shell-conditioned families provide the geometric and statistical basis for characterizing complex LEO satellite constellations whose organization cannot be understood solely through object-by-object inspection or shell-level aggregation. These families supply a mesoscale description of satellite constellation architecture, linking the orbital geometry recovered through

clustering to the reduced representations required for propagation, inference, and optimization. Accurate future propagation further requires thermospheric forcing informed by time-dependent solar and geomagnetic drivers, along with residual structures that remain visible after deterministic modeling. Heterogeneity in drag response across satellites requires posterior distributions over effective, object-specific parameters. Although the inverse problem remains rank deficient, reduced and phase-conditioned actuation behavior can still be recovered with sufficient fidelity for predictive modeling, uncertainty propagation, and operationally meaningful scenario analysis.

The custom simulator integrates these analytical results into a cohesive framework. Without this propagation environment, clustering, forecasting, Bayesian inference, and system identification would remain isolated methodological components. The simulator converts those outputs into forecast-conditioned state ensembles, posterior-predictive uncertainty descriptions, and family-aware propagation products applicable to future orbital evolution. This integration constitutes a central contribution of the thesis because it transforms public orbital records from retrospective catalog histories into structured inputs for predictive simulation and reduced-order optimization.

The framework also defines the scope of knowledge that can be extracted from public data. The archive supports recovery of behavioral regularities, effective parameter distributions, shell-conditioned family structure, and structured dynamical priors. It cannot, by itself, determine proprietary maneuver commands, attitude histories, covariance products, or operator-maintained truth states. This boundary is methodologically important because it separates scientifically defensible inference from unsupported reconstruction. The inferred effective parameters become the primary entities for simulation, uncertainty propagation, and reduced-order optimization. Their interpretation is behavioral and catalog-constrained, representing the portion of satellite and environmental response that is identifiable from the available public record.

The synthesis of results demonstrates that environmental, drag, and control uncertainties are central elements of the satellite constellation-analysis problem. Large LEO satellite constellations evolve under time-dependent atmospheric density, heterogeneous satellite responses, partially observed actuation, and operational phase structure. Frameworks that omit these effects may support simplified sensitivity studies, yet provide limited insight into the mechanisms that govern long-horizon satellite constellation evolution. The present work incorporates these uncertainties directly within the model architecture by connecting forecast-conditioned thermospheric forcing, posterior drag inference, phase-conditioned actuation recovery, and high-fidelity propagation.

The primary advancement in optimization methodology is the transition from brute-force fleet-scale analysis to robust, family-conditioned control. Shell-conditioned clustering reveals mesoscale structures, while posterior inference quantifies dominant uncertainties at the scale where optimiza-

tion becomes computationally tractable. Optimization over families, medoids, or representative satellites enhances computational efficiency while preserving the geometric organization and uncertainty structure governing the fleet. This thesis therefore provides the reduced state descriptions, effective parameter distributions, and forecast-conditioned propagation tools required for physically grounded satellite constellation optimization.

8.2 Future Research Directions

Several research directions emerge from the present findings. The first concerns temporal family tracking. Dynamic clustering, sequential embedding alignment, or latent-state modeling can link synchronized shell-conditioned snapshots across epochs, enabling quantification of family persistence, cluster formation and dissolution, deployment sequences, occupancy transitions, and topology evolution associated with maneuvers and disposal. This extension would advance the static taxonomy developed in the present analysis toward a temporal characterization of satellite constellation morphology, in which family structure becomes an evolving state variable rather than a single-epoch descriptor.

The second direction focuses on hierarchical drag inference. The ballistic-coefficient results reveal significant object-level heterogeneity and discernible group structure. A multilevel Bayesian model that incorporates family membership, shell altitude, lifecycle phase, and environmental regime as partial-pooling predictors for effective drag response would preserve satellite-specific inference while utilizing structural relationships across the satellite constellation. Such a model would be especially valuable for satellites with short, noisy, or sparsely sampled decay arcs because information could be shared across physically and geometrically related objects without eliminating individual variability.

The third direction involves joint inference over density and drag. The current workflow conditions on forecasted density realizations before aerodynamic inference. A more advanced formulation would model the thermospheric forcing field and satellite drag state as coupled latent processes. Ensemble thermospheric scenarios, residual-informed stochastic forcing, or simulator-based joint posterior inference could support a more principled separation of atmospheric uncertainty from satellite-specific uncertainty. This extension would directly address one of the central identifiability challenges in LEO orbit prediction, where density error and drag-parameter error can produce similar orbital signatures.

The fourth direction addresses more expressive propulsion inversion. The reduced phase-conditioned model reveals identifiability structure while maintaining computational tractability.

Incorporating maneuver-mode transitions, duty-cycle scheduling, power and attitude dependencies, and shell- or family-conditioned priors would enhance predictive adequacy and refine the mapping between orbital signatures and plausible control actions. Such extensions should explicitly preserve uncertainty over weakly identifiable parameter regimes, since the inverse problem remains underdetermined when only public mean-element histories are available.

The fifth direction entails direct integration with conjunction and disposal analysis. Family- and uncertainty-aware states established in this work provide a foundation for covariance-aware conjunction screening, maneuver triage, and disposal reliability assessment. A unified framework could integrate posterior-predictive propagation, family-conditioned maneuver priors, and probabilistic disposal forecasts within a single decision-support environment. This would support operational analysis by identifying families or representatives with elevated propagation uncertainty, anomalous maneuver behavior, or increased disposal-risk sensitivity, while also supporting regulatory evaluation through repeatable family-level metrics.

The sixth direction is the full implementation of cluster-level optimization. The conceptual framework has been established, and the remaining challenge is to develop robust solvers that propagate medoid- or family-level decisions through high-fidelity simulation under forecast-conditioned density and posterior uncertainty. Model-predictive control, stochastic programming, and distributionally robust optimization are promising methodologies because these can accommodate receding-horizon decisions, uncertain forcing, and risk-sensitive constraints. The reduced-order representation developed here provides computational efficiency and physical interpretability well suited for this application, particularly when optimization is performed over families, medoids, or representative subpopulations instead of the full satellite constellation state.

The final and potentially most consequential direction concerns data fusion. The public archive enables substantial analysis; however, even limited integration with high-fidelity ephemerides, maneuver metadata, covariance products, phase labels, or structured operational data products would significantly expand inferential capacity. The methods presented here provide a foundation designed for sparse public data while remaining compatible with richer trajectory and uncertainty products. This compatibility is important because the same family-aware inference framework can operate under public-data constraints, partial operator disclosure, or higher-fidelity exchange protocols.

Collectively, these directions converge on a long-term objective: the development of a family-aware digital representation of large satellite constellations that is dynamically calibrated, uncertainty aware, and directly actionable for robust planning. This thesis establishes the methodological foundation necessary to pursue that objective, including shell-conditioned geometric decomposi-

tion, forecast-conditioned environmental forcing, Bayesian estimation of effective drag response, reduced identification of phase-conditioned actuation, and high-fidelity propagation integrated within a unified simulation and optimization framework.

