

UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

STOCHASTIC METHODS FOR SLIP DETECTION IN A SHEARED  
GRANULAR SYSTEM

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

PHILIP BRETZ  
Norman, Oklahoma  
2023

STOCHASTIC METHODS FOR SLIP DETECTION IN A SHEARED  
GRANULAR SYSTEM

A DISSERTATION APPROVED FOR THE  
DEPARTMENT OF MATHEMATICS

BY THE COMMITTEE CONSISTING OF

Dr. Miroslav Kramar, Chair

Dr. Nikola Petrov

Dr. Alexander Grigo

Dr. Dimitrios Diochnos

© Copyright by PHILIP BRETZ 2023

All Rights Reserved.

## DEDICATION

I dedicate this dissertation to my wife Alissa, without whom I would have never completed it, and my daughters Lily and Edith, without whom I would have finished far sooner.

## **Acknowledgements**

I am deeply indebted to my professor and chair of my committee for his tireless support and feedback. I also would like to express my deepest appreciation to my defense committee for their knowledge and time. Additionally, this endeavor would not have been possible without the expertise of Dr. Lou Kondic.

I am also grateful to my cohort members for their moral support along this journey. Lastly, I would like to acknowledge my family, especially my wife, for their patience and belief in me. Their support kept my motivation high throughout all the stages of this process.

# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Granular media and the considered granular system</b>	<b>7</b>
2.1	Dynamics of sheared granular media . . . . .	8
2.2	Description of the sheared granular system . . . . .	11
2.3	Delineation of regimes in the considered system . . . . .	13
<b>3</b>	<b>Quantifying topological change between force networks</b>	<b>18</b>
3.1	Simplicial representation of the contact and force networks . . . . .	22
3.2	Homology of simplicial complexes . . . . .	26
3.3	Persistent homology of filtered simplicial complexes . . . . .	31
3.4	Distances between persistence diagram representations of successive force networks . . . . .	35
3.5	Additional measures of change between successive force networks . . .	41
<b>4</b>	<b>Bayesian framework for modeling considered measures</b>	<b>45</b>
4.1	Introduction to Bayesian statistics . . . . .	47
4.2	Stochastic state space models and the Kalman filter . . . . .	62
4.3	Specification of error structure . . . . .	76
4.3.1	Discount factor specification of model error covariance matrices	77
4.3.2	Conjugate-variance Kalman filter for observation error variance	78
4.4	Extended Kalman filter (EKF) . . . . .	82

<b>5</b>	<b>Model specification for considered measures</b>	<b>86</b>
5.1	State construction for a synthetic example . . . . .	89
5.2	Modeling higher frequency oscillations . . . . .	92
5.3	Model specification for considered measures . . . . .	94
5.4	Detections from dynamic linear models (DLMs) . . . . .	102
<b>6</b>	<b>Results</b>	<b>104</b>
6.1	Tri-partite classification of detections . . . . .	105
6.2	Detections for varying error thresholds . . . . .	108
6.3	Insight provided by different measures . . . . .	112
6.4	Predictive power of different measures . . . . .	115
<b>7</b>	<b>Conclusions and future work</b>	<b>119</b>
7.1	Future avenues of research . . . . .	120

## Abstract

We develop a flexible methodology for predicting slip events in a sheared granular system. The considered system consists of two-dimensional soft disks between two rigid horizontal walls, where the top wall is exposed to downward pressure and horizontal elastic shearing force, resulting in intermittent stick-slip regimes. The prediction methodology first uses topological data analysis to compute the persistent homology between successive force networks of the system and then quantifies the topological change by placing a metric between the respective persistence diagrams, resulting in a time series. Next, we construct a Bayesian stochastic state space model, which describes the behavior of the time series during the stick regime. We also create similar models for the stick regime behavior of the time series of more traditional measures on the granular system. A model identifies departure from the stick regime by detecting when the predictive error exceeds a specified threshold. The resulting detections demonstrate that this approach can detect the slip events in advance, with further investigation revealing a rough sequence of events. First, a local change appears in the force network and either dissipates or spreads globally. Next, the global change either triggers a slip event or a much smaller ‘micro-slip,’ depending on if its magnitude exceeds a critical threshold.



# Chapter 1

## Introduction

Granular materials are conglomerations of discrete macroscopic particles, and such materials are ubiquitous in the natural world. Consider the omnipresence of sand, gravel, and dirt; these are all granular materials. There is a deep connection between granular media and geological processes [16, 5, 13]. Granular materials also populate industrial settings, of which seeds, grains, and gravel are just a few examples. Since the storage, transport, and manipulation of such materials is a facet of daily life, much research is devoted to understanding their behavior, which can be quite complex.

We will focus on a particular class of granular media systems, those exposed to shearing force. Fig. 1.1(a) depicts a snapshot of the considered system. The granular system consists of two-dimensional particles in a rectangular domain. The bottom wall is static, while the left and right walls have periodic boundary conditions. Both downward pressure and elastic shearing force are applied to the top wall. The elastic force is induced by a spring attached to the top wall being pulled horizontally at a constant velocity.

Sufficiently chosen values of the system's parameters result in two basic regimes of behavior, classified by the movement of the top wall [2]. The first, the stick regime, is characterized by slow movement of the top wall. In this period of minimal movement, the downward pressure and resulting friction force the top wall to be essentially stuck, hence the terminology. The second type of behavior, the slip regime, is characterized by rapid movement of the top wall (relative to the stick regime). These intermittent

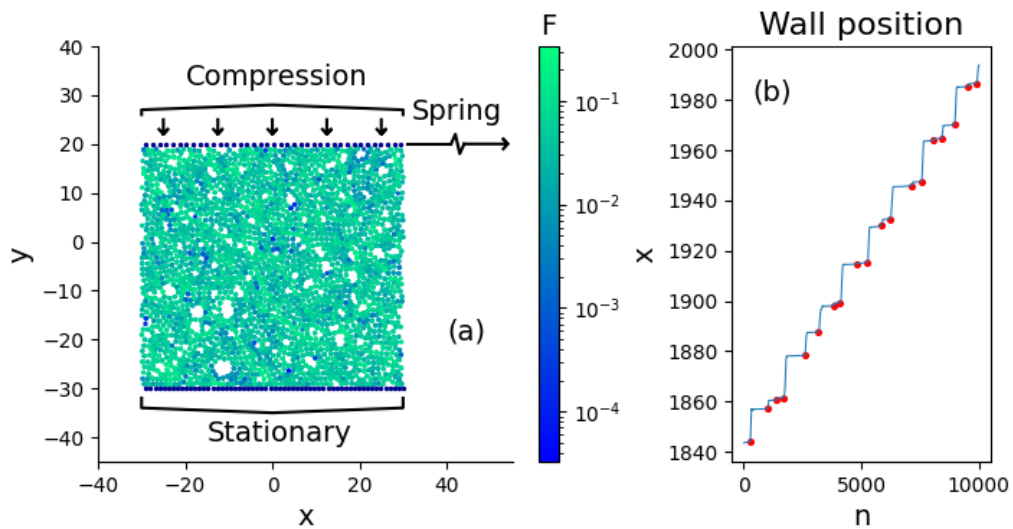


Figure 1.1: (a) DEM simulation of the granular system at a fixed time, with particles colored by force. The rigid row of particles forming the bottom wall are stationary, while elastic shearing force and normal compression force are applied to the rigid row forming the top wall. (b) Horizontal displacement of the top wall for the first 10,000 frames, where  $n$  is the number of frames after data collection initialization. Red points mark the start of slip regimes as determined by the offline method for regime change identification, which is detailed at the end of Ch. 2.

regimes are visible in the temporal evolution of the top wall position shown in Fig. 1.1(b).

At a basic level, the transition between these two regimes is caused by the force built up in the spring attached to the top wall [2, 29]. In the stick regime, the slow movement of the top wall causes the spring to grow taut, eventually resulting in regime change. These are the marked points in Fig. 1.1(b). In contrast, during the slip regime, the rapid movement of the top wall releases the built-up tension of the spring. After enough energy is spent, the wall becomes stuck again, transitioning into the stick regime once more.

In sheared granular systems, the prediction of regime change is of great interest. One component of this interest is the similarity between regime changes in sheared granular media systems and earthquakes/landslides in geophysical systems [6]. In particular, the stick-slip transition in sheared granular media bears a resemblance to the high-level tectonic behavior governing earthquakes [17, 13]. Tectonic plates have pressure forcing them together with the resulting friction temporarily locking them in place. During that period, shearing force builds, culminating in the earthquake event.

With geophysical systems, in some cases, earthquakes have notable precursors: observable phenomena that precede the event. Foreshocks, smaller events that precede a large-scale earthquake event are one example, though their predictive potential is unclear [3, 29]. Another example of precursor behavior is documented in [34], which analyzes a series of displacement measurements taken at the site of a landslide. The authors show that the cliff separates into distinct groupings of similarly moving parts of the cliff, or kinematic clusters, before collapse. The rise in data availability and precision has caused a recent expansion in the research of tools to properly analyze signs of imminent geophysical failure events [33].

Geophysical systems are highly complex, and the prediction of earthquakes is outside the scope of our research. However, geophysical precursors have an analog in sheared granular media systems; signs of regime change are known to appear before the slip transition [3]. In particular, for the considered system depicted in Fig. 1.1, we will show examples of observable internal changes that occur well before the wall begins to move.

To utilize the precursors in the considered system for the prediction of upcoming regime change, we apply a novel fusion of topological data analysis (TDA) and Bayesian statistical models. We focus on precursors seen in the force network, the dynamic network formed by particles in contact, and the forces between them. Fig. 1.2 depicts a snapshot of the force network of the considered system.

The structure of the force network is visibly quite complex. To overcome this complexity we use TDA to encode the important properties of the force network compactly [23]. In Fig. 1.2, we can see a structure of lighter-colored areas, similar to mountains rising above the darker-colored valleys. By considering the force network to be a function defined over the network whose vertices correspond to the particles and edges to the contacts between the particles, we can obtain a powerful description of the overall structure through persistent homology. Persistent homology succinctly describes the changes in the topology of the super-level sets of this function through a persistence diagram. Distance between the persistence diagrams can then be used to quantify the changes in the consecutive force networks.

We will show that the distance between the persistence diagrams of consecutive force networks indeed reflects the precursors to regime change. However, the time series formed by this sequence of persistence diagram distances exhibits complicated behavior that needs to be analyzed with statistical techniques. This is where we apply Bayesian methods. We first show that the typical behavior of the time series




Figure 1.2: Example of the force network of the considered. The force value between interacting particles is the magnitude of the normal force between them. Force chains, linked configurations of particles with high levels of force, are heterogeneously distributed and extend throughout the system.

during the stick regime can be modeled by a particular class of Bayesian models, a dynamic linear model (DLM). Then we will show that during the slip regime, the constructed model performs poorly, with large predictive error. We will further find that the predictive errors of this model can rise well before the wall begins to move

appreciably, giving this method its predictive potential.

Moreover, the Bayesian modeling approach is not restricted to topological measures but applies to any measure on the system whose signal during the stick regime can be modeled by a DLM. We apply the same approach to other more classical measures, such as the wall velocity and the percolation force. Through this approach, we will find that out of the considered measures, TDA gives the earliest warning of an upcoming slip.

This dissertation expands on our research published in the prestigious *Physical Review E* [7], providing more detail on our methodology and the underlying theory. The manuscript is organized as follows. In Chapter 2 we briefly introduce granular materials and the dynamics relevant to the stick-slip regimes. We then describe the simulated granular system. In Ch. 3 we introduce the topological measure used to quantify change between successive force networks, as well as more classical measures on the dynamics of the granular system. In Ch. 4 we describe the Bayesian modeling framework used to analyze the measures. In Ch. 5 we detail some types of behavior a DLM can describe and then specify the model for each of the considered measures. In Ch. 6 we analyze the events detected by our method before offering our conclusions and directions for future work in Ch. 7.

## Chapter 2

### Granular media and the considered granular system

A system of granular media is a large collection of discrete, macroscopic particles. Granular systems are typically modeled as large dynamical systems following Newtonian physics, see [15] for a review. At a high level, the state of the whole system consists of the individual states (position, momentum, etc.) of each of the particles. The dynamics of the system's state are modeled by a system of differential equations that encode the Newtonian evolution of particles, including the result of particle interaction. The solution of this system of differential equations is the trajectory of the system, and the state at each time.

Granular media is often contrasted with continuous media. The former is inherently discrete, while the latter consists of a homogeneous material such as a liquid, gas, etc. Continuous media is typically modeled by partial differential equations; the Navier-Stokes equations [1]. These two classes of models, a large system of ordinary differential equations compared to a partial differential equation, encode the fundamental difference between these types of media: granular media is discrete, while continuous media is, as the name suggests, continuous. This difference results in emergent behaviors in granular systems that do not appear in continuous media [16]. In the case of the experiment depicted in Fig. 1.1, the granular system produces a stick-slip dynamic in the top wall - a dynamic that would not be present if the rectangular box contained continuous media.

## 2.1 Dynamics of sheared granular media

There are two basic regimes of behavior in the sheared system of granular media that we consider: the stick and slip regimes, visible in Fig. 1.1(b) in Ch. 1. These regimes are not unique to our experiment, but have a rich history in sheared granular systems [2]. However, not all sheared granular systems exhibit stick-slip dynamics.

Recall the sheared system described in Ch. 1. The presence of the intermittent stick-slip regimes is primarily a result of the interplay between two parameters: the magnitude of the downward pressure  $N$  and the speed of the spring  $V$ , sometimes called the loading velocity. Fig. 2.1 describes the interplay between these parameters as a phase diagram. For sufficiently large pressure and small loading velocity, the sheared granular system displays stick-slip dynamics. However, if the pressure is too small or the spring moves too fast, the system never enters a stick regime [10, 2]. The behavior of the system under these conditions is referred to as continuous flow, to indicate that the top wall is moving continuously instead of in alternating regimes. To cause the considered system to display the desired stick-slip regimes instead of continuous flow, we performed a numerical investigation to choose appropriate values of these parameters.

So, with proper selection of the loading velocity and the downward pressure, a sheared granular system will result in stick-slip dynamics, as in Fig. 1.1(b). But what is the threshold of shearing force that marks the transition between regimes in a sheared granular system? Perhaps surprisingly, there is often no definitive threshold. The stick and slip regimes are defined through the movement of the top wall by convention. However, the actual causes of regime transitions extend well beyond just the top wall; regime change is a result of the interactions of particles throughout the entire system [10].



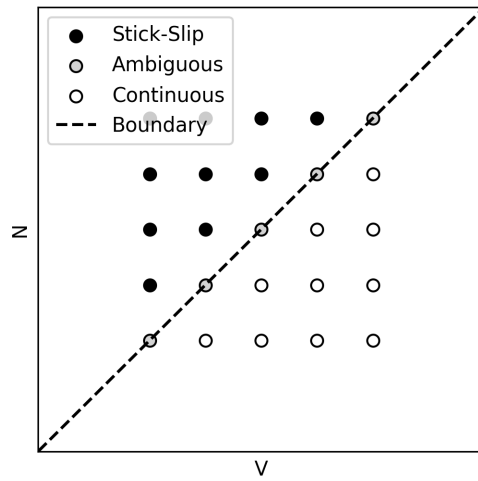


Figure 2.1: Diagrammatic depiction of dependence of stick-slip dynamics on magnitude of downward pressure  $N$  and loading velocity  $V$  in a sheared granular system. Actual values depend on the considered system. For previous experiments on this dependence, see [2].

Due to the complexity of granular systems, the exact causes of regime change remain somewhat mysterious [2]. Indeed, this is why our research is important; we wish to shed light on these critical dynamics. To discuss more of what is currently known about the mechanics of regime change, we first define a fundamental tool used to describe systems of granular media: the force network.

The *force network*, an example of which is depicted in Fig. 1.2, is a weighted graph. Each particle is assigned a node, and edges are placed between particles in contact, this unweighted graph is called the *contact network*. The weight on an edge is a measure of force between the two particles; in our work, we use the magnitude of the normal force.

A common feature of force networks in granular systems is a heterogeneous, or anisotropic distribution of forces [16]. In this context, a heterogeneous distribution refers to the appearance of connected groups of edges with high forces that extend

throughout the material interspersed among edges with lower forces. An example of these groupings, referred to as force chains, is visible in Fig. 1.2.

The force network is deeply connected to the mechanical properties of the material [23]. A quintessential example of this grain in a silo, which has prompted numerous experiments [16]. With continuous media, the pressure at the base of the silo increases with the height of the material. However, this is not the case with a granular material like grains. Instead, chains of particles in contact with each other, and the walls, can lock in place and distribute the pressure, causing the silo walls to support additional weight after some threshold height.

In a dynamic setting, granular materials tend to form temporarily stable micro-structures, with particles that are ‘jammed’ together. Jammed groups, reflected as force chains in the force network, exist as solids, with forces transferring from particle to particle through the self-organized grouping [23, 9].

In sheared granular systems, the stick regime is characterized by such jamming, while the slip regime is accompanied by large-scale rearrangements of the force chains [10]. In our system, we observe that at the end of the slip regime, the decreased shearing force allows for more stable micro-structures to form until eventually the force chains extending throughout the system are stable enough to ‘catch’ the top wall. During the stick regime, the force chains remain relatively stable. Then, as the shearing force builds, it seems that an internal change occurs in the system, triggering the onset of the slip regime. So, regime change does not just depend on the sheared wall, but on complex internal dynamics that are still not fully understood.

Full access to the exact force network to study such micro-mechanics in detail is a strong motivation for the numerical simulation of granular systems. Granular media is frequently studied computationally through discrete element method (DEM) simulations [15]. Such simulations numerically solve the system of ordinary differential

equations to evolve the conglomeration of granular particles forward in time. Since the number of differential equations to solve increases directly with the quantity of particles, the size of a simulated system is limited by computational resources. Despite recent advances in computing, the number of particles remains a practical limitation [15]. Indeed, the considered system from Fig. 1.1 is a DEM simulation that consists of a few thousand particles, a relatively small size. We give the details of the simulated system in the following section.

## 2.2 Description of the sheared granular system

To produce intermittent stick-slip regimes, we use the previously mentioned elastically sheared system of granular media. The system consists of  $N = 2500$  circular particles densely placed in a rectangular domain with periodic boundary conditions on the right and left. The bottom is fixed, while the top wall is exposed to normal compressive force and elastic shearing force, resulting in the desired stick-slip dynamics. Gravitational force is not included. Figure 1.1(a) shows a snapshot of this system, which we further describe here.

The state of the system consists of the position, linear momentum, and angular momentum of the particles, which are modeled by 2-dimensional soft frictional disks. For particles that are not in contact with others, the Newtonian evolution is linear. To model particle interactions, we use the linear spring-dashpot model [21]. This model decomposes the force between two contacting particles into normal and tangential components, each modeled by a spring.

We use a bi-disperse sample of particles, with a 25% / 75% split between large and small particles, where the small particles have diameter  $d = 1.27$  cm and mass  $m = 1.32$  g and the large particles are 25% larger in diameter. The top and bottom

walls consist of equally spaced (small) particles in a rigid row. The spacing of the particles forming the bottom row is exactly  $d$ . The row of top particles can move, subject to an elastic shearing force, a normal compressive force, and the contacting particles below it. The top row also has wider spacing, which allows the internal particles touching the top wall to more easily ‘catch’ it when transitioning to the stick regime. Additionally, we also use a larger static friction coefficient for particle-wall interaction than for particle-particle interaction.

The small particle diameter  $d$  and mass  $m$  provide the length scale and mass scale for the linear spring-dashpot model, respectively. The time scale is the binary collision time  $\tau_c = 1.25 \times 10^{-3}$ . These choices of scale are motivated by previous experiments with photoelastic particles [11]. In particular, the Young modulus of elasticity that corresponds to  $\tau_c$  is  $Y \approx 0.7$  MPa.