References

- [1] David A. Vallado. *Fundamentals of Astrodynamics and Applications*. 5th ed. Hawthorne, CA: Microcosm Press, 2022. ISBN: 9781881883227.
- [2] Oliver Montenbruck and Eberhard Gill. *Satellite Orbits: Models, Methods, and Applications*. Berlin, Heidelberg: Springer, 2000. ISBN: 9783540672807. DOI: 10.1007/978-3-642-58351-3.
- [3] Hanspeter Schaub and John L. Junkins. *Analytical Mechanics of Space Systems*. 4th ed. AIAA Education Series. Reston, VA: American Institute of Aeronautics and Astronautics, 2018. ISBN: 9781624105210. DOI: 10.2514/4.105210.
- [4] Jonathan C. McDowell. “The Low Earth Orbit Satellite Population and Impacts of the SpaceX Starlink Constellation”. In: *The Astrophysical Journal Letters* 892.2 (Apr. 2020), p. L36. DOI: 10.3847/2041-8213/ab8016.
- [5] Aaron C. Boley and Michael Byers. “Satellite mega-constellations create risks in Low Earth Orbit, the atmosphere and on Earth”. In: *Scientific Reports* 11 (2021), p. 10642. DOI: 10.1038/s41598-021-89909-7.
- [6] Donald J. Kessler and Burton G. Cour-Palais. “Collision Frequency of Artificial Satellites: The Creation of a Debris Belt”. In: *Journal of Geophysical Research* 83.A6 (1978), pp. 2637–2646. DOI: 10.1029/ja083ia06p02637.
- [7] United States Government. *U.S. Government Orbital Debris Mitigation Standard Practices*. Tech. rep. United States Government, Nov. 2019. URL: https://orbitaldebris.jsc.nasa.gov/library/usg_orbital_debris_mitigation_standard_practices_november_2019.pdf.
- [8] National Aeronautics and Space Administration. *Process for Limiting Orbital Debris*. NASA Technical Standard NASA-STD-8719.14C. National Aeronautics and Space Administration, Nov. 2021. URL: <https://standards.nasa.gov/standard/NASA/NASA-STD-871914>.

- [9] International Organization for Standardization. *Space Systems: Space Debris Mitigation Requirements*. International Standard ISO 24113:2023. International Organization for Standardization, May 2023. URL: <https://www.iso.org/standard/83494.html>.
- [10] Electronic Code of Federal Regulations. *47 CFR §25.283: End-of-life Disposal*. Electronic Code of Federal Regulations. 2026. URL: <https://www.ecfr.gov/current/title-47/chapter-I/subchapter-B/part-25/subpart-D/section-25.283> (visited on 05/07/2026).
- [11] Hanspeter Schaub and Kyle T. Alfriend. “ J_2 Invariant Relative Orbits for Spacecraft Formations”. In: *Celestial Mechanics and Dynamical Astronomy* 79.2 (2001), pp. 77–95. DOI: 10.1023/a:1011161811472.
- [12] Eelco Doornbos. *Thermospheric Density and Wind Determination from Satellite Dynamics*. Springer Theses. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012. ISBN: 978-3-642-25128-3. DOI: 10.1007/978-3-642-25129-0.
- [13] J. G. Walker. “Satellite Constellations”. In: *Journal of the British Interplanetary Society* 37 (1984), pp. 559–572.
- [14] A. H. Ballard. “Rosette Constellations of Earth Satellites”. In: *IEEE Transactions on Aerospace and Electronic Systems* 16.5 (1980), pp. 656–673. DOI: 10.1109/taes.1980.308932.
- [15] Simone D’Amico. “Autonomous Formation Flying in Low Earth Orbit”. PhD thesis. Delft, The Netherlands: Delft University of Technology, 2010.
- [16] Felix R. Hoots and Ronald L. Roehrich. *Models for Propagation of NORAD Element Sets*. Spacetrack Report No. 3. Colorado Springs, CO: U.S. Air Force Aerospace Defense Command, 1980.
- [17] David A. Vallado et al. “Revisiting Spacetrack Report #3”. In: *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*. AIAA 2006-6753. Keystone, CO, 2006. DOI: 10.2514/6.2006-6753.
- [18] Space-Track.org. *Help Documentation: General Perturbations and Two-Line Element Format*. Online documentation. 2026. URL: <https://www.space-track.org/documentation> (visited on 05/07/2026).
- [19] T. S. Kelso. *A New Way to Obtain GP Data, Also Known as TLEs*. Celestrak documentation. May 2020. URL: <https://celestrak.org/NORAD/documentation/gp-data-formats.php> (visited on 05/07/2026).

- [20] Consultative Committee for Space Data Systems. *Orbit Data Messages*. Recommended Standard CCSDS 502.0-B-3. Washington, DC: Consultative Committee for Space Data Systems, Apr. 2023. URL: <https://ccsds.org/publications/bluebooks/>.
- [21] Inter-Agency Space Debris Coordination Committee. *IADC Space Debris Mitigation Guidelines*. IADC-02-01, Revision 2. Inter-Agency Space Debris Coordination Committee, Mar. 2020. URL: <https://orbitaldebris.jsc.nasa.gov/library/iadc-space-debris-guidelines-revision-2.pdf>.
- [22] Daniele Mortari, Matthew P. Wilkins, and Christian Bruccoleri. “The Flower Constellations”. In: *The Journal of the Astronautical Sciences* 52.1-2 (2004), pp. 107–127. DOI: 10.1007/bf03546424.
- [23] N. H. Crisp et al. “The Benefits of Very Low Earth Orbit for Earth Observation Missions”. In: *Progress in Aerospace Sciences* 117 (2020), p. 100619. DOI: 10.1016/j.paerosci.2020.100619.
- [24] James R. Wertz. *Mission Geometry: Orbit and Constellation Design and Management*. Spacecraft Orbit and Attitude Systems. El Segundo, CA: Kluwer Academic Publishers, 2001. ISBN: 9780792371489.
- [25] Hang Woon Lee et al. “Satellite Constellation Pattern Optimization for Complex Regional Coverage”. In: *Journal of Spacecraft and Rockets* 57.6 (2020), pp. 1309–1327. DOI: 10.2514/1.a34657.
- [26] Simeng Huang, Camilla Colombo, and Franco Bernelli-Zazzera. “Multi-Criteria Design of Continuous Global Coverage Walker and Street-of-Coverage Constellations through Property Assessment”. In: *Acta Astronautica* 188 (2021), pp. 151–170. DOI: 10.1016/j.actaastro.2021.07.002.
- [27] Luca Mencarelli et al. “Mixed Integer (Non)Linear Approaches for the Satellite Constellation Design Problem”. In: *Optimization and Engineering* 24 (2023), pp. 2299–2320. DOI: 10.1007/s11081-022-09774-9.
- [28] Consultative Committee for Space Data Systems. *Navigation Data: Definitions and Conventions*. Report Concerning Space Data System Standards, Green Book CCSDS 500.0-G-4. Washington, DC: Consultative Committee for Space Data Systems, Nov. 2019. URL: <https://ccsds.org/publications/greenbooks/>.
- [29] Dirk Brouwer. “Solution of the Problem of Artificial Satellite Theory without Drag”. In: *The Astronomical Journal* 64 (1959), pp. 378–397. DOI: 10.1086/107958.

- [30] Yoshihide Kozai. “The Motion of a Close Earth Satellite”. In: *The Astronomical Journal* 64 (1959), pp. 367–377. DOI: 10.1086/107957.
- [31] Martin Lara. “On Mean Elements in Artificial-Satellite Theory”. In: *Celestial Mechanics and Dynamical Astronomy* 135.4 (2023), p. 43. DOI: 10.1007/s10569-023-10153-4.
- [32] R. A. Broucke and P. J. Cefola. “On the Equinoctial Orbit Elements”. In: *Celestial Mechanics* 5 (1972), pp. 303–310. DOI: 10.1007/bf01228432.
- [33] P. J. Cefola. “Equinoctial Orbit Elements: Application to Artificial Satellite Orbits”. In: *AIAA/AAS Astrodynamics Conference*. AIAA Paper 72-937. Palo Alto, CA, 1972. DOI: 10.2514/6.1972-937.
- [34] J. M. Picone et al. “NRLMSISE-00 Empirical Model of the Atmosphere: Statistical Comparisons and Scientific Issues”. In: *Journal of Geophysical Research: Space Physics* 107.A12 (2002), p. 1468. DOI: 10.1029/2002ja009430.
- [35] J. T. Emmert et al. “NRLMSIS 2.0: A Whole-Atmosphere Empirical Model of Temperature and Neutral Species Densities”. In: *Earth and Space Science* 8.3 (2021), e2020EA001321. DOI: 10.1029/2020ea001321.
- [36] Sean Bruinsma et al. “Thermosphere and Satellite Drag”. In: *Advances in Space Research* 72.9 (2023), pp. 3714–3730. DOI: 10.1016/j.asr.2023.05.011.
- [37] Nikolaos K. Pavlis et al. “The Development and Evaluation of the Earth Gravitational Model 2008 (EGM2008)”. In: *Journal of Geophysical Research: Solid Earth* 117.B4 (2012), B04406. DOI: 10.1029/2011jb008916.
- [38] Kenneth Moe and Mildred M. Moe. “Gas-Surface Interactions and Satellite Drag Coefficients”. In: *Planetary and Space Science* 53.8 (2005), pp. 793–801. DOI: 10.1016/j.pss.2005.03.005.
- [39] Piyush M. Mehta et al. “New Density Estimates Derived Using Accelerometers on Board the CHAMP and GRACE Satellites”. In: *Space Weather* 15.4 (2017), pp. 558–576. DOI: 10.1002/2016sw001562.
- [40] Bruce R. Bowman et al. “A New Empirical Thermospheric Density Model JB2008 Using New Solar and Geomagnetic Indices”. In: *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*. AIAA 2008-6438. Honolulu, HI, 2008. DOI: 10.2514/6.2008-6438.
- [41] Sean Bruinsma and Christine Boniface. “The Operational and Research DTM-2020 Thermosphere Models”. In: *Journal of Space Weather and Space Climate* 11 (2021), p. 47. DOI: 10.1051/swsc/2021032.

- [42] A. D. Richmond, E. C. Ridley, and R. G. Roble. “A Thermosphere/Ionosphere General Circulation Model with Coupled Electrodynamics”. In: *Geophysical Research Letters* 19.6 (1992), pp. 601–604. DOI: 10.1029/92gl00401.
- [43] Cyrus Foster et al. “Constellation Phasing with Differential Drag on Planet Labs Satellites”. In: *Journal of Spacecraft and Rockets* 55.2 (2018), pp. 473–483. DOI: 10.2514/1.a33927.
- [44] Riccardo Bevilacqua and Marcello Romano. “Rendezvous Maneuvers of Multiple Spacecraft Using Differential Drag Under J_2 Perturbation”. In: *Journal of Guidance, Control, and Dynamics* 31.6 (2008), pp. 1595–1607. DOI: 10.2514/1.36362.
- [45] Ciara N. McGrath and Malcolm Macdonald. “General Perturbation Method for Satellite Constellation Deployment Using Nodal Precession”. In: *Journal of Guidance, Control, and Dynamics* 43.4 (2020), pp. 814–824. DOI: 10.2514/1.g004443.
- [46] Ciara N. McGrath and Malcolm Macdonald. “General Perturbation Method for Satellite Constellation Reconfiguration Using Low-Thrust Maneuvers”. In: *Journal of Guidance, Control, and Dynamics* 42.8 (2019), pp. 1676–1692. DOI: 10.2514/1.g003739.
- [47] W. H. Clohessy and R. S. Wiltshire. “Terminal Guidance System for Satellite Rendezvous”. In: *Journal of the Aerospace Sciences* 27.9 (1960), pp. 653–658. DOI: 10.2514/8.8704.
- [48] Koji Yamanaka and Finn Ankersen. “New State Transition Matrix for Relative Motion on an Arbitrary Elliptical Orbit”. In: *Journal of Guidance, Control, and Dynamics* 25.1 (2002), pp. 60–66. DOI: 10.2514/2.4875.
- [49] J. Tschauner and P. Hempel. “Rendezvous zu einem in elliptischer Bahn umlaufenden Ziel”. In: *Astronautica Acta* 11 (1965), pp. 104–109.
- [50] Simone D’Amico and Oliver Montenbruck. “Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation”. In: *Journal of Guidance, Control, and Dynamics* 29.3 (2006), pp. 554–563. DOI: 10.2514/1.15114.
- [51] Gabriella Gaias and Simone D’Amico. “Impulsive Maneuvers for Formation Reconfiguration Using Relative Orbital Elements”. In: *Journal of Guidance, Control, and Dynamics* 38.6 (2015), pp. 1036–1049. DOI: 10.2514/1.g000189.
- [52] Russell P. Patera. “General Method for Calculating Satellite Collision Probability”. In: *Journal of Guidance, Control, and Dynamics* 24.4 (2001), pp. 716–722. DOI: 10.2514/2.4771.