The linear spring-dashpot model with spring constants for the normal and tangential components that are dependent on  $\tau_c$ . The normal spring constant is  $k_n = m\pi^2/2\tau_c^2 \approx 4.17$  N/m. Static friction is modeled by the Cundall-Strack method [15], where a damped tangential spring with spring constant  $k_t = 6k_n/7$  appears between particles upon contact. For particle-particle interaction, the static friction coefficient is  $\mu_p = 0.7$ , while as mentioned previously, particle-wall interaction uses a larger coefficient of  $\mu_w = 2$ .

The simulation results in slip-stick dynamics when a sufficiently large pressure  $p$  is applied to the top wall in the  $-y$  direction and a spring attached to the top wall moves at a sufficiently small speed  $v_s$  in the  $+x$  direction. The spring coefficient of the spring inducing the elastic shearing force is  $k_s = k_n/400$  (with a restitution coefficient of 0.5) and is two orders of magnitude smaller than the spring coefficients that govern particle interactions. The chosen pressure  $p = 0.02$  and spring speed  $v_s = 1.5 \times 10^{-3}$  were selected via numerical investigation as values that induced slip-stick dynamics,

as in [2, 10].

The above model for the dynamical system is numerically simulated as a DEM simulation. The simulation progresses by integrating the Newtonian equations of motion: basic linear evolution for non-interacting particles and the linear spring-dashpot model for particles in contact. This is done with a fourth-order predictor-corrector method using a time step of  $dt = 0.02$  s. We store the results of the simulation at times spaced by  $\delta t = 10dt$ . The simulation is initialized by first applying pressure  $p$  to the top wall until the particles have ‘settled.’ This is determined by a sufficiently small ratio of kinetic energy to potential energy. Then, the spring starts to move in the  $+x$  direction at speed  $v_s$ , causing the system to enter an initial transient regime, characterized by erratic stick-slip dynamics. After an initial period, the transient regime transitions to the intermittent stick-slip regimes shown in Fig. 1.1(b), at which point we begin collecting data.

**Remark 2.1.** *In this manuscript, all time-dependent figures are shown in units of  $\delta t$ , using the index  $n \in \mathbb{N}$ , which denotes the number of  $\delta t$  units that have passed since data collection was initialized. The corresponding time is  $t_n = n\delta t + t_0$  where  $t_0$  is the first data collection time. We will refer to  $n$  as the ‘number of frames,’ while any reference to ‘time’ uses the variable  $t$  [seconds].*

### 2.3 Delineation of regimes in the considered system

By design, the considered system detailed in the previous section exhibits intermittent stick-slip regimes. This can be seen in the sample of top wall position  $x$ , and velocity,  $v_x$ , shown in Fig. 2.2. The top wall, which is the focus of the downward pressure and elastic shearing force, undergoes periods of small  $v_x$ , the stick regime, punctuated by shorter periods of larger  $v_x$ , the slip events.

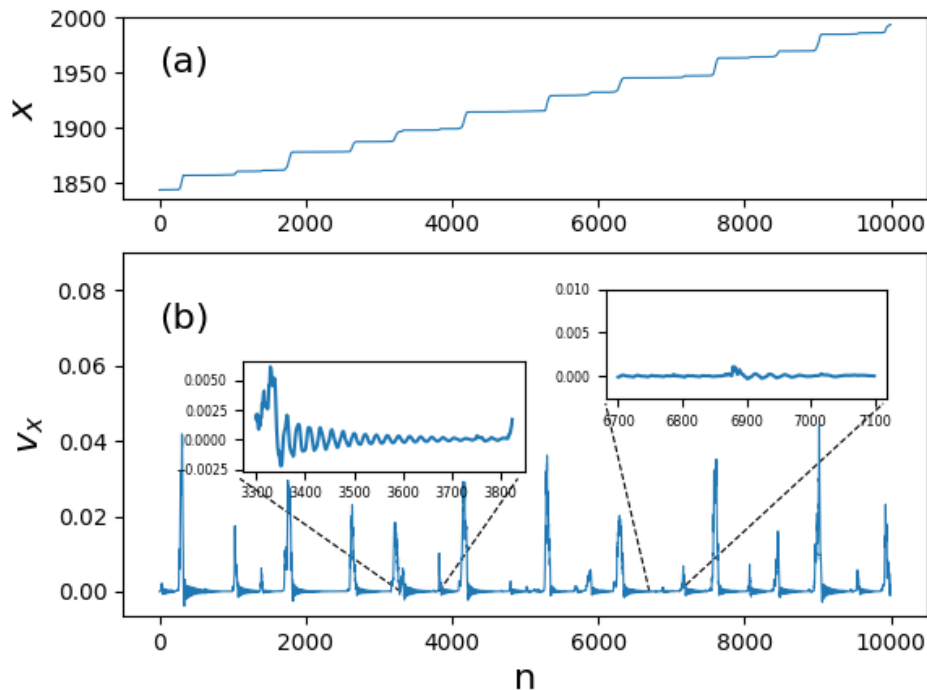


Figure 2.2: (a) The top wall horizontal position  $x$  and (b) the horizontal top wall velocity  $v_x$  for the first 10,000 frames of the considered system. The first inset depicts a sample of  $v_x$  during the stick regime; the oscillations are typical. The second inset shows another sample of the stick regime with a micro-slip.

To begin our investigation into the intermittent stick-slip dynamics, we must first delineate the stick and slip regimes. The two regimes are defined by the speed of the top wall: low  $v_x$  corresponds to the stick regime while large  $v_x$  identifies the slip, as seen in Fig. 2.2(b). Under this principle, a basic approach for identifying the slip regime would consist of setting a threshold and if  $v_x$  surpasses it, identifying the beginning of the slip event. Then, once the velocity falls below that threshold, the new stick regime begins.

However, the problem is choosing this threshold. As shown in the inset of Fig. 2.2(b), the wall often oscillates during the stick regime. The amplitude of these oscillations can remain relatively high even at the end of stick regimes. Consequently,

to apply this basic approach, we must choose a threshold that is larger than this amplitude. However, if we do so, the point where the wall velocity rises above this relatively large threshold occurs well after the wall has begun to move. So, to more accurately identify the point at which  $v_x$  begins to increase, we must use a second, smaller threshold.

A further complication in the choice of large threshold is the existence of so-called ‘micro-slips,’ events in which the wall moves, but on a smaller scale than large slip events. An example of a micro-slip in the considered system can be seen in the second inset in Fig. 2.2(b). This nomenclature is motivated by the literature [28, 9, 18, 19], though the distinction between slip events and micro-slips is often subjective, based on the focus of the researcher. In our case, the demarcation is based on a critical velocity threshold. Events in which the wall movement exceeds this threshold are defined as slip events and are the focus of our predictive efforts. We will later show that our choice of critical threshold appears to separate two distinct classes of events inherent in the system.

With the oscillations and micro-slips in mind, we provide an offline approach for identifying slip events using two thresholds. ‘Offline’ indicates that this method uses data after the slip event has begun to identify it, making such a method not useful for prediction. The algorithm first identifies a slip event when the wall velocity exceeds a critical threshold  $v_L = 2 \times 10^{-3}$ . This is the threshold that distinguishes slip events from micro-slips. It is also sufficiently larger than the oscillations and noise in the stick regime. For the identified slip event, the method uses a second, smaller threshold  $v_S = 2 \times 10^{-4}$ , to determine the beginning of the slip event. We stress that large fluctuations of the velocity make it impossible to use a single threshold to precisely identify when slip events start.

In more detail, the algorithm for offline delineation of the stick-slip regimes is as

follows. First, if the wall velocity exceeds the larger threshold,  $v_L$ , then a slip event is detected. Then, the algorithm moves back in time until the value of  $v_x$  drops below the smaller threshold  $v_S$ . The detected beginning of the slip then occurs at the time when the wall velocity exceeds  $v_S$ . The end of the slip is identified as the time at which  $v_x$  as well as its average over the preceding 50 frames (roughly the period of typical oscillations in the stick regime) is smaller than  $v_L$ . For more details on this method, see [22], where it was introduced.

In this method, the choice of  $v_L$  is critical, as it is the threshold separating slips from micro-slips, oscillations, and noise. Our choice of  $v_L = 2 \times 10^{-3}$  is not arbitrary but instead marks an apparent bipartite classification between slips and micro-slips in the considered system. In a similar experiment of a granular system under elastic shearing force, the authors also found a bipartite classification [10]. So, it is not surprising that we observe evidence of this classification in the frequency plot shown in Figure 2.3. This panel shows (on the logarithmic scale) the number of detections made by the offline method for different values of  $v_L$ . The marked change in slope indicates a bipartite classification: detections made at that threshold form one class while the additional detections made at lower thresholds form the other class.

Due to the oscillations in the stick regime, this offline approach was necessary. If we wish to identify an upcoming regime change as it occurs (i.e., an online approach), we will need to develop a more sophisticated approach that accounts for them. The online method we developed is our answer to this. Moreover, the flexibility of the resulting framework allows us to examine not only the wall velocity  $v_x$  but also other measures on the system. Since the transition to the slip regime is deeply connected to the complex internal structure of the granular system, we hope that appropriate measures of internal change may allow for the prediction of an upcoming slip. It is with this motivation that we introduce the topological measures in the following






Figure 2.3: The count of offline detections for varying values of  $v_L$ . The dashed line marks the selected  $v_L = 2 \times 10^{-3}$ , roughly located at the changepoint between differing rates of decrease. The different slopes (on the logarithmic scale) are indicative of a bipartite classification: events detected at the selected  $v_L$  form one class (the slip events) while the additional detections made at smaller  $v_L$  form another class (the micro-slips).

chapter.

## Chapter 3

### Quantifying topological change between force networks

In the previous chapter, we defined the *contact network* of a granular system. We begin by making the definition more precise. At time  $t$ , the contact network  $\text{CN}_t$  is a graph where each vertex corresponds to a particle in the system, and edges are placed between particles in contact. Each particle has a position, a point in  $\mathbb{R}^2$  located at its center<sup>1</sup>. So,  $\text{CN}_t$  has a natural embedding in  $\mathbb{R}^2$ : each vertex is mapped to the position of its particle, while each edge is mapped to the line segment connecting the positions of the particles it connects. Fig. 3.1(a) shows a simple contact network, which we will use for illustration throughout this chapter.

In particular, since the contact network  $\text{CN}_t$  is a collection of vertices and edges where the embedded edges are straight, non-intersecting line segments, it is a *simplicial complex*. Fig. 3.1 shows that this is the case for a simple contact network. For more details on the precise definition of a simplicial complex and the proof that  $\text{CN}_t$  is always a simplicial complex, see Sec. 3.1.

Since  $\text{CN}_t$  has a simplicial complex representation, its topological structure can be captured by *simplicial homology*, which describes the number of connected components and distinct loops in the contact network. Moreover, simplicial homology can be computed algorithmically [35].

As introduced in Ch. 2, we further associate a value to the edges in the contact

---

<sup>1</sup>All the theory in this chapter applies equally to a three-dimensional granular system, where the position of each particle is a point in  $\mathbb{R}^3$

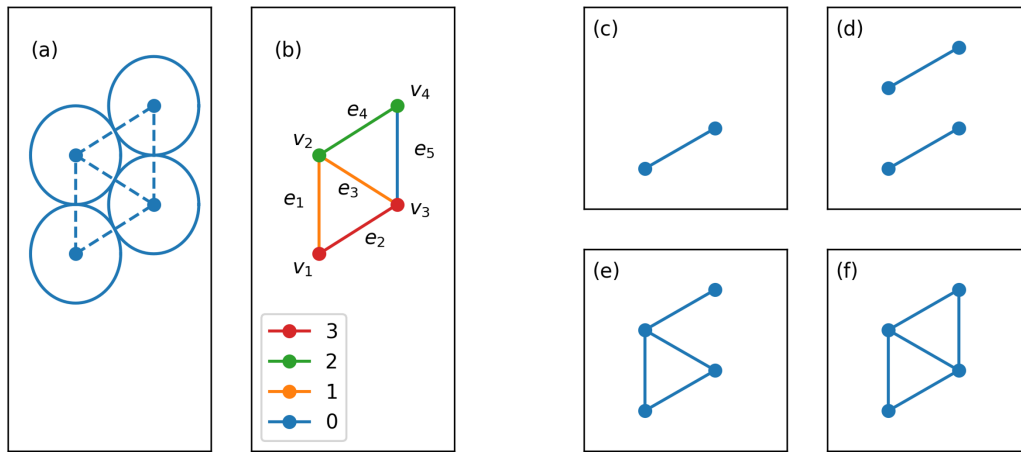


Figure 3.1: Example of particles (a) with a force network (b): an embedded contact network  $\text{CN}_t$  of four vertices and five edges, with a force function  $f : \text{CN}_t \rightarrow \mathbb{R}$  that takes values 0, 1, 2, 3. The super-level sets for the decreasing sequence of thresholds 3, 2, 1, 0, depicted in (c)-(f), are a nested sequence of sub-complexes, a filtration. The persistence diagrams at the zero and one-level are shown in (g)-(h).

network that quantifies the force between each edge's two particles, the *force network*. The force network is a real-valued function on the contact network  $f_t : \text{CN}_t \rightarrow \mathbb{R}$ . Recall in Ch. 2 that we define  $f_t$  on an edge in  $\text{CN}_t$  to be the magnitude of the

normal force between the edge's particles at time  $t$ . The force network is extended to each vertex  $v$  by defining  $f_t(v)$  to be the maximum of  $f_t$  over the edges with  $v$  as an endpoint. Fig. 3.1(b) depicts an example force network over the contact network in panel (a).

This extension to the vertices guarantees that each super-level set of the force network is also a simplicial complex. That is, the super-level sets are *sub-complexes* of the contact network. Moreover, as depicted in Fig. 3.1(c)-(f), the super-level sets for a decreasing sequence of force thresholds result in a nested sequence of sub-complexes. This sequence of sub-complexes is called the *filtration* of the force network.

As each super-level set is a simplicial complex, we can associate a simplicial homology group with each force threshold. For example, in Fig. 3.1(d) the super-level set has two connected components and no loops while in (e) the super-level set has one connected component and one loop. To understand the topological structure of the force network, we study how homology changes with the force threshold. This is quantified by *persistent homology* which, like simplicial homology, can be computed algorithmically [35].

Persistent homology describes how the homology of the super-level sets of a force network changes with a threshold value  $r$  decreasing from  $\infty$  to  $-\infty$ . Note that due to the finite number of edges, the force function takes a finite set of values that can be listed in descending order, e.g.,  $r_1 > r_2 > \dots > r_m$ . At  $r = \infty$ , the super-level set is empty, so there are no connected components or loops. This continues until  $r = r_1$ , the maximal value in the force network, upon which a super-level set with connected component(s) and possibly loops appears (e.g., at  $r = 3$  in Fig. 3.1).

As the force threshold  $r$  decreases from  $r_1$  to the next discrete force value  $r_2$ , the super-level set is unchanged until the threshold reaches  $r_2$ . In this new super-level set, new connected components and loops may appear, but connected components

may also merge (as in Fig. 3.1 at  $r = 1$ ). As  $r$  continues to decrease, loops appear<sup>2</sup> and connected components appear/merge whenever  $r$  reaches the next discrete force value until the threshold reaches the lowest force value. At this point, the super-level set is the entire contact network, and this remains the case as the threshold decreases to  $-\infty$ .

The appearance and disappearance of topological features (connected components and loops, captured by the homology groups) as the force value decreases can be encoded in a collection of pairs  $(b, d)$ . Each topological feature is assigned a pair  $(b, d)$  where  $b$ , called the ‘birth’, is the force value at which the feature appeared and  $d \in [-\infty, b)$ , called the ‘death’, is the force value at which the feature disappeared (e.g., when connected components merge, or when a loop is filled). A death of  $d = -\infty$  indicates a feature present at the lowest force value, i.e., a feature in the contact network, and the associated point is called a ‘point at infinity.’

The collection of these pairs for the zero-level homology is called the zero-level *persistence diagram*, denoted  $\text{PD}_0$ . Similarly, the collection of pairs for the one-level homology is the one-level persistence diagram,  $\text{PD}_1$ . As a convention, persistence diagrams also include a countably infinite set of copies of birth-death pairs of the form  $(x, x)$  for every  $x \in \mathbb{R}$ , referred to as the ‘diagonal’ of the diagram. A persistence diagram can be visually represented by plotting the birth-death pairs (and diagonal) in the plane, with points at  $-\infty$  placed just below the lowest force level.

Fig. 3.1(g) and (h) depict the zero and one-level persistence diagrams of the example force network. At the zero level, a connected component appears at threshold  $r = 3$ , a second at  $r = 2$ , the two merge at  $r = 1$ , and that single connected component is present at the lowest level. Hence, the diagram  $\text{PD}_0$  in Fig. 3.1(g) has two points:

---

<sup>2</sup>If higher dimensional simplices are allowed, loops may also disappear by being ‘filled in’ by 2-simplices, and in general, an element in the  $k$ -th homology group can be filled in by  $k + 1$ -simplices [23].

$(2, 1)$  represents the component that merges, and  $(3, -\infty)$  represents the component that is still present at the lowest level. Similarly, the diagram  $\text{PD}_1$  has two points,  $(1, -\infty)$  and  $(0, -\infty)$ , representing the loops that appear at thresholds 1 and 0, respectively, and remain at the lowest threshold.

Recall from Ch. 2 that we will be examining the force network of a sheared granular system designed to exhibit intermittent stick-slip regimes. Based on similar experiments [2, 10], we expect the force network of the sheared granular system to exhibit signs of regime change before the onset of a slip. To identify these precursory changes, we first quantify the complex structure of the force network at fixed times by computing its zero and one-level persistent diagrams. Next, we quantify changes in the persistence diagrams of the force network over time, providing us with a measure of topological change. To do this, we define a distance between two persistence diagrams in Sec. 3.4. This allows us to quantify topological changes between successive force networks.

At the end of this section, we will introduce more conventional measures for quantifying changes between successive force networks. We will show in Ch. 6 that the measure based on persistent homology allows us to detect an upcoming slip event earlier.

### 3.1 Simplicial representation of the contact and force networks

We began this chapter by stating that the contact network is a simplicial complex, as are the super-level sets of the force network. To make this claim precise, we first properly define *simplices* and *simplicial complexes*.

**Definition 3.1.** *Let  $\{v_0, \dots, v_k\}$  be a collection of affinely independent points in  $\mathbb{R}^n$  and  $\sigma \subset \mathbb{R}^n$  be their convex hull. We then say that  $\sigma$  is a **simplex** of dimension  $k$ .*

If  $\tau$  is the convex hull of a subset of  $\{v_0, \dots, v_k\}$ , we say that  $\tau$  is a **face** of  $\sigma$ . We refer to  $\{v_0, \dots, v_k\}$  as the **vertices** of  $\sigma$  and use the shorthand  $[v_0, \dots, v_k]$  to denote  $\sigma$ .

Lower-dimensional simplices are familiar objects: points, line segments, filled triangles, and tetrahedrons for 0, 1, 2, and 3-dimensional simplices, respectively. By construction, the contact network is indeed a collection of simplices, where the particles are 0-simplices and the edges, line segments connecting particles, are 1-simplices. Later, we will show that the contact network is a simplicial complex, as defined below.

**Definition 3.2.** Let  $X$  be a collection of simplices in  $\mathbb{R}^N$ . We say that  $X$  is a **simplicial complex** if it satisfies the following conditions:

1. For each  $\sigma \in X$ , if  $\tau$  is a face of  $\sigma$  then  $\tau \in X$
2. For each  $\sigma_1, \sigma_2 \in X$ , if  $\sigma_1 \cap \sigma_2 \neq \emptyset$  then  $\sigma_1 \cap \sigma_2 \in X$

We then call the subset of  $\mathbb{R}^N$  consisting of the union of the simplices of  $X$  the **carrier** of  $X$ , and denote it by  $|X|$ . A simplicial complex  $Y$  is a **sub-complex** of  $X$  if  $Y \subseteq X$ .