- [53] Salvatore Alfano. “Relating Position Uncertainty to Maximum Conjunction Probability”. In: *The Journal of the Astronautical Sciences* 53.2 (2005), pp. 193–205. doi: 10.1007/bf03546350.
- [54] J.-C. Liou and N. L. Johnson. “Risks in Space from Orbiting Debris”. In: *Science* 311.5759 (2006), pp. 340–341. doi: 10.1126/science.1121337.
- [55] J.-C. Liou and N. L. Johnson. “A Sensitivity Study of the Effectiveness of Active Debris Removal in LEO”. In: *Acta Astronautica* 64.2–3 (2009), pp. 236–243. doi: 10.1016/j.actaastro.2008.07.009.
- [56] Heiner Klinkrad. *Space Debris: Models and Risk Analysis*. Berlin and Heidelberg: Springer Praxis Books, 2006. doi: 10.1007/3-540-37674-7.
- [57] Salvatore Alfano. “Satellite Collision Probability Enhancements”. In: *Journal of Guidance, Control, and Dynamics* 29.3 (2006), pp. 588–592. doi: 10.2514/1.15523.
- [58] Federal Communications Commission. *Mitigation of Orbital Debris in the New Space Age, Second Report and Order*. FCC 22-74, IB Docket No. 18-313. Sept. 2022. URL: <https://docs.fcc.gov/public/attachments/FCC-22-74A1.pdf>.
- [59] Srinivas J. Setty et al. “Application of Semi-analytical Satellite Theory Orbit Propagator to Orbit Determination for Space Object Catalog Maintenance”. In: *Advances in Space Research* 57.10 (2016), pp. 2218–2233. doi: 10.1016/j.asr.2016.02.028.
- [60] Federal Communications Commission. *Space Exploration Holdings, LLC: Application for Approval for Orbital Deployment and Operating Authority for the SpaceX NGSO Satellite System, Memorandum Opinion, Order and Authorization*. FCC 18-38, IBFS File Nos. SAT-LOA-20161115-00118 and SAT-LOA-20170726-00110. Adopted March 28, 2018; released March 29, 2018. Washington, DC, Mar. 2018. URL: <https://docs.fcc.gov/public/attachments/FCC-18-38A1.pdf>.
- [61] Federal Communications Commission. *Space Exploration Holdings, LLC, Request for Modification of the Authorization for the SpaceX NGSO Satellite System, Order and Authorization and Order on Reconsideration*. Tech. rep. FCC 21-48. 36 FCC Rcd 7995. Federal Communications Commission, Apr. 2021. URL: <https://docs.fcc.gov/public/attachments/FCC-21-48A1.pdf>.
- [62] T. E. Berger et al. “The Thermosphere Is a Drag: The 2022 Starlink Incident and the Threat of Geomagnetic Storms to Low Earth Orbit Space Operations”. In: *Space Weather* 21.3 (2023), e2022SW003330. doi: 10.1029/2022sw003330.

- [63] Airong Liu et al. “Maneuver Strategies of Starlink Satellite Based on SpaceX-Released Ephemeris”. In: *Advances in Space Research* 74.7 (2024), pp. 3157–3169. DOI: 10.1016/j.asr.2024.06.038.
- [64] David P. Shorten, Wathsala Karunaratne, and Matthew Roughan. “How Is Starlink Manoeuvring? An Analysis of Patterns in the Manoeuvres of Starlink Satellites”. In: *Proceedings of the 9th International Conference on Internet of Things, Big Data and Security*. SCITEPRESS, 2024, pp. 174–184. DOI: 10.5220/0012619900003705.
- [65] Bin Li, Lei Liu, and Ji-Zhang Sang. “Tracklet-to-Object Matching for Climbing Starlink Satellites through Recursive Orbit Determination and Prediction”. In: *Research in Astronomy and Astrophysics* 22.11 (2022), p. 115010. DOI: 10.1088/1674-4527/ac91be.
- [66] Denny M. Oliveira, Eftyhia Zesta, and Katherine Garcia-Sage. “Tracking Reentries of Starlink Satellites during the Rising Phase of Solar Cycle 25”. In: *Frontiers in Astronomy and Space Sciences* 12 (2025), p. 1572313. DOI: 10.3389/fspas.2025.1572313.
- [67] LLC Space Exploration Holdings. *Application for Modification of Authorization for the SpaceX NGSO Satellite System*. IBFS File No. SAT-MOD-20200417-00037, Narrative. Filed April 17, 2020. Apr. 2020. URL: <https://fcc.report/IBFS/SAT-MOD-20200417-00037/2274315.pdf>.
- [68] Anqi Lang and Yu Jiang. “Orbit Determination for Continuously Maneuvering Starlink Satellites Based on an Unscented Batch Filtering Method”. In: *Sensors* 25.13 (2025), p. 4079. DOI: 10.3390/s25134079.
- [69] LLC Space Exploration Holdings. *SpaceX Non-Geostationary Satellite System: Attachment A, Technical Information to Supplement Schedule S*. IBFS File No. SAT-LOA-20161115-00118, Attachment A. Filed November 15, 2016. Nov. 2016. URL: <https://fcc.report/IBFS/SAT-LOA-20161115-00118/1158350.pdf>.
- [70] LLC Space Exploration Holdings. *SpaceX Non-Geostationary Satellite System: Attachment A, Technical Information to Supplement Schedule S*. IBFS File No. SAT-MOD-20181108-00083, Attachment A. Filed November 8, 2018; request to modify the authorized SpaceX NGSO system. Nov. 2018. URL: <https://fcc.report/IBFS/SAT-MOD-20181108-00083/1569860.pdf>.
- [71] International Earth Rotation and Reference Systems Service. *IERS Conventions (2010)*. IERS Technical Note 36. Frankfurt am Main: International Earth Rotation and Reference Systems Service, 2010. URL: <https://www.iers.org/IERS/EN/Publications/TechnicalNotes/tn36>.

- [72] JPL Solar System Dynamics Group. *NASA/JPL Horizons On-Line Ephemeris System Manual*. Jet Propulsion Laboratory, California Institute of Technology. Version 4.98d, November 21, 2025; accessed May 7, 2026. 2025. URL: <https://ssd.jpl.nasa.gov/horizons/manual.html>.
- [73] Navigation and Ancillary Information Facility. *SPK Required Reading*. NASA Jet Propulsion Laboratory, SPICE Toolkit Documentation. Last revised August 31, 2021. 2021. URL: https://naif.jpl.nasa.gov/pub/naif/toolkit_docs/C/req/spk.html.
- [74] Ernst Hairer, Syvert P. Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*. 2nd ed. Vol. 8. Springer Series in Computational Mathematics. Springer, 1993. ISBN: 978-3-540-56670-0. DOI: 10.1007/978-3-540-78862-1.
- [75] SciPy Developers. *scipy.integrate.solve_ivp*. SciPy v1.17.0 Manual. Accessed May 7, 2026. 2026. URL: https://docs.scipy.org/doc/scipy/reference/generated/scipy.integrate.solve_ivp.html.
- [76] Pauli Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17.3 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.
- [77] SciPy Developers. *scipy.integrate.DOP853*. SciPy v1.17.0 Manual. Accessed May 7, 2026. 2026. URL: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.integrate.DOP853.html>.
- [78] Andrea Milani, Anna Maria Nobili, and Paolo Farinella. *Non-Gravitational Perturbations and Satellite Geodesy*. Bristol: Adam Hilger, 1987. ISBN: 0-85274-538-9.
- [79] Marek Ziebart. “Generalized Analytical Solar Radiation Pressure Modeling Algorithm for Spacecraft of Complex Shape”. In: *Journal of Spacecraft and Rockets* 41.5 (2004), pp. 840–848. DOI: 10.2514/1.13097.
- [80] U.S. Committee on Extension to the Standard Atmosphere. *U.S. Standard Atmosphere, 1976*. Tech. rep. NOAA-S/T 76-1562; NASA-TM-X-74335. Washington, DC: National Oceanic et al., 1976. URL: https://www.ngdc.noaa.gov/stp/space-weather/online-publications/miscellaneous/us-standard-atmosphere-1976/us-standard-atmosphere_st76-1562_noaa.pdf.
- [81] Community Coordinated Modeling Center. *NRLMSIS 2.0*. NASA Goddard Space Flight Center, CCMC Model Catalog. Accessed May 7, 2026. 2026. URL: <https://ccmc.gsfc.nasa.gov/models/NRLMSIS~2.0/>.