Observe how the contact network and super-level sets depicted in Fig. 3.1 immediately fit this definition: the vertices on the endpoints of each edge are always in contact network/super-level sets (condition 1) and edges only meet at vertices (condition 2). For the general case, we give the following proposition and its proof.

**Proposition 3.3.** Let  $CN$  be the contact network of a granular system and  $f : CN \rightarrow \mathbb{R}$  be a force function defined on the edges and each vertex  $v$  by defining  $f(v)$  to be the maximum of  $f$  over the edges with  $v$  as an endpoint. Then  $CN$  is a simplicial

complex and each super-level set at a threshold  $r \in \mathbb{R}$ , denoted by

$$X_r = \{\sigma \in CN : f(\sigma) \geq r\}$$

is a sub-complex of  $CN^3$ . Further,  $r_1 \geq r_2$  implies  $X_{r_1} \subseteq X_{r_2}$ .

*Proof.* As already stated,  $CN$  is a collection of 0 and 1-simplices (the vertices and edges). Property 1 in Def. 3.2 is quickly satisfied. For any edge  $e \in CN$ , its (proper) faces are the vertices it connects, which are in  $CN$ .

Property 2 is satisfied due to the physical nature of contacting particles. As shown in Fig. 3.2, two particles in contact meet along a line segment of positive length. We impose the physical restriction that particles cannot overlap, e.g., each point on the interior of a particle cannot also belong to another. Observe in Fig. 3.2 how a  $\delta/2$ -neighborhood of the line segment connecting the centers of the particles (i.e., the edge in the contact network) belongs to either the first particle, the second, or the boundary at which they meet.

Suppose, for contradiction, that two edges in the contact network, connecting intersect at a non-vertex point. Let  $\delta_1$  be the length of the line segment along which the contacting particles of the first edge meet, and let  $\delta_2$  be the same for the second edge. For  $\delta = \min(\delta_1, \delta_2)$ , consider a point in the  $\delta/2$  neighborhood of the non-vertex intersection that is not on the boundary of any of the four particles. This point must simultaneously be in the interior of one of the first edge's particles and one of the second edge's particles, which contradicts the non-overlapping restriction. So, two edges cannot intersect at a non-vertex. Since properties 1 and 2 of Def. 3.2 are satisfied,  $CN$  is a simplicial complex.

---

<sup>3</sup>The empty set is considered a simplicial complex, so the claim that  $X_r$  is a sub-complex is trivially true for  $r > \max f$ .



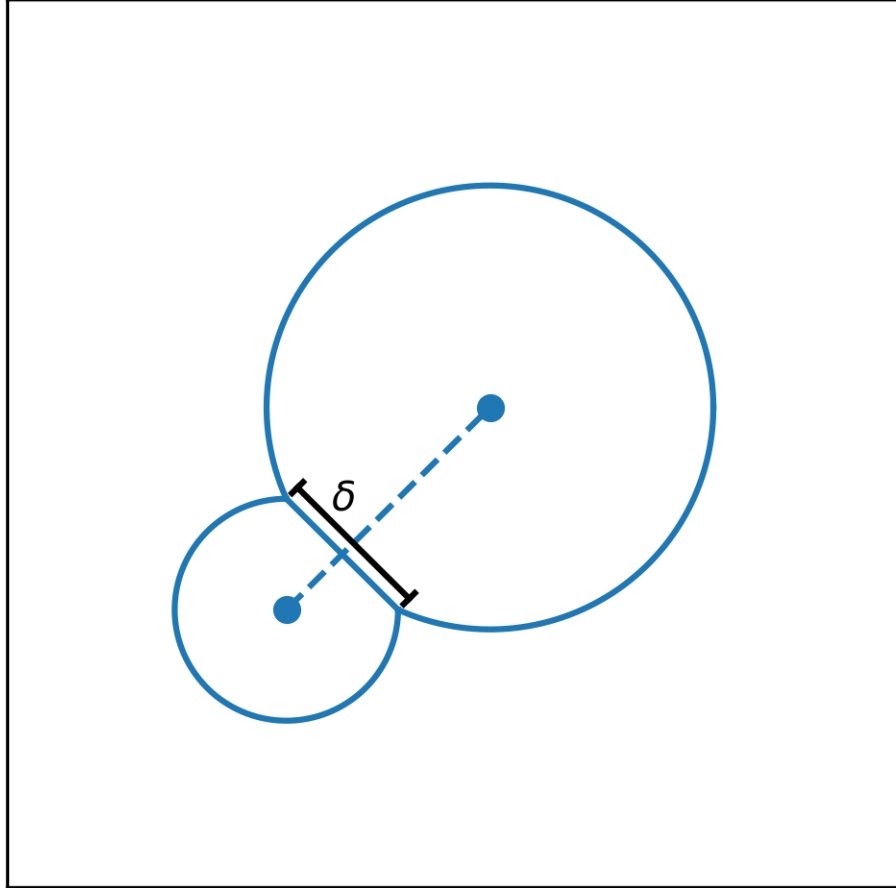


Figure 3.2: Diagram of two particles in contact meeting along a line segment of length  $\delta$ .

Next, let  $r \in \mathbb{R}$ . The super-level set  $X_r$  is a subset of CN and is therefore a collection of simplices where property 2 in Def 3.2 is satisfied. Property 1 is satisfied by defining  $f$  on each vertex as the maximum of  $f$  on the vertex's edges. By this definition, for any edge  $e \in \text{CN}$  and one of its vertices  $v$ ,  $f(v) \geq f(e)$ . Consequently, if  $e \in X_r$  then  $v \in X_r$ , since  $f(v) \geq f(e) \geq r$ . So, each  $X_r$  is a sub-complex of CN.

Lastly, suppose  $r_1 \geq r_2$ . Let  $\sigma \in X_{r_1}$ ; then  $f(\sigma) \geq r_1$ . But then  $f(\sigma) \geq r_2$ , so  $\sigma \in X_{r_2}$ . As a result, a decreasing sequence of thresholds corresponds to a nested sequence of sub-complexes.  $\square$

### 3.2 Homology of simplicial complexes

A simplicial complex has a definitive topological structure. In the case of the contact network and super-level sets of the force network, the topology is characterized by two quantities: the number of connected components and the number of loops. In this section, we will show how a purely algebraic characterization of simplicial complexes allows us to extract those quantities through the *simplicial homology* of a simplicial complex.

We begin by defining the  $k$ -th *chain group* of a simplicial complex, the abstract algebraic characterization of the simplices of dimension  $k$ .

**Definition 3.4.** *Let  $X$  be a simplicial complex. The  $k$ -th **chain group** over a ring  $R$ , denoted by  $C_k$ , is the collection of all formal sums of the form*

$$c = \sum_i a_i \sigma_i$$

where each  $\sigma_i$  is a  $k$ -simplex of  $X$  and  $a_i \in R$ . Each  $c \in C_k$  is called a **chain**. We will only consider the case where  $R = \mathbb{Z}_2$  so that  $C_k$  is a vector space for which the collection of  $k$ -simplices forms a basis.

To extend the definition of the chain groups of  $X$  to all  $k \in \mathbb{Z}$ , we use the convention that the  $k$ -th chain group is considered trivial if there are no  $k$ -simplices, which is trivially the case for  $k < 0$ .

With  $\mathbb{Z}_2$ -valued coefficients a  $k$ -chain  $c$  has a straightforward interpretation:  $c$

represents a collection of  $k$ -simplices where a coefficient of one for a simplex indicates the simplex is present in the collection and a coefficient of zero indicates the simplex is not.

In addition to the sequence of chain groups  $\{\dots, C_{k-1}, C_k, C_{k+1}, \dots\}$ , the purely algebraic formulation of a simplicial complex includes a sequence of homomorphisms, called the *boundary maps*, that keep track of the faces of the simplices. The  $k$ -th boundary map is simply a linear operation from  $C_k$  to  $C_{k-1}$  that maps each  $k$ -simplex to its  $k - 1$ -dimensional faces, or more precisely, the chain representing the collection of its faces.

**Definition 3.5.** *The  $k$ th boundary operator for chain groups over  $\mathbb{Z}_2$  is the linear operator  $\partial_k : C_k \rightarrow C_{k-1}$  defined on a  $k$ -simplex  $\sigma = [v_0, \dots, v_k]$  by*

$$\partial_k \sigma = \sum_i [v_0, \dots, \hat{v}_i, \dots, v_k]$$

where  $\hat{v}_i$  indicates the removal of the vertex  $v_i$ . The boundary operator extends linearly to general  $k$ -chains. For the definition of the boundary operator on chain groups over a general ring  $R$ , see [35].

The boundary maps have the property that their composition is trivial, which is made precise in the following proposition and its proof.

**Proposition 3.6.** *Let  $\{\dots, C_{k-1}, C_k, C_{k+1}, \dots\}$  be the chain groups (over  $\mathbb{Z}_2$ ) of a simplicial complex. Then the composition  $\partial_k \partial_{k+1} : C_{k+1} \rightarrow C_{k-1}$  is trivial:*

$$\partial_k \partial_{k+1} = 0$$

*Proof.* To see why the composition of boundary maps is trivial, consider a  $k + 1$ -simplex  $\sigma = [v_0, \dots, v_{k+1}]$ . If  $k < 1$ ,  $\partial_k \partial_{k+1} \sigma = 0$  because  $\partial_0$  is trivial. So, suppose

$k \geq 1$ . From Def. 3.5, the  $k$ -simplex  $\partial_{k+1}\sigma$  is the formal sum of the  $k + 2$  simplices made by removing each of the vertices  $v_0, \dots, v_{k+1}$ . Note that each particular vertex,  $v_j$ , is present in exactly  $k + 1$  terms in the sum. Next, when examining  $\partial_k\partial_{k+1}\sigma$ , note that by the linearity of the boundary operators,

$$\partial_k \left( \sum_i [v_0, \dots, \hat{v}_i, \dots, v_{k+1}] \right) = \sum_i \partial_k [v_0, \dots, \hat{v}_i, \dots, v_{k+1}]$$

Each term in the sum on the right,  $\partial_k [v_0, \dots, \hat{v}_i, \dots, v_{k+1}]$ , is a  $k - 1$  simplex that is the formal sum of  $k + 1$  simplices made by removing each of the vertices  $v_0, \dots, v_{k+1}$ , except for  $v_i$ , which was already removed. Further, any vertex  $v_j \neq v_i$  is present in exactly  $k$  terms of this sum.

So,  $\partial_k\partial_{k+1}\sigma$  is a double formal sum of  $k - 1$ -simplices composed of all possible permutations of removing exactly two vertices. More importantly, each vertex  $v_j$  is present in exactly  $k(k + 1)$  terms of the double sum. Since the coefficients of the chain groups are taken over  $\mathbb{Z}_2$  and  $k(k + 1)$  is an even number, the coefficient of each vertex  $v_j$  is zero, resulting in  $\partial_k\partial_{k+1}\sigma = 0$ . By extending linearly to general  $k + 1$ -chains, we can conclude that the composition of boundary maps,  $\partial_k\partial_{k+1}$  is trivial<sup>4</sup>.  $\square$

Let us return to Fig. 3.1 to illustrate how chain groups and boundary maps capture topological structure and the importance of Prop. 3.6. If we consider the simplicial complex in panel (c), we can state all the chain groups and boundary maps. There are two vertices,  $v_1$  and  $v_3$ , and a single edge,  $e_2$ , connecting them. So,  $C_0 = \{0, v_1, v_3, v_1 + v_3\}$ ,  $C_1 = \{0, e_2\}$ , and the remaining chain groups are trivial. The only non-trivial boundary map is  $\delta_1$ , which maps  $e_2$  to its vertices, the chain  $v_1 + v_3$ .

Observe how the connected component can be viewed as the quotient in  $C_0$  of the kernel of  $\partial_0$ , all of  $C_0$ , and image of  $\partial_1$ , the subgroup  $\{0, v_1 + v_3\}$ . This quotient

---

<sup>4</sup>The composition of boundary maps on chain groups over a general ring is also trivial, see [35].

identifies  $v_1$  with  $v_3$ , effectively ‘collapsing’ the two vertices along their connecting edge. The triviality of  $\partial_0\partial_1$  guarantees that this quotient group is well-defined since  $\partial_0\partial_1 = 0$  implies the kernel of  $\partial_0$  is a subgroup of the image of  $\partial_1$ .

In general, a sequence of groups and homomorphisms where quotient groups of the kernel and image of subsequent homomorphisms are well-defined is called a *chain complex*, defined below.

**Definition 3.7.** Let  $C_* = \{\dots, C_{k+1}, C_k, C_{k-1}, \dots\}$  be a sequence of groups and let  $\partial_* = \{\dots, \partial_{k+1}, \partial_k, \dots\}$  be sequence of homomorphisms between them with each  $\partial_k : C_k \rightarrow C_{k-1}$ . We say that  $(C_*, \partial_*)$  is a **chain complex** when

$$\partial_k\partial_{k+1} = 0 \tag{3.1}$$

In a chain group, every quotient group

$$H_k := \frac{\ker \partial_k}{\text{im } \partial_{k+1}} \tag{3.2}$$

is well-defined. The quotient group  $H_k$  is called the  $k$ -th **homology group** of the chain complex.

For a simplicial complex  $X$ , we call the homology groups of its chain complex the *simplicial homology* of the complex. As we have alluded to previously, the simplicial homology of a complex is the same as the singular homology of its carrier as a subset of  $\mathbb{R}^2$ , which captures the subset’s topology [35].

**Proposition 3.8.** Let  $X$  be a simplicial complex. Then the homology groups  $H_k$  of its chain complex are equivalent to the respective singular homology groups of its carrier  $|X|$  up to homomorphism. The dimension  $\dim(H_k)$  is called the  $k$ -th **Betti number**, denoted by  $\beta_k$ .

From this remarkable proposition, we can rest assured that the purely algebraic homology groups of the chain complex of a simplicial complex return information about its topology. In particular,  $\beta_0$  is the number of connected components and  $\beta_1$  is the number of loops. The advantage of simplicial homology is that it can be computed algorithmically, allowing us to calculate the Betti numbers of a simplicial complex.

The algorithmic computation of homology is based on matrix manipulation of matrix representations of the boundary maps, for which we provide a summary here. With respect to the bases of the chain groups, the boundary maps can be represented as  $\mathbb{Z}_2$ -valued matrices. For example, in Fig. 3.1(f), the matrix representation of  $\partial_1$  with respect to the basis of  $\{v_1, v_2, v_3, v_4\}$  for  $C_0$  and  $\{e_1, e_2, e_3, e_4, e_5\}$  for  $C_1$  is

$$M_1 = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix}$$

To compute  $H_1$ , we would need to find the kernel of  $\partial_1$  and the image of  $\partial_2$  in  $C_1$  in a format where we can take their quotient. In the case of the former, we can extract the kernel of  $\partial_1$  from  $M_1$  by using elementary row and column operations to manipulate the matrix into its Jordan-Normal form. The column operations correspond to a change in the basis of  $C_1$ . With the Jordan-Normal form of  $M_1$  and the corresponding new basis of  $C_1$ , we can identify the kernel of  $\partial_1$  as the span of a subset of this new basis, as per the Jordan-Normal form.

We can similarly apply elementary operations to  $M_2$  to find the image of  $\partial_2$ . In this example,  $\partial_2$  is trivial, so its image is simply  $\{0\}$ . However, in general, we first

apply row operations to change the basis of  $C_1$  to that obtained from the column operations on  $M_1$ . Next, we apply column operations to  $M_2$  to manipulate it into row-echelon form, from which we can extract the image as the span of a subset of the new basis of  $C_1$ . Because the kernel of  $\partial_1$  and the image of  $\partial_2$  are taken over the same basis, we can easily compute their quotient.

In general, this is the process for computing  $H_k$ : apply row and column operations to put  $M_k$  in its Jordan-Normal form, then apply row and column operations to put  $M_{k+1}$  in the row-echelon form with the same basis in  $C_k$ . A full description of the algorithmic computation of homology can be found in [35].

### 3.3 Persistent homology of filtered simplicial complexes

We have shown that in a force network, the super-level sets are simplicial complexes. So, for each force threshold, we can associate simplicial homology groups where the elements of  $H_0$  correspond to connected components and the elements of  $H_1$  correspond to loops of the super-level set above that threshold. In this section, we will examine how to track these elements over the homology groups associated with a decreasing sequence of thresholds. The *persistent homology* of the force network captures at what thresholds these elements first appear and when/if they disappear.

Persistent homology requires a *filtration*, a nested sequence of simplicial complexes.

**Definition 3.9.** *Let  $X$  be a simplicial complex and  $f : X \rightarrow \mathbb{R}$  such that each super-level set*

$$X^r = \{\sigma \in X : f(\sigma) \geq r\}$$

*is a sub-complex of  $X$ . For a decreasing sequence  $r_0 > r_1 > \dots > r_m$  we call the nested*

sequence of simplicial complexes

$$\emptyset \subseteq X^{r_0} \subseteq \dots \subseteq X^{r_m} \subseteq X$$

a **filtration** of  $X$ . Equivalently, we also refer to  $X$  as a **filtered simplicial complex** when the filtration is understood.

The force network has a natural filtration over its contact network. Since its force function only takes a finite set of values (defined on the finite set of edges), we can sort these values in descending order into a unique sequence of thresholds  $r_0 > \dots > r_m = 0$ . We refer to the resulting sequence of super-level sets as *the* filtration of a force network. For example, the filtration in Fig. 3.1 is induced by the thresholds 3, 2, 1, and 0 and is the depicted complexes in (b)-(e).

For the remainder of this section, let us consider a general filtered simplicial complex

$$\emptyset \subseteq X^{r_0} \subseteq \dots \subseteq X^{r_m} = X$$

and let  $(C_*^i, \partial_*^i)$  denote the chain complex of the simplicial complex  $X^{r_i}$ .

First, note that the natural inclusion map from each  $X^{r_i} \rightarrow X^{r_{i+1}}$  maps each simplex at filtration level  $r_i$  to the same simplex at level  $r_{i+1}$ . Based on this natural inclusion, we can construct an inclusion map between chain groups, the map  $\iota^i : C_*^i \rightarrow C_*^{i+1}$  which maps each simplex  $\sigma \in C_*^i$  to the same simplex in  $C_*^{i+1}$  and extends linearly to general chains.

Next, observe that the inclusion maps commute with the boundary maps. To see why, consider a  $k$ -simplex  $\sigma \in C_k^i$ . Applying the boundary map results in the



$k - 1$ -simplex  $\partial_k \sigma \in C_{k-1}^i$  consisting of the formal sum

$$\partial_k \sigma = \sum_i \sigma_{\hat{i}}$$

where  $\sigma_{\hat{i}}$  is shorthand for removing the  $i$ -th vertex, as in Def. 3.5. Applying the inclusion map results in the same formal sum, now at the filtration level  $i + 1$ ,  $\iota^i \partial_k \sigma \in C_{k-1}^{i+1}$

$$\iota^i \left( \sum_i \sigma_{\hat{i}} \right) = \sum_i \iota^i \sigma_{\hat{i}} = \sum_i \sigma_{\hat{i}}$$

Conversely, applying the inclusion map first yields the same formal sum in  $C_{k-1}^{i+1}$ :

$$\partial_k \iota^i \sigma = \partial_k \sigma = \sum_i \sigma_{\hat{i}}$$

Because of this commutativity, we can consider an inclusion-induced homomorphism from the homology groups at the  $i$ -th filtration level to the homology groups at the  $i + 1$ -th filtration level<sup>5</sup>. This is made precise in the following proposition.

**Proposition 3.10.** *For a filtered simplicial complex, define the map  $\phi_k^i : H_k^i \rightarrow H_k^{i+1}$  by*

$$\phi_k^i([c_k]^i) = [\iota^i(c_k)]^{i+1}$$

where  $[c_k]^j$  indicates the element of the quotient group  $H_k^j$  with a representative  $c_k \in \ker \partial_k \subseteq C_k^i$ . Then the map  $\phi_k^i$  is a well-defined homomorphism.

*Proof.* Suppose  $[c_k]^i \in H_k^i$ , then  $c_k \in \ker \partial_k \subseteq C_k^i$ . The first condition for the map in Prop. 3.10 to be well-defined specifies that the mapping of the representative  $c_k$  at the  $i + 1$ -th level must be in the kernel of the boundary map. By commutativity of the inclusion and boundary maps,  $\partial_k(\iota^i(c_k)) = 0$ , so  $\iota^i(c_k) \in \ker \partial_k \subseteq C_k^{i+1}$ .

---

<sup>5</sup>For more details on the rich theory behind the structure of homology groups over a filtration, see [35].

Second, we must show that the map, which is stated using a representative of the equivalence class of the quotient groups, is well-defined over the equivalence class. Suppose  $[\tilde{c}_k]^i = [c_k]^i$ . That is,  $\tilde{c}_k - c_k$  is in the image of  $\partial_{k+1}$  in  $C_k^i$ . By commutativity of the inclusion and boundary maps,  $\iota^i(\tilde{c}_k - c_k)$  is in the image of  $\partial_{k+1}$  in  $C_k^{i+1}$ . Then by linearity of the inclusion map,  $\iota_i(\tilde{c}_k) - \iota_i(c_k)$  is in the image of  $\partial_{k+1}$ , so  $[\iota_i(\tilde{c}_k)]^{i+1} = [\iota_i(c_k)]^{i+1}$ .

So,  $\phi_k^i$  is well-defined. Lastly, by linearity of the inclusion map,  $\phi_k^i$  is a homomorphism. □

With the sequence of homomorphisms between the homology groups over the filtration, we can track the appearance and disappearance of different elements. For each element  $m_k \in H_k^i$ , if  $m_k$  is not in the image of the homomorphism from the previous filtration level,  $\phi_k^{i-1} : H_k^{i-1} \rightarrow H_k^i$ , then we say that the element is *born* at the filtration threshold  $r_i$ . For such an element, if we can identify the first level  $j$  where  $\phi_k^j(\dots\phi_k^i(m_k)) = 0$ , then we say the element *dies* at the filtration threshold  $r_j$ . If this does not occur, then we say the element  $m_k$  persists through the entire filtration.

Using this approach, we can parameterize the  $k$ -level persistent homology by a multiset of birth-death pairs [35]. In addition to the persistence pairs, we also consider a countable set of abstract elements representing copies of the *diagonal*,  $\{\Delta_1, \Delta_2, \dots\}$ . Each copy  $\Delta_i$  of the diagonal represents the diagonal line  $b = d$  where birth is equal to the death, and can be thought of as a topological feature that instantly disappears [26]. The distance between a persistence pair and a copy of the diagonal is the shortest distance from that point to the diagonal line, while the distance between two copies of the diagonal is zero. This collection is called the *persistence diagram* of the filtered complex, defined below.

**Definition 3.11.** *Let  $X$  be a filtered simplicial complex. For each  $k \in \mathbb{Z}$ , the  $k$ -level **persistence diagram** of a filtered simplicial complex, denoted  $PD_k$ , is the multiset of ordered pairs of the form  $(r_i, r_j)$  with  $r_i > r_j \in \mathbb{Z} \cup \{-\infty\}$  be the birth-death pairs of the  $k$ -th level persistent homology of  $X$ , along with a countable set of abstract elements  $\{\Delta_1, \Delta_2, \dots\}$ . Each  $\Delta_i$  is called a copy of the **diagonal**. Each pair represents a  $k$ -dimensional homology element that first appears at filtration threshold  $r_i$  and persists until disappearing at filtration threshold  $r_j$  (or through the entire filtration if  $r_j = -\infty$ ). Further, a metric  $d$  is placed on the space of persistence diagrams, where  $d(x, y)$  is the Euclidean distance for two persistence pairs, the shortest distance to the diagonal line if  $x$  or  $y$  is a persistence pair and the other is a copy of the diagonal, and zero if both  $x$  and  $y$  are copies of the diagonal.*

Similarly to the computation of homology of simplicial complexes, persistent homology can be calculated by manipulation of matrix representations. In this case, the boundary maps across filtration levels are represented by  $\mathbb{Z}_2[t]$ -valued matrices that encode the filtration levels of simplices. Details on an algorithm for this computation can be found in [35]. The described algorithm, similar to Gaussian elimination, is  $O(N^3)$  in the worst case, where  $N$  is the number of simplices. The authors note that the typical behavior appears to be  $O(N)$ .

### 3.4 Distances between persistence diagram representations of successive force networks

To quantify the topological change between successive force networks in the granular system, we will consider metrics that quantify the difference between the persistence diagrams of successive force networks. While there are numerous such metrics, we will consider the family of *Wasserstein distance for persistence diagrams*, which are based

on the general principle of measuring the minimal ‘work’ required to move the points in one diagram to the other. These metrics are closely related to the Wasserstein distance for probability measures, whose historical context we explore below. We will show the distinction between the two, which is intrinsically connected to the role of the ‘diagonal’ in persistence diagrams.

The Wasserstein distance for probability measures is historically formulated in the context of optimal transport [26]. The classical optimal transport problem is as follows: for a distribution  $\mu$  on a space  $X$ , we wish to find the minimal work to ‘transport’ the mass represented by  $\mu$  to a second distribution  $\nu$  on  $X$ . In the Kantorovich formulation, a transport plan is a measure on the product  $X \times X$  indicating the mass transported. More precisely,  $\gamma$  is joint distribution with marginals  $\mu$  and  $\nu$ . If  $w : X \times X \rightarrow \mathbb{R}$  is the function where  $w(x, y)$  is the work to transport mass from  $x$  to  $y$ , the minimal work is the infimum of  $\int_{X \times X} wd\gamma$  over transport plans.

The  $p$ -th Wasserstein distance for probability measures is defined as this quantity raised to the  $1/p$ -th power where the work is the distance raised to the  $p$ -th power. For example, the 1-th Wasserstein distance for probability measures is exactly the minimal distance over transport plans where the work required is simply distance. This is formulated precisely in the definition below.

**Definition 3.12.** *Let  $(X, d)$  be a metric space, and  $\mu, \nu$  be two distributions on  $X$ . The  $p$ -th **Wasserstein distance** is given by*

$$W_p(\mu, \nu) = \left( \inf_{\gamma \in \Gamma(\mu, \nu)} \int d(x, y)^p d\gamma \right)^{1/p} \tag{3.3}$$

where  $\Gamma(\mu, \nu)$  is the set of joint distributions with marginals  $\mu$  and  $\nu$ , the allowable transport plans.

Before moving on, we note that when  $\mu$  and  $\nu$  are empirical distributions with

the same number of elements, then the Wasserstein distance reduces to an infimum over permutations of the elements [14]. Without the diagonal, a persistence diagram can be viewed as an empirical distribution on  $\mathbb{R}^2$ . However, the classical formulation of the optimal transport problem does not apply, since we wish to allow points (i.e., persistence pairs) to be transported to the diagonal. Nevertheless, the underlying principle is still the same; the Wasserstein distance between two persistence diagrams can be interpreted as the minimal work to move the points from one diagram to the other.

With this in mind, we will define the Wasserstein distance for persistence diagrams through bijections, the standard approach in the literature [26]. To build up the intuition behind a transport plan for persistence diagrams, consider the two example persistence diagrams in Fig. 3.3(a)-(b), where (a) is the zero-level persistence diagram of the example force network in Fig. 3.1. Panel (b) depicts a slightly different zero-level persistence diagram for a fictional force network. There are an infinite number of bijections between the two diagrams, since each persistence pair in (a) can be mapped to not only another pair in (b) but any point on the diagonal. Fig. (c) shows one example of such a bijection in which the points at infinity are mapped together, the short-lived point in (b) is mapped to the diagonal, and the remaining pairs are mapped together.

Recall that a persistence diagram PD is a finite multi-set of persistence pairs along with countable copies of the diagonal,  $\{\Delta_1, \Delta_2, \dots\}$ . Consequently, a bijection can always be found between two persistence diagrams. Next, recall that the distance  $d$  between a persistence pair and any  $\Delta_i$  is the shortest distance from the point to the diagonal line. Additionally, the distance from two copies of the diagonal, some  $\Delta_i$  and  $\Delta_j$  is zero. So, for a bijection between two persistence diagrams, we can consider a metric of distance based on the distances between mapped elements in the bijection.




Figure 3.3: Panel (a) shows the zero-level persistence diagram of the force network of Fig. 3.1, while (b) is the zero-level persistence diagram of a different example force network. Panel (c) depicts a possible bijection between the diagrams.

This is the basis of the Wasserstein distance for persistence diagrams, defined below.

**Definition 3.13.** *Let  $PD$  and  $PD'$  be persistence diagrams and  $p \in [1, \infty)$ . The  $p$ -th Wasserstein distance for persistence diagrams is defined as the following infimum over bijections between  $PD$  and  $PD'$ ,  $\Gamma(PD, PD')$ :*

$$W_p(PD, PD') = \inf_{\gamma: PD \rightarrow PD'} \left( \sum_{x \in PD} d(x, \gamma(x))^p \right)^{\frac{1}{p}} \quad (3.4)$$

*After this point, the term Wasserstein distance will be used exclusively to refer to the Wasserstein distance for persistence diagrams.*

Note that for any  $p \in [1, \infty)$  the term inside the infimum in Def. 3.13 is non-negative. So, the Wasserstein distance is always non-negative, though it need not be finite. Indeed, if two diagrams have a different number of points at infinity, their Wasserstein distance is infinite.

To guarantee a finite Wasserstein distance, we follow the convention that the lowest filtration level of a force network is the high-dimensional simplex composed of all the vertices in the network. E.g., if the network has  $k + 1$  vertices  $v_0, \dots, v_k$ , this is

the  $k$ -simplex  $[v_0, \dots, v_k]$ . While generally this simplex cannot be embedded in  $\mathbb{R}^2$ , it still fits the abstract algebraic formulation of simplicial complexes in Sec. 3.1. This level is artificially placed at a threshold below the contact network.

Of course, from the introduction of this artificial level, the persistent homology is no longer that of the original force network. However, the alterations are fairly straightforward since this complex is a single component topologically equivalent to a point. Applying this convention results in all components being connected and all loops being filled in. So, all zero-level diagrams have a single point at infinity (which can be mapped together in bijections) and all higher-level diagrams have no points at infinity. Fig. 3.4 depicts the persistence diagram of the simple example in Fig. 3.1 when this convention is applied. Observe the minimal difference to Fig. 3.1(f)-(g).

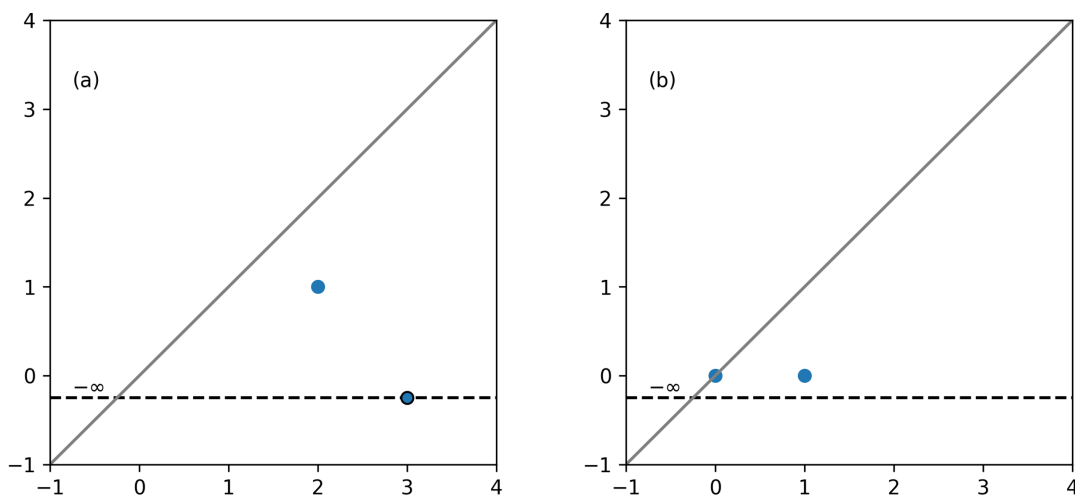


Figure 3.4: (a) The zero-level persistence diagram and (b) the one-level persistence diagram of the simple force network in Fig. 3.1(b) under the convention that an artificial lowest level is the four-dimensional simplex  $[v_1, v_2, v_3, v_4]$ . Compared to the true persistence diagrams in Fig. 3.1(g)-(h),  $PD_0$  is unchanged, and in  $PD_1$  the two points at infinity now have a death at the threshold 0.

One of the motivating factors for considering the Wasserstein distances between

filtered simplicial complexes is stability: small changes in the force network induce small changes in the persistence diagrams, as measured by the Wasserstein distances. In the inequality below, the functions  $f$  and  $g$  are abstractions of the filtration of a filtered simplicial complex. Note that the convention of having the lowest level of all force network filtrations artificially defined as the simplex of all vertices results in a uniform domain, allowing successive force networks to be directly comparable, as in the following theorem. For the proof of a more general version, see [8].

**Theorem 3.14.** *Let  $X$  be a triangulable, compact metric space and  $f, g : X \rightarrow \mathbb{R}$  be tame Lipschitz functions. Then the  $p$ -th Wasserstein distance between the persistence diagrams<sup>6</sup> of the respective functions, denoted  $W_p(f, g)$  satisfies*

$$W_p(f, g) \leq C^{\frac{1}{p}} \cdot \|f - g\|_{\infty}^{1 - \frac{k}{p}}$$

for constants  $C$  and  $k$  that depend on the Lipschitz constants of  $f$  and  $g$  and the space  $X$ .

So, the  $p$ -th Wasserstein distance is stable in the sense that an increase in the value of  $W_p$  between successive force networks implies a corresponding increase in the maximal difference between those force networks. Conversely, if the maximal force differential is smaller, the value of  $W_p$  will be correspondingly smaller.

The  $p$ -th Wasserstein distance is analogous to  $l_p$  norms. At the extreme, the absolute Wasserstein distance  $W_1$  is the sum of all changes between the infimal transport plan. So, the bottleneck distance is highly sensitive to outliers in differences between the persistence diagrams (similar to the  $l_{\infty}$  norm). In contrast, the  $W_1$  is dominated by the average difference between persistence pairs, and larger changes can be obscured when there is a large number of pairs, which is the case for the considered

---

<sup>6</sup>This inequality holds for the persistence diagrams at any level.



granular system.

In our work, we use the second Wasserstein distance, a middle ground between those two extremes. Large changes are less obscured than with the  $W_1$ , but outliers do not dominate the measure. We define the following univariate measure on the granular system, a sample of which is shown at the end of this chapter in Fig. 3.6(c).

**Definition 3.15.** *For the considered sheared granular system, let  $PD_0(n)$  denote the zero-level persistence diagram of the force network at frame  $n$ . The second Wasserstein distance, hereafter denoted by  $W_2$ , is given by*

$$W_2(n) = W_2(PD_0(n-1), PD_0(n)).$$

### 3.5 Additional measures of change between successive force networks

To analyze the evolution of the sheared granular system, we focus on measures that provide some metric of change between successive force networks. One simple measure which we have already encountered in Ch. 2, is the horizontal wall velocity  $v_x$ . We have also introduced a persistent homology-based measure, the zero-level second Wasserstein distance  $W_2$ , which is a metric on the changes between the persistence diagrams of successive force networks of the granular system.

However, Wasserstein distances, and persistent homology in general, are far from the only way to capture changes between successive force networks of a granular system. A more conventional metric we will consider is based on the *differential force network* (DFN) of the sheared granular system. The DFN at frame  $n$ , denoted by  $DFN(n)$ , is a function over the union of the contact networks at frames  $n-1$  and  $n$ . For an edge present in both, the differential force is the difference between the force values at frames  $n-1$  and  $n$ , while for an edge present at only one of the frames, the

differential force is the force value at that frame.

The measure on the differential force network that we consider is the *percolation force*. The percolation force is the largest differential force level  $df^*$  such that the super-level set of  $DFN(n)$  above  $df^*$  is a graph that ‘percolates’ through the system. More precisely, we give the following definition for the left-right percolation force.

**Definition 3.16.** *The **left-right percolation force**  $f_{plr}(n)$  is the maximal value  $df^*$  such that the super-level set of  $DFN(n)$  above  $df^*$  has a connected component in contact with both the left and right boundaries of the domain of the granular system<sup>7</sup>.*




Figure 3.5: Example of the differential force network  $DFN$  in (a), and the resulting super-level set at the level of the left-right percolation force  $f_{plr}$  in (b). The largest connected component in (b) indeed connects the left and right walls.

The value  $f_{plr}(n)$  marks the force at which a force chain percolates across the entire differential force network. An example of this is shown in Figure 3.5, with the corresponding super-level set shown in panel (b). The left and right walls are

<sup>7</sup>A similar measure is the top-bottom percolation force  $f_{ptb}$ . In the analyses we performed, we found that these measures are interchangeable, so we focus exclusively on  $f_{plr}$  for brevity.

connected by that super-level set, while the super-level set for a larger differential force threshold would not connect the walls. So, higher values of this measure indicate larger changes in the successive force networks that percolate across the entire system.

Note that localized changes in the differential force network are not necessarily reflected in the percolation force value. For example, in Fig. 3.5 suppose that at one of the edges, the force value at frame  $n$  is significantly larger, increasing the differential force of that edge accordingly. In all likelihood, with such a localized change the left and right walls would remain unconnected for values larger than  $f_{plr}$ , and the percolation force would remain unchanged. In contrast, the persistent homology at frame  $n$  would almost certainly reflect this changed force value, changing the value of the  $W_2$  metric accordingly.

We refer to the percolation force as a *global measure* in the sense that it is relatively unaffected by localized changes in the force network. Similarly, we refer to the  $W_2$  as a *local measure* as it is more sensitive to localized changes in the force network. For more examples of global and local measures on the force network, see [11]. A sample of the left-right percolation force, along with the  $W_2$  and horizontal top wall velocity  $v_x$  is shown in Fig. 3.6.

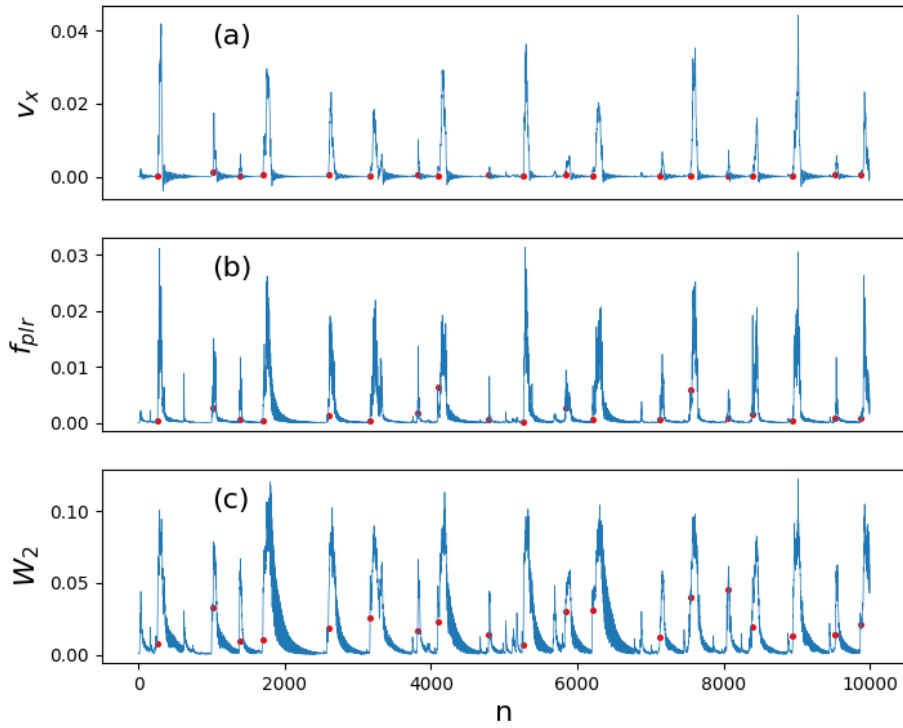


Figure 3.6: Sample of the evolution for first 10,000 frames in (a) the horizontal velocity of the top wall  $v_x$ , (b) the right left percolation force  $f_{plr}$ , and (c) the  $W_2$ . The red dots mark the start of large slip times, as defined using the offline method at the beginning of this chapter. The horizontal wall movement  $v_x$  shows two alternating regimes of behavior: long periods of stick, in which the wall moves slowly, and shorter periods of slip, where the wall shifts rapidly. In many of these instances, the  $f_{plr}$  and  $W_2$  measures already show movement before the identified slip start times.