- [82] National Geospatial-Intelligence Agency. *Department of Defense World Geodetic System 1984: Its Definition and Relationships with Local Geodetic Systems*. Tech. rep. NGA.STND.0036_1.0.0_WC Version 1.0.0. Springfield, VA: National Geospatial-Intelligence Agency, July 2014. URL: <https://nsgreg.nga.mil/doc/view?i=4085>.
- [83] Dan M. Goebel and Ira Katz. *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*. JPL Space Science and Technology Series. Hoboken, NJ: John Wiley & Sons, 2008. DOI: 10.1002/9780470436448.
- [84] George P. Sutton and Oscar Biblarz. *Rocket Propulsion Elements*. 9th ed. Hoboken, NJ: John Wiley & Sons, 2017. ISBN: 978-1-118-75365-1.
- [85] Greg Wilson et al. “Good Enough Practices in Scientific Computing”. In: *PLOS Computational Biology* 13.6 (2017), e1005510. DOI: 10.1371/journal.pcbi.1005510.
- [86] Python Software Foundation. *concurrent.futures: Launching Parallel Tasks*. Python 3.14 Documentation. Accessed May 7, 2026. 2026. URL: <https://docs.python.org/3/library/concurrent.futures.html>.
- [87] Mark D. Wilkinson et al. “The FAIR Guiding Principles for Scientific Data Management and Stewardship”. In: *Scientific Data* 3 (2016), p. 160018. DOI: 10.1038/sdata.2016.18.
- [88] United States Space Command. *Space-Track Documentation*. <https://www.space-track.org/documentation>. General Perturbation data, TLE, OMM, SATCAT, and historical element-set documentation; accessed 2026-05-07. 2026.
- [89] Anil K. Jain. “Data Clustering: 50 Years Beyond K-Means”. In: *Pattern Recognition Letters* 31.8 (2010), pp. 651–666. DOI: 10.1016/j.patrec.2009.09.011.
- [90] Christian Hennig. “What Are the True Clusters?” In: *Pattern Recognition Letters* 64 (2015), pp. 53–62. DOI: 10.1016/j.patrec.2015.04.009.
- [91] Martin Ester et al. “A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise”. In: *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*. KDD’96. AAAI Press, 1996, pp. 226–231. URL: <https://dl.acm.org/doi/10.5555/3001460.3001507>.
- [92] Ricardo J. G. B. Campello et al. “Hierarchical Density Estimates for Data Clustering, Visualization, and Outlier Detection”. In: *ACM Transactions on Knowledge Discovery from Data* 10.1 (2015), pp. 1–51. ISSN: 1556-4681. DOI: 10.1145/2733381.

- [93] Leland McInnes, John Healy, and Steve Astels. “hdbscan: Hierarchical Density Based Clustering”. In: *Journal of Open Source Software* 2.11 (2017), p. 205. doi: 10.21105/joss.00205.
- [94] Leland McInnes, John Healy, and James Melville. *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*. 2020. doi: 10.48550/arxiv.1802.03426. arXiv: 1802.03426 [stat.ML].
- [95] Leland McInnes et al. “UMAP: Uniform Manifold Approximation and Projection”. In: *Journal of Open Source Software* 3.29 (2018), p. 861. doi: 10.21105/joss.00861.
- [96] United States Space Force. *Global Positioning System Standard Positioning Service Performance Standard*. Tech. rep. Defines the baseline 24-slot constellation in six orbital planes with four slots per plane. United States Space Force, Apr. 2020. URL: <https://www.gps.gov/sites/default/files/2025-07/2020-SPS-performance-standard.pdf>.
- [97] Federal Communications Commission. *WorldVu Satellites Limited, Petition for Declaratory Ruling to Modify the U.S. Market Access Grant for the OneWeb Ku-Band and Ka-Band NGSO FSS System, Order and Declaratory Ruling*. Tech. rep. DA 23-362. 38 FCC Rcd 5366; revised Phase 1 configuration includes 49 satellites per plane in 12 orbital planes at 87.9 degrees and 16 satellites per plane in 8 orbital planes at 55 degrees. Federal Communications Commission, Apr. 2023. URL: <https://docs.fcc.gov/public/attachments/DA-23-362A1.pdf>.
- [98] Howard D. Curtis. *Orbital Mechanics for Engineering Students*. Fourth Edition. Butterworth-Heinemann, 2021, pp. 1–751. ISBN: 978-0-12-824025-0. doi: 10.1016/b978-0-12-824025-0.00005-2.
- [99] Roger R. Bate, Donald D. Mueller, and Jerry E. White. *Fundamentals of Astrodynamics*. Dover Books on Aeronautical Engineering. New York: Dover Publications, 1971. ISBN: 0486600610.
- [100] John M. Lee. *Introduction to Smooth Manifolds*. 2nd ed. Vol. 218. Graduate Texts in Mathematics. New York, NY: Springer, 2013. ISBN: 9781441999818. doi: 10.1007/978-1-4419-9982-5.
- [101] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. New York, NY: Springer, 2006. ISBN: 9780387310732. doi: 10.1007/978-0-387-45528-0.

- [102] Fabian Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830. URL: <https://jmlr.org/papers/v12/pedregosa11a.html>.
- [103] Ricardo J. G. B. Campello, Davoud Moulavi, and Jörg Sander. “Density-Based Clustering Based on Hierarchical Density Estimates”. In: *Advances in Knowledge Discovery and Data Mining*. Vol. 7819. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2013, pp. 160–172. DOI: 10.1007/978-3-642-37456-2_14.
- [104] Jarkko Venna and Samuel Kaski. “Neighborhood Preservation in Nonlinear Projection Methods: An Experimental Study”. In: *Artificial Neural Networks — ICANN 2001*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 485–491. ISBN: 978-3-540-44668-2. DOI: 10.1007/3-540-44668-0_68.
- [105] Peter J. Rousseeuw. “Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis”. In: *Journal of Computational and Applied Mathematics* 20 (1987), pp. 53–65. DOI: 10.1016/0377-0427(87)90125-7.
- [106] David L. Davies and Donald W. Bouldin. “A Cluster Separation Measure”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1.2 (1979), pp. 224–227. DOI: 10.1109/tpami.1979.4766909.
- [107] Davoud Moulavi et al. “Density-Based Clustering Validation”. In: *Proceedings of the 2014 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics, 2014, pp. 839–847. DOI: 10.1137/1.9781611973440.96.
- [108] National Coordination Office for Space-Based Positioning, Navigation, and Timing. *Space Segment*. <https://www.gps.gov/space-segment>. GPS.gov page, accessed 2026-04-13. 2026.
- [109] United States Coast Guard Navigation Center. *GPS Constellation*. <https://www.navcen.uscg.gov/gps-constellation>. Official NAVCEN constellation status and slant-chart page; accessed 2026-05-07. 2026.
- [110] Karl Pearson. “Notes on Regression and Inheritance in the Case of Two Parents”. In: *Proceedings of the Royal Society of London* 58.347–352 (1895), pp. 240–242. DOI: 10.1098/rsp1.1895.0041.
- [111] Patrick Schober, Christa Boer, and Lothar A. Schwarte. “Correlation Coefficients: Appropriate Use and Interpretation”. In: *Anesthesia & Analgesia* 126.5 (2018), pp. 1763–1768. DOI: 10.1213/ane.0000000000002864.

- [112] Ken F. Tapping. “The 10.7-cm Solar Radio Flux (F10.7)”. In: *Space Weather* 11.7 (2013), pp. 394–406. doi: 10.1002/swe.20064.
- [113] Liying Qian and Kalevi Mursula. “Evaluating F10.7 and F30 Radio Fluxes as Long-Term Solar Proxies of Energy Deposition in the Thermosphere”. In: *Annales Geophysicae* 43 (2025), pp. 175–182. doi: 10.5194/angeo-43-175-2025.
- [114] Natural Resources Canada. *About the Solar Flux Data*. <https://www.spaceweather.gc.ca/forecast-prevision/solar-solaire/solarflux/sx-3-en.php>. Official documentation for observed, adjusted, and URSI Series D F10.7 values; accessed 2026-04-13. Dec. 2025.
- [115] P. N. Mayaud. “Planetary Indices Derived from K Indices (Kp, am, and aa)”. In: *Derivation, Meaning, and Use of Geomagnetic Indices*. Washington, DC: American Geophysical Union, 1980. Chap. 5, pp. 40–85. doi: 10.1029/gm022.
- [116] Magali Menvielle and Alain Berthelier. “The K-Derived Planetary Indices: Description and Availability”. In: *Reviews of Geophysics* 29.3 (1991), pp. 415–432. doi: 10.1029/91rg00994.
- [117] Jürgen Matzka et al. *Geomagnetic Kp Index, Version 1.0*. GFZ Data Services, 2021. doi: 10.5880/kp.0001.
- [118] Masahisa Sugiura. *Hourly Values of Equatorial Dst for the IGY*. Vol. 35. Annals of the International Geophysical Year. Oxford: Pergamon Press, 1964.
- [119] Kyoto WDC for Geomagnetism. *Geomagnetic Equatorial Dst Index Home Page*. <https://wdc.kugi.kyoto-u.ac.jp/dstdir/>. Official Dst index distribution portal, accessed 2026-04-13. 2026.
- [120] François Clette and Laure Lefèvre. *SILSO Sunspot Number V2.0*. WDC-SILSO, Royal Observatory of Belgium. July 2015. doi: 10.24414/qnza-ac80.
- [121] Saining Xie et al. “Aggregated Residual Transformations for Deep Neural Networks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 1492–1500. doi: 10.1109/cvpr.2017.634.
- [122] Yong Liu et al. “iTransformer: Inverted Transformers Are Effective for Time Series Forecasting”. In: *International Conference on Learning Representations*. Also available as arXiv:2310.06625. 2024. doi: 10.48550/arxiv.2310.06625.
- [123] Emma Stevenson et al. “A Deep Learning Approach to Solar Radio Flux Forecasting”. In: *Acta Astronautica* 193 (2022), pp. 595–606. doi: 10.1016/j.actaastro.2021.08.004.

- [124] Junqi Luo et al. “Forecasting the 10.7-cm Solar Radio Flux Using Deep CNN-LSTM Neural Networks”. In: *Processes* 10.2 (2022), p. 262. doi: 10.3390/pr10020262.
- [125] Hua Zhang et al. “Deep Temporal Convolutional Networks for F10.7 Radiation Flux Short-Term Forecasting”. In: *Annales Geophysicae* 42 (2024), pp. 91–101. doi: 10.5194/angeo-42-91-2024.
- [126] CelesTrak. *Space Weather Data Documentation*. <https://celestrak.org/SpaceData/SpaceWx-format.php>. Documentation for SW-All.csv and related formats. May 2023.
- [127] NOAA National Centers for Environmental Information. *Geomagnetic Indices and Data*. <https://www.ncei.noaa.gov/products/geomagnetic-indices>. Official documentation for geomagnetic indices, including K, Kp, ap, and Dst definitions; accessed 2026-05-07. 2026.
- [128] Masahito Nose et al. *Geomagnetic Dst Index*. Data set, World Data Center for Geomagnetism, Kyoto. 2015. doi: 10.17593/14515-74000.
- [129] Mike Lockwood et al. “An Assessment of Sunspot Number Data Composites over 1845–2014”. In: *The Astrophysical Journal* 824.1 (2016), p. 54. doi: 10.3847/0004-637x/824/1/54.
- [130] William D. Pesnell. “Lessons Learned from Predictions of Solar Cycle 24”. In: *Journal of Space Weather and Space Climate* 10 (2020), p. 60. doi: 10.1051/swsc/2020060.
- [131] Chris Chatfield. *The Analysis of Time Series: An Introduction*. 6th. Boca Raton, FL: Chapman & Hall/CRC, 2003, pp. 1–329. doi: 10.4324/9780203491683.
- [132] Sean J. Taylor and Benjamin Letham. “Forecasting at Scale”. In: *The American Statistician* 72.1 (2018), pp. 37–45. doi: 10.1080/00031305.2017.1380080.
- [133] David H. Hathaway. “The Solar Cycle”. In: *Living Reviews in Solar Physics* 12 (2015), p. 4. doi: 10.1007/lrsp-2015-4.
- [134] K. Mursula and Th. Ulich. “A New Method to Determine the Solar Cycle Length”. In: *Geophysical Research Letters* 25.11 (1998), pp. 1837–1840. doi: 10.1029/98gl151317.
- [135] Ashish Vaswani et al. “Attention Is All You Need”. In: *Advances in Neural Information Processing Systems*. Vol. 30. 2017. URL: <https://papers.nips.cc/paper/7181-attention-is-all-you-need>.
- [136] Haoyi Zhou et al. “Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 12. 2021, pp. 11106–11115. doi: 10.1609/aaai.v35i12.17325.

- [137] Li Shen and Yangzhu Wang. “TCCT: Tightly-Coupled Convolutional Transformer on Time Series Forecasting”. In: *Neurocomputing* 480 (2022), pp. 131–145. doi: 10.1016/j.neucom.2022.01.039.
- [138] Bryan Lim et al. “Temporal Fusion Transformers for Interpretable Multi-Horizon Time Series Forecasting”. In: *International Journal of Forecasting* 37.4 (2021), pp. 1748–1764. doi: 10.1016/j.ijforecast.2021.03.012.
- [139] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. “Layer Normalization”. In: *arXiv preprint arXiv:1607.06450* (2016). doi: 10.48550/arxiv.1607.06450.
- [140] Souhaib Ben Taieb et al. “A Review and Comparison of Strategies for Multi-Step Ahead Time Series Forecasting Based on the NN5 Forecasting Competition”. In: *Expert Systems with Applications* 39.8 (2012), pp. 7067–7083. doi: 10.1016/j.eswa.2012.01.039.
- [141] Christoph Bergmeir and José M. Benítez. “On the Use of Cross-Validation for Time Series Predictor Evaluation”. In: *Information Sciences* 191 (2012), pp. 192–213. doi: 10.1016/j.ins.2011.12.028.
- [142] Vitor Cerqueira, Luís Torgo, and Igor Mozetič. “Evaluating Time Series Forecasting Models: An Empirical Study on Performance Estimation Methods”. In: *Machine Learning* 109.11 (2020), pp. 1997–2028. doi: 10.1007/s10994-020-05910-7.
- [143] Ilya Loshchilov and Frank Hutter. “Decoupled Weight Decay Regularization”. In: *International Conference on Learning Representations*. 2019. URL: <https://openreview.net/forum?id=Bkg6RiCqY7>.
- [144] Ilya Loshchilov and Frank Hutter. “SGDR: Stochastic Gradient Descent with Warm Restarts”. In: *International Conference on Learning Representations*. 2017. URL: <https://openreview.net/forum?id=Skq89Scxx>.
- [145] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. “On the Difficulty of Training Recurrent Neural Networks”. In: *Proceedings of the 30th International Conference on Machine Learning*. Vol. 28. Proceedings of Machine Learning Research. 2013, pp. 1310–1318. URL: <https://proceedings.mlr.press/v28/pascanu13.html>.
- [146] Nitish Srivastava et al. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. URL: <https://jmlr.org/papers/v15/srivastava14a.html>.