## Chapter 4

### Bayesian framework for modeling considered measures

In Ch. 3 we introduced three measures that we hope have the potential to predict slips: the wall velocity  $v_x$ , the percolation force  $f_{plr}$ , and the Wasserstein distance  $W_2$ . The wall velocity is the baseline against which slip is defined, while the latter two measures often increase before the wall has begun to move appreciably. An example of such an increase can be seen in the sample shown in Figure 4.1. In particular, the  $W_2$  increases well before the start of the slip near  $n = 1000$ .

Ideally, we wish to identify a slip right at the beginning of such increases, which in the case of  $f_{plr}$  and  $W_2$  are before the start of the slip. However, recall from Ch. 2 that the oscillations of  $v_x$  during the stick regime, highlighted in the inset of Fig. 4.1(a), are an obstacle to using a single threshold to identify the start of a slip. Consequently, we developed an offline method for slip identification, in which the start of a slip is identified using later data. In the same way, the oscillations in the  $f_{plr}$  and  $W_2$  measures, highlighted in the insets of (b)-(c), impede the identification of their increase before the slip regime.

To identify when the measures start changing their behavior and increasing, we construct models. For each measure, the model captures its dynamic behavior during the stick regime: the oscillations and trend (slow decrease). We consider stochastic state space models, which consist of a deterministic element that captures the modeled behavior (e.g., the oscillations and overall decrease) and a stochastic element that models uncertainty in a Bayesian manner. Each model produces predictions of its

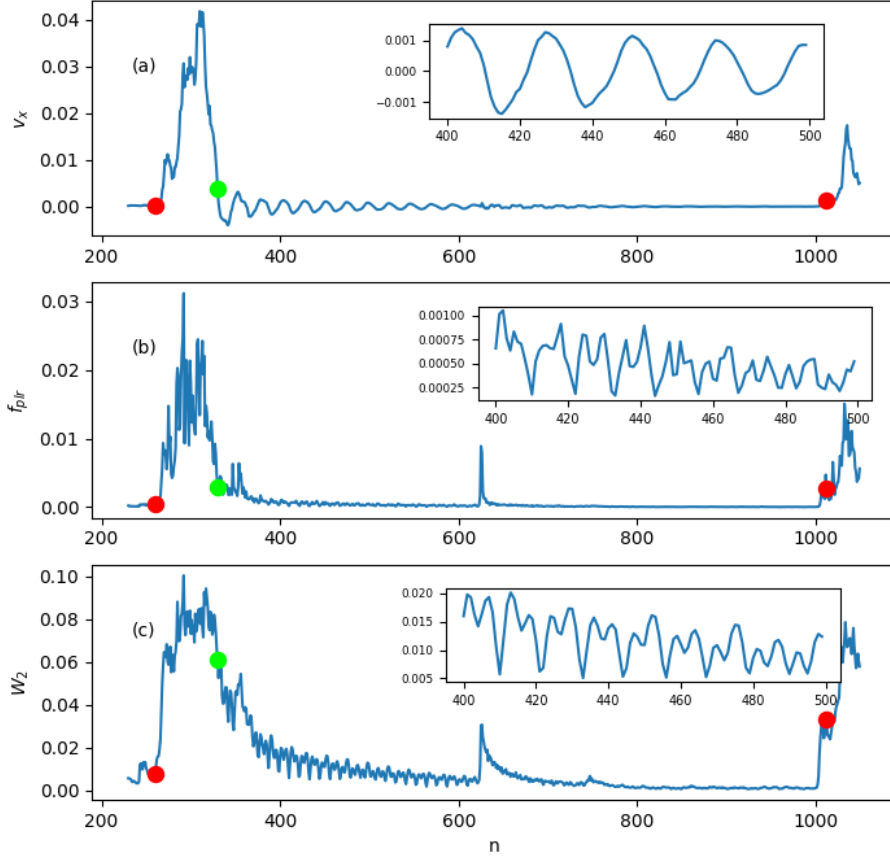


Figure 4.1: Small sample of the evolution of (a) the horizontal velocity of the top wall  $v_x$ , (b) the right left percolation force  $f_{plr}$ , and (c) the  $W_2$  from just before the onset of a slip event to shortly after the start of the next slip. The insets show a sample of the stick regime for each of the measures.

measure's future value during the stick regime, and since the  $v_x$ ,  $f_{plr}$ , and  $W_2$  all have different behavior during the stick and slip regimes, the predictive error (the difference between the prediction and the observed value) rises sharply upon the onset of the slip regime. The predictive error can then be used to predict slips by identifying sharp increases of the type in Fig. 4.1 as they occur.

## 4.1 Introduction to Bayesian statistics

The Bayesian approach models unknown quantities through probability distributions. These distributions are updated as our knowledge about the unknown quantity(ies) increases with observed data. While this is a fairly simple principle, it has a rich mathematical basis in formal probability theory. To make the abstract ideas more concrete, throughout this section we will use a toy example to describe and illustrate the Bayesian approach.

In this example, we will suppose that we have recorded a series of temperature measurements  $y_1, \dots, y_N$  taken by a thermometer in a room, where the true room temperature at the time of the  $n$ -th measurement is  $\theta_n$ . First, we will suppose that the true room temperature is constant over the measurement period. However, to illustrate stochastic state space models in the next section, we will expand the example to consider the case where the room temperature increases due to a heater being turned on. The corresponding stochastic state space model will then allow us to statistically infer information about each unknown temperature  $\theta_n$  given the measurements up to that time,  $y_1, \dots, y_n$ . But for now, suppose  $\theta_1 = \dots = \theta_N$ , and denote the unknown temperature by  $\theta$ .

Fig. 4.2 shows an example of what this series of temperature measurements could look like, with the (unknown to us) true temperature  $\theta$  marked by the dashed line. If the thermometer perfectly measured the room temperature, then each  $y_n$  would be equal to  $\theta$ . However, measurements are imprecise, resulting in observation errors. For simplicity, we will suppose that we are given a statistical model for each  $y_n$  that describes the imprecision in measurement for the thermometer as well as interactions between the measurements, or lack thereof. This statistical model is given by a *parameterized distribution* for each measurement. To make this more precise, we




Figure 4.2: Simulated observed thermometer measurement values, with the true room temperature marked by the dashed line.

need to introduce the following nomenclature. We start with the measure-theoretic definition of a probability space below, and assume a level of familiarity with basic measure theory.

**Definition 4.1.** *Let  $\Omega$  be a set,  $\mathcal{F}$  be a  $\sigma$ -algebra on  $\Omega$ , and  $P$  be a measure on  $(\Omega, \mathcal{F})$ . We say that the triple  $(\Omega, \mathcal{F}, P)$  is a **probability space** if*

$$P(\Omega) = 1$$

The above definition of a probability space formally captures the notion of outcomes, events, and the probability of an event. In that definition, an outcome is any element  $\omega \in \Omega$ , an event is a subset  $A \subseteq \Omega$  in the  $\sigma$ -algebra  $A \in \mathcal{F}$ , and the probability of event  $A$  is defined as  $P(A)$ . For example, in the context of the toy scenario, outcomes are the temperature measurements, and an appropriately selected probability measure will describe the relative likelihood of different temperature measurements. To begin to formalize the notion of relative likelihood, we first define a real-valued *random variable* and its *cumulative distribution function*.



**Definition 4.2.** For a probability space  $(\Omega, \mathcal{F}, \hat{P})$ , we say that a measurable function  $X : \Omega \rightarrow \mathbb{R}$  is a real-valued **random variable** and define a probability measure  $P$  on  $\mathbb{R}$  by

$$P(E) = \hat{P}(\{\omega \in \Omega : X(\omega) \in E\}) \quad (4.1)$$

for each measurable  $E \subseteq \mathbb{R}$ . The probability measure  $P$  is called the **distribution** of  $X$  and  $P(E)$  is interpreted to be the probability  $X$  takes a value in  $E$ . The **cumulative distribution function** of  $X$  is the function  $F_X : \mathbb{R} \rightarrow [0, 1]$  given by

$$F_X(x) = P((-\infty, x]) \quad (4.2)$$

In the toy scenario, the statistical model for observations with observation error will consist of a sequence of random variables  $Y_1, \dots, Y_N$  where the random variable  $Y_n$  is the real-valued temperature measurement of the thermometer when the room temperature is  $\theta$ . If we are given real-valued random variables describing the observation results, we can quantify the probability of each measurement  $Y_n$  taking certain values through its distribution. In fact, the distribution of a real-valued random variable can be fully specified through a cumulative distribution function, which we will make clear shortly.

The cumulative distribution function of any real-valued random variable satisfies certain properties. Suppose  $F_X : \mathbb{R} \rightarrow [0, 1]$  is a cumulative distribution function. Then  $F_X(x) \rightarrow 0$  as  $x \rightarrow -\infty$ ,  $F_X(x) \rightarrow 1$  as  $x \rightarrow \infty$ , it is non-decreasing, and it is right-continuous [20]. Conversely, for any function  $F$  satisfying those four properties, a random variable can be constructed whose cumulative distribution function is exactly  $F$ . We make this claim explicit in the following proposition, whose proof can be found in [20].

**Proposition 4.3.** *Let  $F : \mathbb{R} \rightarrow [0, 1]$  be a function that is non-decreasing, right-continuous, and where*

$$\lim_{x \rightarrow -\infty} F(x) = 0$$

and

$$\lim_{x \rightarrow \infty} F(x) = 1$$

*Then there exists a random variable  $X$  with a cumulative distribution function  $F_X(x) = F(x)$  for all  $x \in \mathbb{R}$ . Moreover, two random variables with the same cumulative distribution function have the same distribution.*

Because of Prop. 4.3, it is common practice to characterize a real-valued random variable with its cumulative distribution function, or, as we will show shortly, an equivalent. While the cumulative distribution function is defined for any real-valued random variable, the temperature measurement random variables  $Y_1, \dots, Y_N$  provided to us will be of a particular class. Namely, they are continuous random variables, since the measurements take values continuously rather than discretely over  $\mathbb{R}$ . More precisely, each  $Y_n$  is *absolutely continuous*, according to the definition given below.

**Definition 4.4.** *We say that a real-valued random variable  $X$  is **absolutely continuous** if its cumulative distribution function  $F_X$  is absolutely continuous. In this case, the derivative of  $F_X$ , which we will denote by  $f_X$ , is defined almost everywhere and*

$$P([a, b]) = \int_a^b f_X(x) dx$$

*The function  $f_X$  is called the **probability density function** of  $X$ . When the random variable in reference is understood, we typically leave off the subscript. We further*

say that the **mean** of  $X$  is the quantity

$$\bar{X} = \int_{\mathbb{R}} x f_X(x) dx$$

and the **variance** of  $X$  is the quantity

$$V(X) = \int_{\mathbb{R}} (x - \bar{X})^2 f_X(x) dx$$

when these integrals are defined.

The distribution of an absolutely continuous random variable is fully specified by its probability density function. An example of a simple but extremely useful distribution is the *normal distribution*, also sometimes referred to as the *Gaussian distribution*. The probability density function of the normal distribution is given below.

**Definition 4.5.** For  $\theta, V \in \mathbb{R}$ , the probability density function of a normally distributed random variable  $X$  with mean  $\theta$  and variance  $V$  is given by

$$f_X(x) = \frac{1}{\sqrt{2\pi V}} e^{-\frac{(x - \theta)^2}{2V}}$$

To denote that  $X$  is distributed normally with mean  $\theta$  and variance  $V$ , we write

$$X \sim \mathcal{N}(\theta, V)$$

A Gaussian distribution is fully characterized by two quantities, its mean and variance. Values closer to the mean are more likely than those that are further away, and the distribution is symmetric around its mean (an outcome  $x$  less than the mean

is exactly as likely as an outcome  $x$  greater than the mean). For further details on Gaussian distributions, see [20] pp. 45-47, 51-53.

In our example, we suppose that each of the thermometer's measurements  $Y_n$  is distributed normally with a mean of  $\theta_n$  and a specified variance  $V$ . Due to the translation invariance of the Gaussian distribution [20], this is equivalent to supposing that the observation errors,  $Y_n - \theta_n$ , are distributed normally with a mean of zero and a variance  $V$ . Consequently, we will refer to  $V$  as the *observation error variance*. Under this model, the thermometer does not systematically over or underestimate the true room temperature, and the scale of the measurement error is the same with each measurement, given by the observation error variance  $V$ .

Note that the probability density function of each random variable  $Y_n$  depends on the value of the (unknown) room temperature  $\theta_n$ . When the distribution of a random variable depends on an unknown value as  $Y_n$  does on  $\theta_n$ , we say that the random variable is *parameterized* by the unknown value. The precise definition of a parameterized distribution is given below.

**Definition 4.6.** Consider a multivariate function  $f : \mathbb{R} \times \mathbb{R}^d \rightarrow [0, \infty)$  where, for each fixed value of  $\boldsymbol{\theta} \in \mathbb{R}^d$ , the function  $f(\cdot, \boldsymbol{\theta})$  is a probability density function of an absolutely continuous random variable denoted by  $X; \boldsymbol{\theta}$ . We denote this family of probability density functions by

$$f_{X; \boldsymbol{\theta}}(x; \boldsymbol{\theta})$$

and say that the corresponding family of random variables  $X; \boldsymbol{\theta}$  is **parameterized** by  $\boldsymbol{\theta}$ .

So,  $Y_n$  is parameterized by the unknown temperature  $\theta_n$ <sup>1</sup> with the distribution

$$Y_n; \theta_n \sim \mathcal{N}(\theta_n, V)$$

In addition to providing this parameterized distribution function, the statistical model must consider potential interactions between the random variables  $Y_1, \dots, Y_N$ . That is, it is certainly possible that the taking of one measurement could influence the value of another. In our example, the model will assume that any such influence is negligible, a concept known as *independence* between the random variables. To define independence, we first define a multivariate random variable to formally describe a collection of random variables that could potentially influence each other.

**Definition 4.7.** For a probability space  $(\Omega, \mathcal{F}, \hat{P})$ , we say that a measurable function  $\mathbf{X} : \Omega \rightarrow \mathbb{R}^n$  is a ***multivariate random variable*** and define a probability measure  $P$  on  $\mathbb{R}^n$  by

$$P(E) = \hat{P}(\{\omega \in \Omega : \mathbf{X}(\omega) \in E\}) \quad (4.3)$$

for each measurable  $E \subseteq \mathbb{R}^n$ . The probability measure  $P$  is called the ***joint distribution*** of  $\mathbf{X}$ . The ***joint cumulative distribution function***  $F_{\mathbf{X}} : \mathbb{R}^n \rightarrow \mathbb{R}$  is given by

$$F_{\mathbf{X}}(\mathbf{x}) = P((-\infty, x_1] \times \dots \times (-\infty, x_n])$$

where  $\mathbf{x} = (x_1, \dots, x_n)$ . Lastly, we say that  $\mathbf{X}$  is ***absolutely continuous*** if  $F_{\mathbf{X}}$  is absolutely continuous, and in that case the ***joint probability density function***  $f_{\mathbf{X}}$

---

<sup>1</sup>We do not consider  $Y_n$  to be parameterized by  $V$ , as the value of  $V$  is a fixed quantity provided to us as part of the statistical model.

is defined almost everywhere and is characterized by

$$P(E) = \int_E f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}$$

for each measurable  $E \subseteq \mathbb{R}^n$ .

Independence of random variables can be characterized through *conditional probability*, a concept based on the joint distribution. The definition below defines conditional probability for two absolutely continuous random variables and can naturally be extended to a collection of  $n$  random variables.

**Definition 4.8.** Suppose two absolutely continuous real-valued random variables  $X$  and  $Y$  have a joint probability density function denoted by  $f_{X,Y}(x, y)$ . The **conditional probability density function** of  $Y$  given  $X$ , denoted by  $f_{Y|X}$ , is defined as

$$f_{Y|X}(y|x) = \frac{f_{X,Y}(x, y)}{f_X(x)}$$

For each fixed value of  $x$ , the resulting univariate density function is interpreted to be the density function of the random variable  $Y$  after observing the value  $X = x$ . The conditional density is undefined when  $f_X(x) = 0$ .

Using the terms in the above definition, along with its interpretation, we say that  $Y$  is *independent* of  $X$  if the members of the family of conditional densities  $f_{Y|X}(y|x)$  are the same regardless of the value of  $x$ ;

$$f_{Y|X}(y|x) = f_Y(y)$$

This is equivalent to the joint density function factoring into the product of the

individual density functions,

$$f_{X,Y}(x,y) = f_X(x)f_Y(y)$$

In the same way, we say that a collection of absolutely continuous real-valued random variables  $X_1, \dots, X_n$  captured in a multivariate random variable  $\mathbf{X} = (X_1, \dots, X_n)$  are *jointly independent* when

$$f_{\mathbf{X}}(\mathbf{x}) = f_{X_1}(x_1) \cdot \dots \cdot f_{X_n}(x_n)$$

In this sense, the provided statistical model of the observation variables  $Y_1, \dots, Y_N$  assumes that they are jointly independent by supposing that the joint probability density function of  $\mathbf{Y} = (Y_1, \dots, Y_N)$  is the product of the individual density functions of each  $Y_1, \dots, Y_N$ .

With the provided statistical model for the observation error, we now turn to the fundamental question of statistics: what information can we infer about the unknown room temperature  $\theta$  after observing the thermometer values  $y_1, \dots, y_N$ ? The Bayesian approach to this question relies on two main components: the *likelihood function* and the *prior distribution*. We will begin by defining the likelihood function, a function of the unknown parameter  $\theta$  that returns the relative likelihood of observing the values  $y_1, \dots, y_N$  for each fixed value of  $\theta$ .

**Definition 4.9.** Let  $X; \boldsymbol{\theta}$  be a continuous random variable parameterized by  $\boldsymbol{\theta}$ . Recall that the family of density functions  $f_{X; \boldsymbol{\theta}}(x; \boldsymbol{\theta})$  is a function  $f : \mathbb{R} \times \mathbb{R}^d \rightarrow [0, \infty)$ . For the fixed value of  $X; \boldsymbol{\theta} = x$ , the real-valued function  $f(x, \cdot)$  of  $\boldsymbol{\theta}$  is called the **likelihood function**, denoted by

$$L(x; \boldsymbol{\theta})$$

Similarly, the joint likelihood function  $L(\mathbf{x}; \boldsymbol{\theta})$  of a parameterized multivariate continuous random variable  $\mathbf{X}; \boldsymbol{\theta}$  for the fixed value of  $\mathbf{X}; \boldsymbol{\theta} = \mathbf{x}$  is the real-valued function of the parameter given by the joint probability density function evaluated at  $\mathbf{x}$ .

As the name suggests, the likelihood function  $L(x; \boldsymbol{\theta})$  for a parameterized random variable  $X; \boldsymbol{\theta}$  describes how likely the outcome  $x$  is for different values of  $\boldsymbol{\theta}$ . One important characteristic of the likelihood function is that, for a collection of outcomes  $\mathbf{x} = (x_1, \dots, x_n)$  of independent random variables  $\mathbf{X}; \boldsymbol{\theta}$  parameterized by  $\boldsymbol{\theta}$ , the likelihood function factors into the product of the individual likelihood functions:

$$L(\mathbf{x}; \boldsymbol{\theta}) = L(x_1; \boldsymbol{\theta}) \cdot \dots \cdot L(x_n; \boldsymbol{\theta})$$

So, in the toy example, the statistical model of independent normal distributions yields the following likelihood function:

$$L(y_1, \dots, y_N; \theta) = \prod_{n=1}^N \frac{1}{\sqrt{2\pi V}} e^{-\frac{1}{2} \left( \frac{y_n - \theta}{\sqrt{V}} \right)^2}$$

Since the quantities  $y_1, \dots, y_N$  and  $V$  are fixed and known, this truly is a function of the single variable  $\theta$ .