- [147] Paulius Micikevicius et al. “Mixed Precision Training”. In: *International Conference on Learning Representations*. 2018. URL: <https://openreview.net/forum?id=r1gs9JgRZ>.
- [148] Lutz Prechelt. “Automatic Early Stopping Using Cross Validation: Quantifying the Criteria”. In: *Neural Networks* 11.4 (1998), pp. 761–767. DOI: 10.1016/s0893-6080(98)00010-0.
- [149] Rob J. Hyndman and Anne B. Koehler. “Another Look at Measures of Forecast Accuracy”. In: *International Journal of Forecasting* 22.4 (2006), pp. 679–688. DOI: 10.1016/j.ijforecast.2006.03.001.
- [150] Jeffrey D. Scargle. “Studies in Astronomical Time Series Analysis. II. Statistical Aspects of Spectral Analysis of Unevenly Spaced Data”. In: *The Astrophysical Journal* 263 (1982), pp. 835–853. DOI: 10.1086/160554.
- [151] N. R. Lomb. “Least-Squares Frequency Analysis of Unequally Spaced Data”. In: *Astrophysics and Space Science* 39.2 (1976), pp. 447–462. DOI: 10.1007/bf00648343.
- [152] Jacob T. VanderPlas. “Understanding the Lomb–Scargle Periodogram”. In: *The Astrophysical Journal Supplement Series* 236.1 (2018), p. 16. DOI: 10.3847/1538-4365/aab766.
- [153] George E. P. Box et al. *Time Series Analysis: Forecasting and Control*. 5th ed. Hoboken, NJ: John Wiley & Sons, 2015. ISBN: 9781118675021. URL: <https://www.wiley.com/en-us/Time+Series+Analysis%3A+Forecasting+and+Control%2C+5th+Edition-p-9781118675021>.
- [154] Arrun Saunders, Graham G. Swinerd, and Hugh G. Lewis. “Deriving Accurate Satellite Ballistic Coefficients from Two-Line Element Data”. In: *Journal of Spacecraft and Rockets* 49.1 (2012), pp. 175–184. DOI: 10.2514/1.a32023.
- [155] Jizhang Sang, James C. Bennett, and Craig H. Smith. “Estimation of Ballistic Coefficients of Low Altitude Debris Objects from Historical Two Line Elements”. In: *Advances in Space Research* 52.1 (2013), pp. 117–124. DOI: 10.1016/j.asr.2013.03.010.
- [156] Zhejun Lu and Weidong Hu. “Estimation of Ballistic Coefficients of Space Debris Using the Ratios Between Different Objects”. In: *Chinese Journal of Aeronautics* 30.3 (2017), pp. 1204–1216. DOI: 10.1016/j.cja.2017.03.009.
- [157] Changyong He et al. “Impact of Thermospheric Mass Density on the Orbit Prediction of LEO Satellites”. In: *Space Weather* 18.1 (2020), e2019SW002336. DOI: 10.1029/2019sw002336.

- [158] G. E. Cook. “Satellite Drag Coefficients”. In: *Planetary and Space Science* 13.10 (1965), pp. 929–946. DOI: 10.1016/0032-0633(65)90150-9.
- [159] Piyush M. Mehta et al. “Satellite drag coefficient modeling for thermosphere science and mission operations”. In: *Advances in Space Research* 72.12 (2023), pp. 5443–5459. DOI: 10.1016/j.asr.2022.05.064.
- [160] Erin Kahr, Oliver Montenbruck, and Kyle O’Keefe. “Estimation and Analysis of Two-Line Elements for Small Satellites”. In: *Journal of Spacecraft and Rockets* 50.2 (2013), pp. 433–439. DOI: 10.2514/1.a32352.
- [161] Diogo M. Sanchez, Antonio F. B. A. Prado, and Tadashi Yokoyama. “On the effects of each term of the geopotential perturbation along the time I: Quasi-circular orbits”. In: *Advances in Space Research* (2014). DOI: 10.1016/j.asr.2014.06.003.
- [162] Lawrence F. Shampine and Mark W. Reichelt. “The MATLAB ODE Suite”. In: *SIAM Journal on Scientific Computing* 18.1 (1997), pp. 1–22. DOI: 10.1137/s1064827594276424.
- [163] Emmanuel Tsukerman and Ellen Veomett. “Brunn-Minkowski Theory and Cauchy’s Surface Area Formula”. In: *The American Mathematical Monthly* 124.10 (2017), pp. 922–929. DOI: 10.4169/amer.math.monthly.124.10.922.
- [164] Andrew Gelman et al. *Bayesian Data Analysis*. 3rd ed. Boca Raton, FL: Chapman and Hall/CRC, 2013. ISBN: 9781439840955. DOI: 10.1201/b16018.
- [165] Nicholas Metropolis et al. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092. DOI: 10.1063/1.1699114.
- [166] W. K. Hastings. “Monte Carlo Sampling Methods Using Markov Chains and Their Applications”. In: *Biometrika* 57.1 (1970), pp. 97–109. DOI: 10.1093/biomet/57.1.97.
- [167] Peter J. Rousseeuw and Christophe Croux. “Alternatives to the Median Absolute Deviation”. In: *Journal of the American Statistical Association* 88.424 (1993), pp. 1273–1283. DOI: 10.1080/01621459.1993.10476408.
- [168] Herbert Robbins and Sutton Monroe. “A Stochastic Approximation Method”. In: *The Annals of Mathematical Statistics* 22.3 (1951), pp. 400–407. DOI: 10.1214/aoms/1177729586.
- [169] Gareth O. Roberts and Jeffrey S. Rosenthal. “Optimal Scaling for Various Metropolis-Hastings Algorithms”. In: *Statistical Science* 16.4 (2001), pp. 351–367. DOI: 10.1214/ss/1015346320.