The likelihood function is the first component of Bayesian statistics. It is the result of the statistical model connecting the unknown parameter  $\theta$  to the observation(s). The second component of Bayesian statistics is the *prior distribution*, which describes our initial knowledge about the unknown quantity  $\theta$  before observing the data. That is, in the Bayesian approach the unknown parameter  $\theta$  is treated as a random variable, which is the key difference between Bayesian and classical frequentist statistics [20].

Bayesian inference combines the likelihood function and the prior distribution to produce a new distribution of  $\theta$  called the *posterior distribution*. The posterior dis-



tribution captures the updated knowledge of the unknown quantity  $\theta$  after observing the data. It is formed by taking the product of the likelihood function and the prior distribution's probability density function (both functions of  $\theta$ ) and then normalizing so that the result is a probability density function of  $\theta$ . This process is formally expressed by Bayes' formula, which is the foundation of Bayesian statistics.

**Definition 4.10 (Bayes' formula).** *Let  $x$  be an outcome of an absolutely continuous random variable  $X; \theta$  that is parameterized by  $\theta$ . Given the **prior distribution** on  $\theta$  with probability density function  $f(\theta)$  the **posterior distribution** on  $\theta$  is given by*

$$\underbrace{f(\theta|x)}_{\text{Posterior}} = \frac{\underbrace{f(\theta)}_{\text{Prior}} \underbrace{L(x; \theta)}_{\text{Likelihood}}}{\int_{\theta} f(\theta) L(x; \theta) d\theta} \quad (4.4)$$

**Remark 4.11.** *For a series of independent observations  $y_1, \dots, y_N$ , Bayes' formula can be applied with the joint likelihood of the vector of observations or sequentially, using the likelihood of each observation and the posterior from the previous observation(s) as a prior.*

The posterior distribution is a compromise between the prior and the likelihood. Values of  $\theta$  for which the observed data is more likely are assigned a higher value in the likelihood function. Similarly, values of  $\theta$  which are initially believed to be more likely are assigned a higher value in the prior distribution. By taking the product of the two, values of  $\theta$  that are relatively [un]likely in both are relatively [un]likely in the posterior and when the two disagree, the result is a compromise. This connects to the previously mentioned vague prior; in a vague prior, the values of  $\theta$  are relatively equally likely, so the posterior will largely be a reflection of the likelihood function.

The specification of the prior is a critical task, and is, to a degree, subjective. In the toy example, we suppose that the a priori information we have about the

room temperature  $\theta$  is minimal, and we wish to select a prior distribution reflecting that. Here we reach a conundrum: a complete lack of a priori information would imply that all values in  $\mathbb{R}$  are equally likely, however, a uniform distribution on  $\mathbb{R}$  is impossible<sup>2</sup>. Further, due to the property of *prior conjugacy*, which we will explain in detail shortly, a normal distribution is a desirable choice for the prior when the observations are modeled with normal distributions.

We will use a common method, the so-called *vague* prior approach [20] pp. 56-57 in combination with the prior conjugate choice of a normal distribution<sup>3</sup>. Roughly speaking, a vague prior is a distribution that has a negligible effect on the posterior distribution. For a normal distribution, this takes the form of extremely large variance, in comparison to the typical values the parameter could take. In the toy example, we would be quite surprised to see temperatures outside a reasonable range, even a range like  $20^\circ F$  to  $120^\circ F$  stretches the bounds of credibility. To express such a vague prior with a normal distribution, we pick a mean somewhere in that range, say  $70^\circ F$ , and a variance a few orders of magnitude larger than the range, say  $10^4$ . In practice, the exact values are unimportant, since a sufficiently large variance results in a distribution that is roughly uniform within the reasonable range, and, as a result, a posterior distribution that is primarily influenced by the observed data.

We now explain why a normal distribution is desirable in this example. The choice of prior distribution is closely related to the practical implementation of obtaining the posterior distribution. Bayes' formula holds for any pair of a likelihood function and prior distribution. However, for a given likelihood function, there is often a

---

<sup>2</sup>A uniform probability density function on  $\mathbb{R}$  must be constant, however the integral of such a function over  $\mathbb{R}$  can either be  $\infty$  or 0, never 1. That is, due to lack of compactness, a uniform distribution on  $\mathbb{R}$  does not exist.

<sup>3</sup>In this example, we use a fixed prior variance for simplicity. However, the prior conjugacy framework can model the variance of the normal prior distribution as a random variable with a gamma distribution, a concept explored in Sec. 4.3.2.

corresponding parameterized distribution that, when selected as a prior with specific parameters, results in the same parameterized distribution for the posterior, simply with different specific parameters<sup>4</sup>. This is called *prior conjugacy*. In this case, the process of finding the posterior distribution is reduced to determining how the hyperparameters change between the prior and posterior in response to the observed data. Without prior conjugacy, numerical methods like Monte Carlo Markov chains are typically used to approximate the posterior distribution [20] ch. 6. A prior conjugate model drastically reduces the computation load and avoids the introduction of error from numerical methods.

In our toy example, we assumed that each  $Y_n$  is normally distributed. If we also select a normal distribution for the prior on the unknown temperature  $\theta$ , then the likelihood function and prior distribution are prior conjugate. That is, the posterior distribution resulting from Bayes' formula in Eq. 4.4 is normally distributed. Further, the posterior mean and variance can be quickly calculated from the prior mean and variance, the observed values  $y_1, \dots, y_N$ , and the observation error variance  $V$ , all of which are known. Further, the posterior mean can be expressed as a weighted average of the prior mean and the observed data, with a weight  $K$  that depends on the relative uncertainties of the observed data and the prior mean, expressed through the observation error variance  $V$  and the prior variance, respectively. The posterior distribution of a normally distributed observation using a normally distributed prior is given explicitly in the proposition below, and a proof can be found in [20] pp. 59-60.

**Proposition 4.12.** *Let  $y$  be an outcome of a normally distributed random variable*

---

<sup>4</sup>These specific parameter values are typically called *hyperparameters* to distinguish them from unknown parameters in the context of Bayesian inference.

$Y$  parameterized by its unknown mean  $\theta$  and known variance  $V$

$$Y; \theta \sim \mathcal{N}(\theta, V)$$

Suppose the prior distribution for the parameter  $\theta$  is normal with mean  $\hat{\theta}$  and variance  $R$ :

$$\theta \sim \mathcal{N}(\hat{\theta}, R)$$

The posterior distribution defined by Eq. 4.4 is also distributed normally

$$\theta|y \sim \mathcal{N}(\bar{\theta}, C)$$

with the posterior mean  $\bar{\theta}$  and variance  $C$  given by

$$\bar{\theta} = \hat{\theta} + K(y - \hat{\theta}) = (1 - K)\hat{\theta} + Ky \tag{4.5}$$

$$C = (1 - K)R \tag{4.6}$$

where  $K$  is given by

$$K = \frac{R}{R + V}.$$

The value  $K$  is a quantity in  $(0, 1)$  that indicates the relative weight to assign to the observation  $y$  over the prior mean  $\hat{\theta}$ . This value depends on the relative uncertainties of the prior (expressed through the prior variance  $R$ ) and the likelihood (expressed through the observation error variance  $V$ ). As the ratio  $V/R$  approaches  $\infty$ ,  $K \rightarrow 0$ . The interpretation is that when the observation error variance is significantly larger than the prior variance, there is correspondingly more trust in the prior, and less weight is assigned to the observation  $y$  when calculating the posterior mean in Eq. 4.5. On the other hand, as the ratio  $R/V$  approaches  $\infty$ ,  $K \rightarrow 1$ . In this case,

significantly more trust in the data corresponds to more weight assigned to  $y$  in Eq. 4.5.

Further, note that since  $K$  is always between zero and one, the posterior variance  $C$  is always less than the prior variance  $R$ . This can be seen in Eq. 4.6, since the quantity  $1 - K$  is also in  $(0, 1)$ . The basic interpretation is that incorporating data reduces uncertainty in the Bayesian framework.

In the toy example, we choose a normal distribution for the prior on  $\theta$  to illustrate the prior conjugacy in Prop. 4.12. This prior distribution is specified by the mean  $\hat{\theta}$  and variance  $R$ . These are called hyperparameters to distinguish them from  $\theta$ , the unknown parameter of interest. The hyperparameters are fixed, chosen values that are selected to represent prior knowledge of  $\theta$ . In our example, we use a vague prior with large variance  $R$  to model a lack of prior knowledge. Recall that if  $R/V \gg 1$  the prior mean has little effect on the posterior mean in Eq. 4.5.

To make the Bayesian process concrete, we return to our example with specific numbers, as depicted in Fig. 4.2. Suppose we are provided with the measurement model described up to this point, with an observation error variance of 4 (i.e., an observation error standard deviation of  $2^\circ F$ ). That is, we assume that the measurements are independently normally distributed around the (unknown to us) room temperature  $\theta$  with a variance of  $V = 4$ . In fact, the 50 measurements depicted in Fig. 4.2 were simulated using this exact model; they are independent random samples from a normal distribution around a true temperature of  $\theta = 75^\circ F$ .

With the statistical model and prior specified above, we perform the Bayesian inference from Prop. 4.12. To do so, we select a vague prior. As already mentioned, in this case, the prior mean has minimal influence on the posterior mean. We consider the following prior,  $\theta \sim (70, 10^4)$ . The mean of  $70^\circ F$  was chosen simply as a reasonable value for the temperature of a room, while the variance  $10^4$  results in a sufficiently

vague prior, as it is a few orders of magnitude larger than typical temperature values.

The resulting distributions of sequentially applying Bayes' formula to the measurements are depicted in Fig. 4.3 with the mean highlighted and two standard deviations shaded. We can see that despite the initial mean of  $70^\circ F$  being fairly different from the true room temperature of  $75^\circ F$ , the posterior means quickly approaches the room temperature as more data is incorporated. By the final measurement, the posterior mean is  $74.92^\circ F$ , with a fairly small standard deviation of  $0.28^\circ F$ .




Figure 4.3: Posterior mean when sequentially using Normal conjugate prior inference to assimilate the values in Fig. 4.2, along with two standard deviations of the prior distribution shaded.

## 4.2 Stochastic state space models and the Kalman filter

In the previous section, we considered a scenario where a series of temperature measurements, depicted in Fig. 4.2, were taken over a time interval during which the underlying room temperature was constant. Now we will consider a more complex scenario, in which the room temperature changes over time.

The measurements for this are depicted in Fig. 4.4. In addition to the observation errors, we see an increasing trend in the room temperature. This roughly linear trend is one example of *dynamic behavior*, an observable pattern of behavior in a time series.

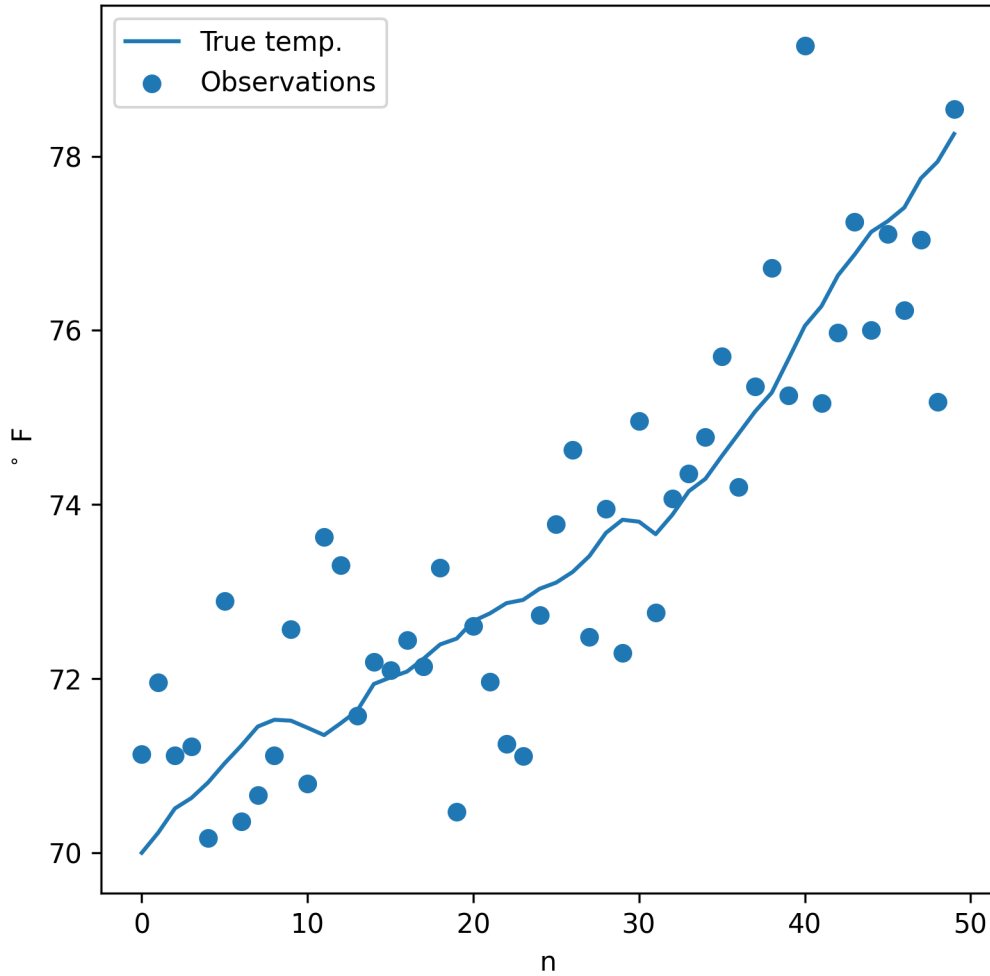


Figure 4.4: 50 noisy temperature measurements over a dynamically changing room temperature.

Later, we will also consider oscillations, shown in Fig. 4.1, a different type of dynamic behavior. For additional types of dynamic behavior that are amenable to the class of models described in this section, see [27].

As before, denote the temperature measurements by  $y_1, \dots, y_N$ , where the unknown temperature at the time of the  $n$ -th measurement is  $\theta_n$ . Like in the previous section,

we will use a statistical model that assigns a parameterized distribution to each measurement. The primary difference, as seen in Fig. 4.4, is that the parameters,  $\theta_1, \dots, \theta_N$ , are different at each measurement time. Consequently, we can no longer use a single unknown  $\theta$ . Instead, when performing Bayesian inference, we will model each unknown temperature in the sequence  $\theta_1, \dots, \theta_N$  as a distinct random variable.

When this is the case, one may wonder how well Bayesian inference can work, since for each unknown temperature  $\theta_n$ , we only have a single data point  $y_n$ . The key is that while  $\theta_1, \dots, \theta_N$  are no longer the same, they may not be completely independent of each other. Suppose we have a *forecast model* that describes how  $\theta_n$  depends on  $\theta_{n-1}$ . In that case, we can use the posterior distribution of  $\theta_{n-1}$  along with the forecast model to obtain a prior distribution on  $\theta_n$ .

We will consider the case where the forecast model has two elements: a deterministic element and a stochastic element. The deterministic element consists of a *forecast function* that returns a predicted value of  $\theta_n$  based on the value of  $\theta_{n-1}$ . The stochastic element describes the *model error*, the difference between this predicted value and the actual value of  $\theta_n$ , by assigning it a probability distribution.

The full statistical model that describes the sequence of measurements will have two components: a forecast model mentioned above, and an *observation model*, a sequence of parameterized distribution that describes the observation error for each  $y_n$ , see 4.1. We have stated that in our forecast model, the variables  $\theta_n$  each depend (stochastically) only on  $\theta_{n-1}$  and are independent of  $\theta_{n-2}, \theta_{n-3}, \dots, \theta_0$ . Similarly, we will only consider observation models where each  $y_n$  depends only on  $\theta_n$ . This notion of dependence can be formalized as *conditional independence*, defined below in terms of conditional probability.

**Definition 4.13.** *Let  $X_1, \dots, X_n, Y$ , and  $Z$  be absolutely continuous random variables.*



If

$$f_{X_1, \dots, X_n, Y|Z}(x_1, \dots, x_n, y|z) = f_{X_1, \dots, X_n|Z}(x_1, \dots, x_n|z) \cdot f_{Y|Z}(y|z)$$

for all  $x_1, \dots, x_n, y, z$ , then we say  $Y$  is **conditionally independent** of  $X_1, \dots, X_n$  given  $Z$ , or alternatively,  $Y$  is **conditionally dependent** on  $Z$  (with respect to the random variables  $X_1, \dots, X_n$ ).




Figure 4.5: Conditional dependency structure of state-space models, where an arrow indicates the target is conditionally independent of the other random variables, given the source. That is, each  $\theta_n$  is conditionally dependent only on  $\theta_{n-1}$  and each  $Y_n$  is conditionally dependent only on  $\theta_n$ .

In the stochastic state space models that we consider, each true temperature  $\theta_n$  is modeled as conditionally dependent on  $\theta_{n-1}$  and each measurement  $Y_n$  is conditionally dependent on  $\theta_n$ . Visually, this dependence is depicted in Fig. 4.5, where an arrow indicates conditional dependence. The following definition formalizes the concept of a stochastic state space model.

**Definition 4.14.** Let  $\{\theta_0, \dots, \theta_N\}$  and  $\{Y_1, \dots, Y_N\}$  be two collections of random variables where each  $\theta_n$  is  $d$ -dimensional<sup>5</sup>. We say that these two collections form a **stochastic state space model** when

<sup>5</sup>While the theory extends to multivariate  $Y_n$ , in our context we will only consider univariate  $Y_n$ .

1. Each  $\boldsymbol{\theta}_n$  is conditionally independent of the other random variables in the sequence, given  $\boldsymbol{\theta}_{n-1}$  and
2. Each  $Y_n$  is conditionally independent of the other random variables in the sequence, given  $\boldsymbol{\theta}_n$ .

The conditional distributions of  $\boldsymbol{\theta}_n|\boldsymbol{\theta}_{n-1}$  for each  $n = 1, \dots, N$  are called the **forecast** distributions, the conditional distributions of  $Y_n|\boldsymbol{\theta}_n$  for each  $n = 1, \dots, N$  are called the **observation** distributions, and the distribution of  $\boldsymbol{\theta}_0$  is called the **initial** distribution. For shorthand, we will typically denote such a stochastic state space model by  $\{\boldsymbol{\theta}_n, Y_n\}$ .

In a stochastic state space model, the sequence of random variables  $\{\boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}_n\}$  is called the state. In many engineering applications,  $\boldsymbol{\theta}_n$  does indeed represent the state of a physical system that produces a measured output  $y_n$ . For example, in meteorology, the state may be the vector consisting of the temperature at a grid of locations on Earth, while the observation is the temperature measured at a particular location [12]. Or, in robotic movement, the state may be the position and velocity of the robot, while the observation is a noisy measurement of the current location [32]. However, in the context of signal processing, the components of the vector  $\boldsymbol{\theta}_n$  model individual dynamic behaviors present in a signal. We will use stochastic state space models in this manner to model the measures on our granular system in the stick regime, as depicted in Fig. 4.1.

In our toy example, mentioned above, we will consider the following model. The state  $\boldsymbol{\theta}_n$  consists of both the unobserved temperature value  $\theta_n^0$  and its time derivative<sup>6</sup>  $\theta_n^1$ . The stochastic element captures model error in the deterministic forecast by

---

<sup>6</sup>More precisely, the second component of the state is the derivative with respect to the time step. This normalization of the derivative, while not required, results in the simple update in Eq. 4.7, where the deterministic forecast temperature value is the sum of the previous value and the normalized time derivative.

adding a normally distributed random variable,  $\omega_n \sim \mathcal{N}_2(0, W)$ , to the deterministic forecast. The deterministic forecast itself consists of updating the temperature value through its time derivative and assuming the time derivative is unchanged.

$$\theta_n^0 | \boldsymbol{\theta}_{n-1} = \theta_{n-1}^0 + \theta_{n-1}^1 + \omega_n^0 \quad (4.7)$$

$$\theta_n^1 | \boldsymbol{\theta}_{n-1} = \theta_{n-1}^1 + \omega_n^1 \quad (4.8)$$

In a model where the full forecast is the sum of a deterministic forecast and normally distributed model error, the matrix  $W$  is called the *model error covariance* matrix. The model error covariance matrix captures the variance of both  $\omega_n^0$  and  $\omega_n^1$ , and their dependence. In this section, we will assume that the matrix  $W$  is known to us a priori, but in practice, the specification of  $W$  can be complicated. We will discuss one approach in Sec. 4.3.1.

The observation distributions are modeled as (conditionally) independent normally distributed random variables centered around the true (unknown) temperature with an error variance  $V$ :

$$Y_n | \boldsymbol{\theta}_n = \theta_n^0 + v_n \quad (4.9)$$

where  $v_n \sim \mathcal{N}(0, V)$ . As before, we will assume that  $V$  is known a priori.

In a stochastic state space model, Bayesian statistics can be applied to sequentially update the distribution of each current state  $\boldsymbol{\theta}_n$  in light of the most recent observation  $y_n$ , while also considering the previous updated state distribution of  $\boldsymbol{\theta}_{n-1}$ . The process of obtaining the distribution  $\boldsymbol{\theta}_n | y_{1:n}$ <sup>7</sup>, often called the current *filtering distribution*, from the previous filtering distribution  $\boldsymbol{\theta}_{n-1} | y_{1:n-1}$  and the observation  $y_n$  is called *sequential filtering*. Sequential filtering consists of two steps: forecasting and data assimilation. In the forecast step, the previous filtering distribution and the forecast

---

<sup>7</sup> $y_{1:n}$  is shorthand for the collection of values  $y_1, \dots, y_n$

distribution  $\boldsymbol{\theta}_n|\boldsymbol{\theta}_{n-1}$  of the model are combined to form the distribution  $\boldsymbol{\theta}_n|y_{1:n-1}$ , the *predictive state distribution*. Next, after observing the value  $y_n$ , the data assimilation step combines the likelihood function  $L(y_n|\boldsymbol{\theta}_n)$  of the model and the predictive state distribution using Bayes' formula. This results in the desired filtering distribution  $\boldsymbol{\theta}_n|y_{1:n}$ , and the process repeats with the next observation. This is made precise in the following proposition and corollary.

**Proposition 4.15.** *Suppose  $\{\boldsymbol{\theta}_n, Y_n\}$  is a stochastic state space model. If  $y_1, \dots, y_n$  are outcomes of the random variables  $Y_1, \dots, Y_n$ , respectively, then the following relations on the probability density functions hold:*

$$f(\boldsymbol{\theta}_n|y_{1:n-1}) = \int f(\boldsymbol{\theta}_n|\boldsymbol{\theta}_{n-1})f(\boldsymbol{\theta}_{n-1}|y_{1:n-1})d\boldsymbol{\theta}_{n-1} \quad (4.10)$$

and

$$f(\boldsymbol{\theta}_n|y_{1:n}) = \frac{f(\boldsymbol{\theta}_n|y_{1:n-1}) \cdot L(y_n|\boldsymbol{\theta}_n)}{\int f(\boldsymbol{\theta}_n|y_{1:n-1}) \cdot L(y_n|\boldsymbol{\theta}_n)d\boldsymbol{\theta}_n} \quad (4.11)$$

**Corollary 4.16.** *Given an initial density function  $f(\boldsymbol{\theta}_0)$ , the above equations can be recursively applied to the sequence of data  $y_1, \dots, y_n$  to obtain  $f(\boldsymbol{\theta}_1|y_1)$  from  $y_1$  and  $f(\boldsymbol{\theta}_0)$ ,  $f(\boldsymbol{\theta}_2|y_{1:2})$  from  $y_2$  and  $f(\boldsymbol{\theta}_1|y_1)$ , etc. This process is **sequential filtering** where the application of Eq. 4.10 is called the **forecast step** and the application of Eq. 4.11 is called the **data assimilation step**.*

For a general state space model, the resulting probability density functions cannot be computed analytically. Consequently, methods such as Monte Carlo algorithms are usually used to approximate the relevant distributions [27] pp. 207-229. A good example of such an implementation is the ensemble Kalman filter utilized in meteorology [12] or the particle filter for target tracking [32]. However, the computations

become much simpler if we consider a special class of stochastic state space models for which sequential filtering uses Bayesian conjugate prior distributions: the dynamic linear model, or DLM.

DLMs are a particular type of stochastic state space model where restrictions on the initial distribution, forecast distributions, and observation distributions allow for efficient sequential filtering. The restrictions require each forecast distribution  $\boldsymbol{\theta}_n|\boldsymbol{\theta}_{n-1}$  to be the sum of a linear function of  $\boldsymbol{\theta}_{n-1}$  and a normally distributed random variable and the initial distribution to be normal. Similarly, each observation distribution  $Y_n|\boldsymbol{\theta}_n$  must be the sum of a linear function of  $\boldsymbol{\theta}_n$  and a normally distributed random variable. These conditions are collectively referred to as Gaussian-linearity. Below is the precise definition of a DLM, explicitly stating the Gaussian-linearity assumptions mentioned above.

**Definition 4.17.** *Let  $\{\boldsymbol{\theta}_n, Y_n\}$  be a stochastic state space model with a  $d$ -dimensional state. We call this model a **dynamic linear model (DLM)** if it has the following dependency structure and initial state distribution. The distributions of  $\boldsymbol{\theta}_n|\boldsymbol{\theta}_{n-1}$  and  $Y_n|\boldsymbol{\theta}_n$  for each  $n$  are given by a pair of equations*

$$\begin{aligned} Y_n|\boldsymbol{\theta}_n &= F_n\boldsymbol{\theta}_n + v_n \\ \boldsymbol{\theta}_n|\boldsymbol{\theta}_{n-1} &= G_n\boldsymbol{\theta}_{n-1} + \boldsymbol{w}_n \end{aligned} \tag{4.12}$$

where  $\{G_n\}$  and  $\{F_n\}$  are known sequences of matrices of order  $1 \times d$  and  $d \times d$  respectively and  $\{v_n\}$  and  $\{\boldsymbol{w}_n\}$  are two independent sequences of normal random variables distributed as  $v_n \sim \mathcal{N}_1(0, V_n)$  and  $\boldsymbol{w}_n \sim \mathcal{N}_d(\mathbf{0}, W_n)$  where the values of the sequence of variances  $\{V_n\}$  and covariance matrices  $\{W_n\}$  are given. The initial state

distribution of  $\boldsymbol{\theta}_0$  is given by a normal  $d$ -dimensional distribution,

$$\boldsymbol{\theta}_0 \sim \mathcal{N}_d(\mathbf{m}_0, C_0) \tag{4.13}$$

for a given vector  $\mathbf{m}_0$  and covariance matrix  $C_0$ .

In the model given by Eqns. 4.7, 4.8, and 4.9, if we further assume that the initial state distribution is normally distributed, the model is a DLM. In this model, the forecast and observation functions are time-independent. While this is not always the case in a DLM, it is fairly common. To indicate the time independence of the forecast and observation functions, we will drop the subscripts when giving the matrix representations of  $G_n$  and  $F_n$ , a convention that we will use when defining other DLMS later. The forecast given by Eqns. 4.7, 4.8 can be represented by the matrix

$$G = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

and the observation function given by Eq. 4.9 can be represented by the matrix

$$F = \begin{bmatrix} 1 & 0 \end{bmatrix}$$

The aforementioned efficiency of sequential filtering of DLMS comes from the fact that *all distributions resulting from the sequential filtering of a DLM are normally distributed*. In the forecast step, this follows from the fact that the linear transformation of a normal random variable is again normal, as is the sum of two normal random variables [20]. In the data assimilation step, the normality of the filtering distribution results from the prior conjugacy of Proposition 4.12. Consequently, sequential filtering of a DLM consists of updating the mean and covariance of the state

distribution, first in the forecast step, and then in the data assimilation step. The relevant formulas describing the computations that track these values are collectively called the Kalman filter. We summarize the Kalman filter in the following proposition from [27].

**Proposition 4.18 (Kalman Filter).** *Consider the DLM specified by (4.12) and (4.13). Let*

$$\boldsymbol{\theta}_{n-1}|Y_{1:n-1} \sim \mathcal{N}_d(\mathbf{m}_{n-1}, C_{n-1}).$$

*Then the following statements hold.*

1. *The predictive state distribution is Gaussian*

$$\boldsymbol{\theta}_n|Y_{1:n-1} \sim \mathcal{N}_d(\mathbf{a}_n, R_n).$$

*where the parameters are given by*

$$\begin{aligned} \mathbf{a}_n &= E(\boldsymbol{\theta}_n|Y_{1:n-1}) = G_n \mathbf{m}_{n-1} \\ R_n &= \text{Var}(\boldsymbol{\theta}_n|Y_{1:n-1}) = G_n C_{n-1} G_n' + W_n. \end{aligned} \tag{4.14}$$

*$G_n'$  denotes the transpose of  $G_n$ .*

2. *The predictive observation distribution is Gaussian*

$$Y_n|Y_{1:n-1} \sim \mathcal{N}_1(f_n, Q_n)$$

*where the parameters are given by*

$$\begin{aligned} f_n &= E(Y_n|Y_{1:n-1}) = F_n \mathbf{a}_n \\ Q_n &= \text{Var}(Y_n|Y_{1:n-1}) = F_n R_n F_n' + V_n. \end{aligned} \tag{4.15}$$

3. *The filtering distribution is Gaussian*

$$\boldsymbol{\theta}_n | Y_{1:n} \sim \mathcal{N}_d(\mathbf{m}_n, C_n)$$

where the parameters are given by

$$\begin{aligned} \mathbf{m}_n &= E(\boldsymbol{\theta}_n | Y_{1:n}) = \mathbf{a}_n + K_n e'_n \\ C_n &= \text{Var}(\boldsymbol{\theta}_n | Y_{1:n}) = (I - K_n F_n) R_n. \end{aligned} \tag{4.16}$$

where  $e'_n = y_n - f_n$  is the forecast error and the  $d \times 1$  matrix  $K_n$  is the Kalman gain matrix given by

$$K_n = R_n F_n' Q_n^{-1}. \tag{4.17}$$

In summary, the forecast step of sequential filtering for a DLM starts with the state mean at the previous time,  $\mathbf{m}_{n-1}$ , obtains the state mean at the current time  $\mathbf{a}_n$ , and then obtains the predicted observation mean  $f_n$ . Next, the data assimilation step obtains an updated estimate of the current state mean  $\mathbf{m}_n$  from the predictive error, the difference between the observed value  $y_n$  and  $f_n$ .

The data assimilation step uses the Kalman gain matrix  $K_n$ , which depends on the predictive state and observation covariances,  $R_n$  and  $Q_n$ . Analogous to the weight  $K$  in Prop. 4.12, each entry in the Kalman gain matrix is the weight assigned to the forecast error  $e'_n$  in order to correct the corresponding component in the predictive state mean  $\mathbf{a}_n$ . As in that proposition, a larger predictive observation error variance  $Q_n$  will lead to less weight placed on the forecast error while a larger predictive observation error variance  $R_n$  leads to greater weight.

The forecast error  $e'_n$  is an important quantity, as it is a measure of the deviation of the model from the observed values. Note that since each observation  $y_n$  is an outcome



of the observation random variable  $Y_n|Y_{1:n-1} \sim \mathcal{N}(f_n, Q_n)$ , each forecast error  $e'_n$  is an independent outcome of a random variable distributed  $\mathcal{N}(0, Q_n)$ . Further, if we define the normalized model error (NME) by

$$e_n = \frac{y_n - f_n}{\sqrt{Q_n}} \quad (4.18)$$

then each NME value is an independent outcome of a random variable distributed  $\mathcal{N}(0, 1)$ , the standard normal distribution. As a result, a posteriori testing of the assumption that the observations follow a specified stochastic state space model is typically reduced to testing that the NME values are independent and normally distributed [27].

To make the Kalman filter concrete, we return to our example of predicting temperature from the observations in Fig. 4.4. For simplicity, we will consider a DLM where the model error covariance matrix  $W_n$  is time-independent, as is the observation error variance  $V_n$ . Further, we will assume that the model error for the temperature and its derivative are independent, with a standard deviation of  $0.1^\circ F$  and  $0.02^\circ F$ , respectively. These relatively small values on the scale of degrees Fahrenheit depict a model with a comparatively small deviation of the true temperatures from a straight line (i.e., the deterministic forecast). In matrix form, we have

$$W = \begin{bmatrix} 0.1^2 & 0 \\ 0 & 0.02^2 \end{bmatrix}$$

In contrast, we will assume the observation error variance has a standard deviation of  $1^\circ F$ . The larger magnitude of the observation error models observed temperature values normally distributed around the relatively straight line implied by the smaller model error values. This is exactly the behavior of the observations in Fig. 4.4.

Indeed, those values were simulated using this exact model.

With this DLM, we can apply the sequential filtering process of the Kalman filter, starting with an initial distribution on  $\boldsymbol{\theta}_0$ . As in standard Bayesian filtering, if we select a vague initial distribution, the initial mean will have minimal influence on the filtered values after a sufficient, typically small, number of observations. This initial period is commonly called a burn-in period [27], and for a vague initialization, it is characterized by filtered values that are nearly identical to the observed values. As in the previous section, the initial temperature mean is chosen to be  $70.0^\circ F$ , simply as a reasonable value. The initial derivative mean is chosen to be zero to indicate no knowledge of the time derivative. The covariance matrix has entries that are sufficiently large for a vague initialization, on the same scale as in the static example. In summary, the initial distribution is  $\boldsymbol{\theta}_0 \sim \mathcal{N}_2(\mathbf{m}_0, C_0)$  where

$$\mathbf{m}_0 = \begin{bmatrix} 70 \\ 0 \end{bmatrix} \text{ and } C_0 = \begin{bmatrix} 10^4 & 0 \\ 0 & 10^4 \end{bmatrix}.$$

The resulting filtered distributions are shown in Fig. 4.6 with the filtered mean temperature highlighted and two standard deviations shaded<sup>8</sup>.

The inset panel shows the resulting NME defined by Eq. 4.18. Recall that if the observations are outcomes of the specified DLM, then the NME values are each distributed  $\mathcal{N}(0, 1)$  and are mutually independent. This does appear to be the case in the depicted NME. In contrast, when we later construct a DLM to model the stick regime behavior of a measure on the granular system, we expect that the constructed DLM will not perfectly model the true signal. The degree to which this is the case will be indicative of the appropriateness of the constructed DLM model.

---

<sup>8</sup>The initial state distribution is not depicted, as it is not a filtered distribution.

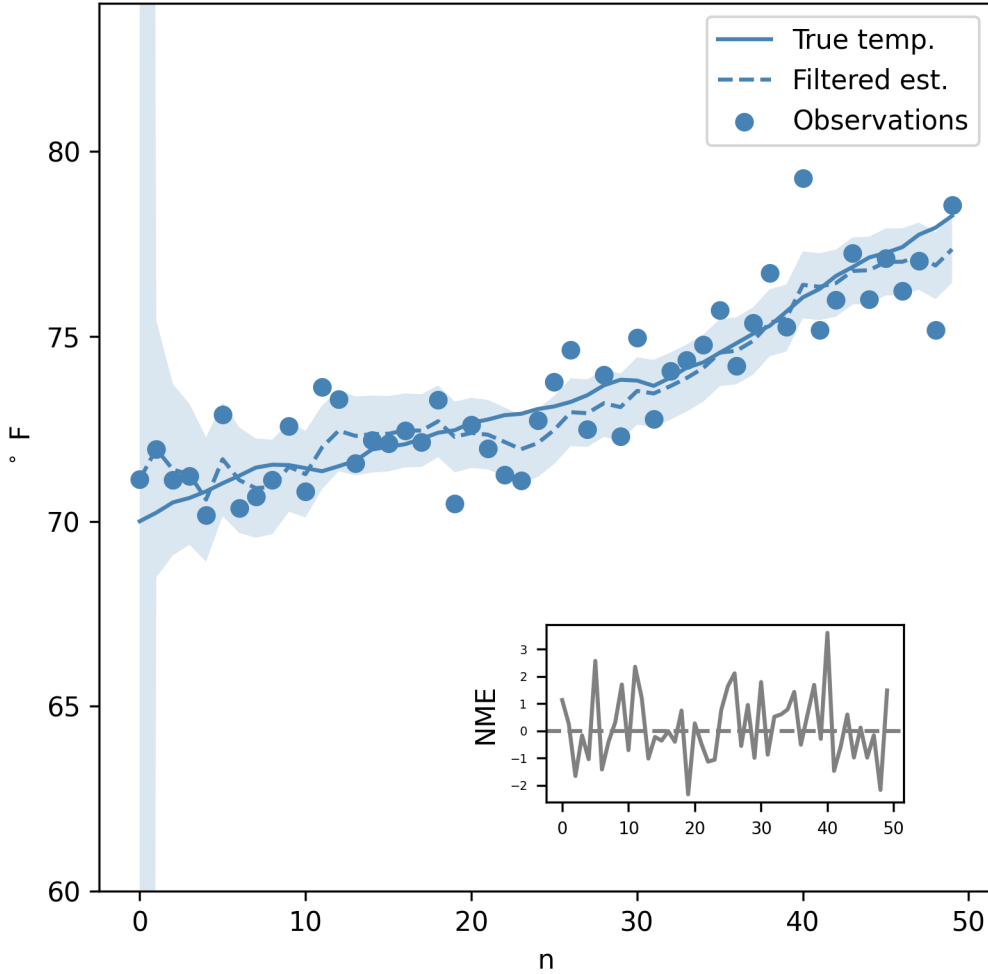


Figure 4.6: The Kalman filter applied to a series of simulated data points with a stochastically varying trend. The 50 data points are shown as points, with the true temperature marked by the solid blue curve and the filtered estimated temperatures given by the dashed blue curve. The shaded region marks the sequence of  $2 \times$  standard deviation credible intervals for the filtered temperature estimates. The inset shows the normalized model error from Equation 4.18.

Note that, unlike the static example in Fig. 4.3, after the initial tuning, the standard deviation is relatively constant. Intuitively, the asymptotic size of the filtered state error is a compromise between the error introduced by the stochastic forecast

at each time and the error reduction that occurs when Bayes' formula incorporates the data at that time. While in this dynamic example, this compromise results in the apparent convergence of the state variance, one can imagine a scenario in which sufficiently large model error variance causes the state variance to grow without bound.

As a simple exercise, consider the dynamic example above, but suppose the observation function maps the state to 0:

$$F = \begin{bmatrix} 0 & 0 \end{bmatrix}$$

Intuitively, this results in the observations giving no information on the state. Formally, the application of the Kalman filter in this case can be performed analytically, resulting in the entries of  $C_n$  growing without bound as  $n \rightarrow \infty$ .

For a time-invariant DLM, convergence of the state variance is called *filter stability*. More precisely a DLM is stable in this sense if there exists a limiting state covariance matrix  $C_\infty$  such that  $C_n \rightarrow C_\infty$  as  $n \rightarrow \infty$ . The existence of this limit is equivalent to the existence of a solution of the Riccati equation [27] pp. 80-82. In practice, we typically assume that a filter is stable if the state covariance appears to converge (as is the case here).