- [170] Andrew Gelman and Donald B. Rubin. “Inference from Iterative Simulation Using Multiple Sequences”. In: *Statistical Science* 7.4 (1992), pp. 457–472. doi: 10.1214/ss/1177011136.
- [171] Aki Vehtari et al. “Rank-Normalization, Folding, and Localization: An Improved R-hat for Assessing Convergence of MCMC”. In: *Bayesian Analysis* 16.2 (2021), pp. 667–718. doi: 10.1214/20-ba1221.
- [172] Rob J. Hyndman. “Computing and Graphing Highest Density Regions”. In: *The American Statistician* 50.2 (1996), pp. 120–126. doi: 10.1080/00031305.1996.10474359.
- [173] Cédric Villani. *Optimal Transport: Old and New*. Vol. 338. Grundlehren der Mathematischen Wissenschaften. Berlin, Heidelberg: Springer, 2009. ISBN: 9783540710493. doi: 10.1007/978-3-540-71050-9.
- [174] Victor M. Panaretos and Yoav Zemel. *An Invitation to Statistics in Wasserstein Space*. Springer, 2020. doi: 10.1007/978-3-030-38438-8.
- [175] Gabriel Peyré and Marco Cuturi. “Computational Optimal Transport”. In: *Foundations and Trends in Machine Learning* 11.5–6 (2019), pp. 355–607. doi: 10.1561/22000000073.
- [176] Andrew Gelman, Xiao-Li Meng, and Hal S. Stern. “Posterior Predictive Assessment of Model Fitness via Realized Discrepancies”. In: *Statistica Sinica* 6.4 (1996), pp. 733–807. URL: <https://www3.stat.sinica.edu.tw/statistica/j6n4/j6n41/j6n41.htm>.
- [177] D. G. King-Hele. *Satellite Orbits in an Atmosphere: Theory and Application*. Dordrecht: Springer Dordrecht, 1987. ISBN: 978-0-216-92252-5.
- [178] Carmen Pardini and Luciano Anselmo. “Orbital re-entries of human-made space objects: Drawbacks for the upper atmosphere and the safety of people”. In: *Journal of Space Safety Engineering* 12.2 (2025), pp. 274–283. doi: 10.1016/j.jsse.2025.04.009.
- [179] Richard R. Hofer and Alec D. Gallimore. “High-Specific Impulse Hall Thrusters, Part 2: Efficiency Analysis”. In: *Journal of Propulsion and Power* 22.4 (2006), pp. 732–740. doi: 10.2514/1.15954.
- [180] Vlad-George Tirila, Alain Demairé, and Charles N. Ryan. “Review of Alternative Propellants in Hall Thrusters”. In: *Acta Astronautica* 212 (2023), pp. 284–306. doi: 10.1016/j.actaastro.2023.07.047.
- [181] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. “The Frontier of Simulation-Based Inference”. In: *Proceedings of the National Academy of Sciences* 117.48 (2020), pp. 30055–30062. doi: 10.1073/pnas.1912789117.

- [182] David S. Greenberg, Marcel Nonnenmacher, and Jakob H. Macke. “Automatic Posterior Transformation for Likelihood-Free Inference”. In: *Proceedings of the 36th International Conference on Machine Learning*. Vol. 97. Proceedings of Machine Learning Research. 2019, pp. 2404–2414. URL: <https://proceedings.mlr.press/v97/greenberg19a.html>.
- [183] Álvaro Tejero-Cantero et al. “sbi: A Toolkit for Simulation-Based Inference”. In: *Journal of Open Source Software* 5.52 (2020), p. 2505. DOI: 10.21105/joss.02505.
- [184] Paul Fearnhead and Dennis Prangle. “Constructing Summary Statistics for Approximate Bayesian Computation: Semi-Automatic Approximate Bayesian Computation”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 74.3 (2012), pp. 419–474. DOI: 10.1111/j.1467-9868.2011.01010.x.
- [185] Kanti V. Mardia and Peter E. Jupp. *Directional Statistics*. Wiley Series in Probability and Statistics. Chichester: John Wiley & Sons, 2000. ISBN: 9780471953333. DOI: 10.1002/9780470316979.
- [186] Abraham Savitzky and Marcel J. E. Golay. “Smoothing and Differentiation of Data by Simplified Least Squares Procedures”. In: *Analytical Chemistry* 36.8 (1964), pp. 1627–1639. DOI: 10.1021/ac60214a047.
- [187] Xue Bai et al. “Mining Two-Line Element Data to Detect Orbital Maneuver for Satellite”. In: *IEEE Access* 7 (2019), pp. 537–549. DOI: 10.1109/access.2019.2940248.
- [188] Jesse A. Linnell and Alec D. Gallimore. “Krypton Performance Optimization in High-Voltage Hall Thrusters”. In: *Journal of Propulsion and Power* 22.4 (2006), pp. 921–929. DOI: 10.2514/1.18531.
- [189] G. S. Voronov. “A Practical Fit Formula for Ionization Rate Coefficients of Atoms and Ions by Electron Impact: $Z = 1-28$ ”. In: *Atomic Data and Nuclear Data Tables* 65.1 (1997), pp. 1–35. DOI: 10.1006/adnd.1997.0732.
- [190] S. D. Loch et al. “Electron-Impact Ionization of All Ionization Stages of Krypton”. In: *Physical Review A* 66.5 (2002), p. 052708. DOI: 10.1103/physreva.66.052708.
- [191] Byron D. Tapley, Bob E. Schutz, and George H. Born. *Statistical Orbit Determination*. Burlington, MA: Elsevier Academic Press, 2004. ISBN: 9780126836301.
- [192] John T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. 2nd ed. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2010. ISBN: 9780898716887. DOI: 10.1137/1.9780898718577.

- [193] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. 2nd ed. Springer Series in Operations Research and Financial Engineering. New York, NY: Springer, 2006. ISBN: 9780387303031. DOI: 10.1007/978-0-387-40065-5.
- [194] Hans Georg Bock and Karl-Josef Plitt. “A Multiple Shooting Algorithm for Direct Solution of Optimal Control Problems”. In: *Proceedings of the 9th IFAC World Congress*. Vol. 17. IFAC Proceedings Volumes 2. Budapest, Hungary, 1984, pp. 1603–1608. DOI: 10.1016/S1474-6670(17)61205-9.
- [195] Jon C. Helton et al. “Survey of Sampling-Based Methods for Uncertainty and Sensitivity Analysis”. In: *Reliability Engineering & System Safety* 91.10–11 (2006), pp. 1175–1209. DOI: 10.1016/j.ress.2005.11.017.
- [196] Andreas Raue et al. “Structural and Practical Identifiability Analysis of Partially Observed Dynamical Models by Exploiting the Profile Likelihood”. In: *Bioinformatics* 25.15 (2009), pp. 1923–1929. DOI: 10.1093/bioinformatics/btp358.
- [197] Ryan N. Gutenkunst et al. “Universally Sloppy Parameter Sensitivities in Systems Biology Models”. In: *PLoS Computational Biology* 3.10 (2007), e189. DOI: 10.1371/journal.pcbi.0030189.
- [198] George Papamakarios, Iain Murray, and Theo Pavlakou. “Masked Autoregressive Flow for Density Estimation”. In: *Advances in Neural Information Processing Systems 30*. 2017, pp. 2338–2347. URL: <https://papers.nips.cc/paper/6828-masked-autoregressive-flow-for-density-estimation>.

Appendix A

Code

The code, data-processing scripts, and supporting files used in this thesis are publicly available through the following GitHub repository:

[`https://github.com/Keeby-Astro/`](https://github.com/Keeby-Astro/)

`On-Analysis-and-Optimization-of-Low-Earth-Orbit-Satellite-Constellations`