### 4.3 Specification of error structure

To specify the DLM one needs to specify the model error covariance matrices  $\{W_n\}$  and observation error variances  $\{V_n\}$ . In our toy example, these were assumed to be known a priori. In practice, minimal information is known about the model error and observation error, which makes specifying these terms in the DLM challenging. In this section, we describe the approaches we use in Ch. 5 to model these two types of error.

### 4.3.1 Discount factor specification of model error covariance matrices

First, consider the matrix  $W_n$ , which is the covariance of the model error arising from the forecast step, Equation 4.14. Without a deep understanding of the structure of the model error, directly specifying this matrix is difficult. The method we use increases the state covariance in the forecast step in the simplest manner: multiplying the forecast state covariance matrix by a scaling factor.

More precisely, we dynamically specify  $W_n$  based only on a single parameter  $\delta \in (0, 1]$ , referred to as the *discount factor*, by

$$W_n = \frac{1 - \delta}{\delta} P_n \quad (4.19)$$

where  $P_n$  refers to the raw forecast state covariance given by  $P_n = G'_n C_{n-1} G_n$  in Eq. 4.19, where the term ‘raw’ is used to indicate that  $P_n$  is the state covariance resulting only from the deterministic component of the forecast. Recall that this equation states that full forecast state covariance,  $R_n$ , is the sum of this term,  $P_n$ , and the stochastic component  $W_n$  that captures model error:

$$R_n = P_n + W_n.$$

Consequently, when specifying  $W_n$  by Eq. 4.19, the resulting forecast state covariance is simply

$$R_n = \frac{1}{\delta} P_n.$$

In summary, the discount factor  $\delta$  specifies  $W_n$  in such a way that its reciprocal  $\frac{1}{\delta}$  inflates the uncertainty on the unknown state. The discount factor is commonly interpreted as a measure of model trust [27]. This interpretation results from the fact

that  $\delta = 1$  corresponds to no increase in state covariance in the forecast step while  $\delta \ll 1$  yields an arbitrarily large state covariance after forecasting.

By applying this approach, we no longer need to specify the  $d^2$  entries of  $W_n$ , which requires a detailed understanding of the error structure of the model. Instead, we set the single value  $\delta$ . Of course, directly specifying the matrix  $W_n$  is preferable when we have specific knowledge of the model's error structure, but since we do not have such information, we use the discount factor methodology. We will discuss how this value is chosen (as well as the specification of other parts of the DLM) in Ch. 5.

### 4.3.2 Conjugate-variance Kalman filter for observation error variance

Now that we have discussed model error, let us turn our attention to the observation error variance,  $V_n$ . As with the model error, we do not have an a priori understanding as to the scale of this error. We have already encountered this type of uncertainty when introducing the time-invariant thermometer example in section 4.1. We mentioned that one common approach is to use a Bayesian model for the observation error variance, and that is the approach we describe here.

To incorporate uncertainty on the observation error term, we treat it as an unknown, in the Bayesian sense. We model the observation error variance as a time-independent random variable  $V$  that is considered to be part of the state. When we do so, the resulting model is not a DLM since  $V$  does not interact linearly with the rest of the state. Consequently, the Kalman filter does not apply.

However, we can still efficiently conduct sequential filtering by applying an extension to the Kalman filter. For prior conjugacy to apply,  $V$  needs to have an inverse-gamma distribution. Equivalently, the random variable  $V^{-1}$ , which we will denote by  $\phi$ , needs to have a gamma distribution. The variable  $\phi$  is typically called the *precision*. Further, the non-linear dependence of the rest of the state on  $\phi$  is

captured through the *joint normal-gamma distribution*, according to the definition below.

**Definition 4.19.** *Let  $(\boldsymbol{\theta}, \phi)$  be a multivariate random variable. That is, the components of the vector  $\boldsymbol{\theta}$  and the value  $\phi$  are each random variables. We say the  $(\boldsymbol{\theta}, \phi)$  is distributed joint normal-gamma with parameters  $(\mathbf{m}, \bar{C}, \alpha, \beta)$  if  $\phi \sim \mathcal{G}(\alpha, \beta)$  (i.e., distributed gamma) and, conditional on  $\phi$ ,  $\boldsymbol{\theta}|\phi \sim \mathcal{N}_d(\mathbf{m}, \bar{C}\phi^{-1})$ . We denote the joint normal-gamma distribution by*

$$(\boldsymbol{\theta}, \phi) \sim \mathcal{NG}(\mathbf{m}, \bar{C}, \alpha, \beta).$$

This distribution is most easily interpreted with a univariate state, in the case where  $C = 1$ . In that case, the joint normal-gamma distribution simply states that, conditional on  $\phi$ , the state  $\theta$  is distributed normally with a variance of  $\phi^{-1}$ . In the more general case, the precision  $\phi$  plays the role of a scaling factor on the error covariance of  $\boldsymbol{\theta}$ . As the precision increases, the error covariance is scaled down. For example, if  $\bar{C} = I$ , then  $\phi^{-1}$  is exactly the variance of each component of  $\boldsymbol{\theta}$ .

Modeling the observation error precision  $\phi = V^{-1}$  as a time-invariant random variable results in a new random variable in addition to the time-varying state random variables  $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N$ . As with the state,  $\phi$  is estimated from the data at each time via Bayes' formula. We utilize the joint multivariate Normal-Gamma distribution described by Def. 4.19 to express the uncertainty around both the state  $\boldsymbol{\theta}_n$  and the unknown observation error precision  $\phi$ , which follows a gamma distribution.

Due to prior conjugacy, this choice of distribution allows for an exact sequential filtering algorithm similar to the Kalman filter for standard DLMS. The primary difference is that the distributions are no longer normal, but joint normal-gamma for the state distributions and student-t for the observation distributions. Further,

all error covariance and variance terms are scaled by the random variable  $\phi$ , which, similar to the above definition, is distributed gamma with parameters  $\alpha_n, \beta_n$ . Those hyperparameters are adjusted as the filter ‘learns’ the distribution of the precision  $\phi$  through the prior conjugate application of Bayes’ formula. We refer to the resulting collection of equations as the conjugate-variance Kalman filter [27] pp. 158-160.

**Proposition 4.20 (Conjugate-Variance Kalman Filter).** *Let the conjugate-variance DLM be specified with a joint normal-gamma prior distribution*

$$\boldsymbol{\theta}_0, V^{-1} \sim \mathcal{NG}(\mathbf{m}_0, \bar{C}_0, \alpha_0, \beta_0)$$

and the dependency structure (4.12), conditional on the unknown  $V$  in the following manner

$$v_n | V \sim \mathcal{N}_1(0, V)$$

$$\mathbf{w}_n | V \sim \mathcal{N}_d(0, V\bar{W}_n)$$

where  $\bar{W}_n$  is specified. Then the following statements hold.

1. The one-step-ahead predictive state distribution is joint normal-gamma

$$\boldsymbol{\theta}_n, V^{-1} | Y_{1:n-1} \sim \mathcal{NG}(\mathbf{a}_n, \bar{R}_n, \alpha_{n-1}, \beta_{n-1}).$$

where the parameters are given by

$$\begin{aligned} \mathbf{a}_n &= G_n \mathbf{m}_{n-1} \\ \bar{R}_n &= G_n \bar{C}_{n-1} G_n' + \bar{W}_n. \end{aligned} \tag{4.20}$$



2. The one-step-ahead predictive observation distribution is scaled student-t

$$Y_n | Y_{1:n-1} \sim \mathcal{T} \left( f_n, \bar{Q}_n \frac{\beta_{n-1}}{\alpha_{n-1} - 1} \right)$$

with  $2\alpha_{n-1}$  degrees of freedom where the parameters are given by

$$\begin{aligned} f_n &= F_n \mathbf{a}_n \\ \bar{Q}_n &= F_n \bar{R}_n F_n' + 1. \end{aligned} \tag{4.21}$$

3. The filtering distribution is joint Normal-Gamma

$$\boldsymbol{\theta}_n, V^{-1} | Y_{1:n} \sim \mathcal{NG}(\mathbf{m}_n, \bar{C}_n, \alpha_n, \beta_n)$$

where the parameters are given by

$$\begin{aligned} \mathbf{m}_n &= \mathbf{a}_n + K_n e_n' \\ \bar{C}_n &= (I - K_n F_n) \bar{R}_n \\ \alpha_n &= \alpha_{n-1} + \frac{1}{2} \\ \beta_n &= \beta_{n-1} + \frac{1}{2} \bar{Q}_n^{-1} e_n^2 \end{aligned} \tag{4.22}$$

where  $e_n' = y_n - f_n$  is the forecast error and the  $n \times 1$  matrix  $K_n$  is the Kalman gain matrix given by

$$K_n = \bar{R}_n F_n' \bar{Q}_n^{-1}. \tag{4.23}$$

**Remark 4.21.** For a conjugate-variance DLM as described above, the normalized

model error (NME) is given by

$$e'_n = \frac{y_n - f_n}{\sqrt{\bar{Q}_n \frac{\beta_{n-1}}{\alpha_{n-1} - 1}}} \quad (4.24)$$

and follows a standard Student- $t$  distribution with  $2\alpha_{n-1}$  degrees of freedom. Note that for large  $\alpha_{n-1}$  (i.e., large  $n$ ) this converges to the standard normal distribution.

The conjugate-variance DLM still allows for discount factor specification, this time of  $\bar{W}_n$ . In this model, if we let  $\bar{P}_n = G_n \bar{C}_n G'_n$  then  $\bar{R}_n = \bar{P}_n + \bar{W}_n$ . As before, when we express  $\bar{W}_n$  as a fraction of  $\bar{P}_n$

$$\bar{W}_n = \frac{1 - \delta}{\delta} \bar{P}_n$$

then  $\bar{R}_n = \frac{1}{\delta} \bar{P}_n$ . The discount factor  $\delta$  has the same interpretation as before.

#### 4.4 Extended Kalman filter (EKF)

In some cases, the forecast function  $G_n$  is non-linear. Since such a stochastic state space model is not a DLM, the Kalman filter can no longer be used for sequential filtering. A field where this frequently occurs is meteorology, where non-linear dynamics are common [12]. In our case, the model we will use in Ch. 5 to describe the signal of the  $W_2$  measure during stick has a non-linear forecast function. To allow for non-linear deterministic evolution, we can consider applying the Kalman filter equations using the linearization of the forecast function  $G_n(\cdot)$  around the state mean  $\mathbf{m}_{n-1}$

$$G_n(\boldsymbol{\theta}_{n-1}) \approx G_n(\mathbf{m}_{n-1}) + dG_n(\mathbf{m}_{n-1})(\boldsymbol{\theta}_{n-1} - \mathbf{m}_{n-1}) \quad (4.25)$$

where  $dG_n(\mathbf{m}_{n-1})$  is the Jacobian matrix of the non-linear  $G_n(\cdot)$  evaluated at the previous state estimate  $\mathbf{m}_{n-1}$ . This approach is called the *extended Kalman filter (EKF)*.

The EKF uses the linearization of the forecast function to obtain an approximate forecast state distribution that is Gaussian with an approximate forecast mean and covariance matrix. Naturally, this approach is accurate only to the level that the linearization approximates  $G_n$ . As is often the case, approximation error can affect the fidelity of the resulting method.

**Definition 4.22 (Extended Kalman Filter).** *Assume a non-linear forecast model*

$$\boldsymbol{\theta}_n | \boldsymbol{\theta}_{n-1} = G_n(\boldsymbol{\theta}_{n-1}) + \boldsymbol{\omega}_n.$$

*Suppose  $\boldsymbol{\theta}_{n-1} | Y_{1:n-1} \sim \mathcal{N}_d(\mathbf{m}_{n-1}, C_{n-1})$ . The extended Kalman filter approximates the one-step ahead predictive state distribution as Gaussian*

$$\boldsymbol{\theta}_n | Y_{1:n-1} \sim \mathcal{N}_d(\mathbf{a}_n, R_n)$$

*with parameters*

$$\begin{aligned} \mathbf{a}_n &= G_n(\mathbf{m}_{n-1}) \\ R_n &= (dG_n(\mathbf{m}_{n-1}))C_{n-1}(dG_n(\mathbf{m}_{n-1}))' + W_n \end{aligned}$$

*The conjugate-variance Kalman filter can be extended to include non-linear  $G_n$  in the same manner.*

The primary concern about using the EKF is that it is an approximation based on the linearization of the forecast function and is therefore prone to numerical error from a variety of sources. To understand these sources, we distinguish the approximate

predictive distribution, a Gaussian distribution obtained using the above linearization, and the true predictive distribution. Recall the basis of the Kalman filter: linear transformations of normal random variables are themselves normal. However, due to the non-linearity of  $G_n(\cdot)$ , the true predictive distribution is not necessarily Gaussian.

In addition to the non-Gaussianity of the true predictive distribution, the forecast point estimate and error covariance are also approximations. More precisely, the forecast point estimate  $\mathbf{a}_n$  is the expected value of the Taylor expansion:

$$\begin{aligned}
 \mathbf{a}_n &= \mathbb{E}[G_n(\boldsymbol{\theta}_{n-1})] \\
 &\approx \mathbb{E}\left[G_n(\mathbf{m}_{n-1}) + dG_n(\mathbf{m}_{n-1})(\boldsymbol{\theta}_{n-1} - \mathbf{m}_{n-1})\right] \\
 &= G_n(\mathbf{m}_{n-1}) + dG_n(\mathbf{m}_{n-1})\left(\mathbb{E}[\boldsymbol{\theta}_{n-1}] - \mathbf{m}_{n-1}\right) \\
 &= G_n(\mathbf{m}_{n-1}) + dG_n(\mathbf{m}_{n-1})(\mathbf{m}_{n-1} - \mathbf{m}_{n-1}) \\
 &= G_n(\mathbf{m}_{n-1}).
 \end{aligned}$$

The difference between  $\mathbf{a}_n$  and the mean of the true predictive distribution depends on the accuracy of the approximation in the second line.

Similarly, the covariance matrix  $R_n$  is an estimate on the second moment of the difference between the true state  $\boldsymbol{\theta}_n = G_n(\boldsymbol{\theta}_{n-1}) + \boldsymbol{\omega}_n$  and forecast point estimate  $\mathbf{a}_n = G_n(\mathbf{m}_{n-1})$

$$\begin{aligned}
R_n &= \mathbb{E} \left[ (G_n(\boldsymbol{\theta}_{n-1}) + \boldsymbol{\omega}_n - G_n(\mathbf{m}_{n-1})) (G_n(\boldsymbol{\theta}_{n-1}) + \boldsymbol{\omega}_n - G_n(\mathbf{m}_{n-1}))' \right] \\
&= \mathbb{E} \left[ (G_n(\boldsymbol{\theta}_{n-1}) - G_n(\mathbf{m}_{n-1})) (G_n(\boldsymbol{\theta}_{n-1}) - G_n(\mathbf{m}_{n-1}))' \right] + \mathbb{E} \left[ \boldsymbol{\omega}_n \boldsymbol{\omega}_n' \right] \\
&\approx \mathbb{E} \left[ (dG_n(\mathbf{m}_{n-1})(\boldsymbol{\theta}_{n-1} - \mathbf{m}_{n-1})) (dG_n(\mathbf{m}_{n-1})(\boldsymbol{\theta}_{n-1} - \mathbf{m}_{n-1}))' \right] + W_n \\
&= (dG_n(\mathbf{m}_{n-1})) \mathbb{E} \left[ (\boldsymbol{\theta}_{n-1} - \mathbf{m}_{n-1})(\boldsymbol{\theta}_{n-1} - \mathbf{m}_{n-1})' \right] (dG_n(\mathbf{m}_{n-1}))' + W_n \\
&= (dG_n(\mathbf{m}_{n-1})) C_{n-1} (dG_n(\mathbf{m}_{n-1}))' + W_n.
\end{aligned}$$

As before, the difference between  $R_n$  and the error covariance of the true predictive distribution is dependent on the approximation (located on the third line). With these sources of error, it is important to have processes in place to identify if the a filter diverges from the observations in the application of the EKF. In our implementation we set lower and upper bounds on the state covariance to identify filter divergence.

## Chapter 5

### Model specification for considered measures

Our goal is to use the theory of dynamic linear models (DLMs) developed in Ch. 4 to predict slips. To do so, we model the signal of the considered measures during the stick regime. For the convenience of the reader, small samples of the measures defined in Ch. 3: the wall velocity  $v_x$ , the left-right percolation force  $f_{plr}$ , and the Wasserstein distance  $W_2$ , are shown again in Fig. 5.1. The models we use are fairly rigid, so that the onset of slip can be identified when the predictive errors grow large.

In this chapter, we develop the theory to show why the measures in Fig. 5.1 are amenable to the structure of a DLM, before specifying each of the precise models. Recall from Def. 4.17 what it means to specify a DLM. We must specify each of the following: the forecast matrix  $G_n$ , the observation matrix  $F_n$ , the forecast error covariance matrix  $W_n$ , the observation error variance  $V_n$ , and the initial prior mean  $\mathbf{m}_0$  and covariance matrix  $C_0$ . In the previous chapter, we described how to specify the error structure with  $W_n$  and  $V_n$ , and we will be using a vague initial prior. Here we focus on the selection of  $G_n$  and  $F_n$ . Colloquially, we refer to setting these two matrices as the *construction of the state*.

In the synthetic example in the previous chapter, shown in Fig. 4.6, the structure of the state was provided, based on the physical assumption that the time derivative was relatively constant. In many other contexts, such as meteorology [12] and target tracking [32] the state represents physical quantities, and the DLM structure is derived from the underlying physics. However, the measures in Fig. 3.6 are more complex;

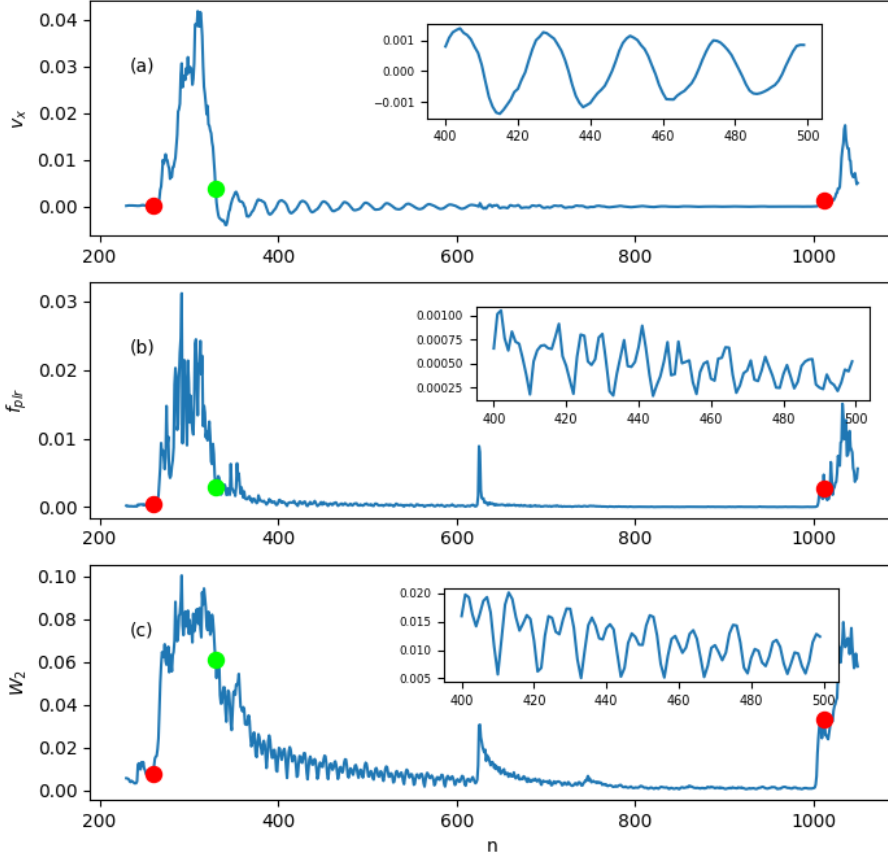


Figure 5.1: Small sample of the evolution of (a) the horizontal velocity of the top wall  $v_x$ , (b) the right left percolation force  $f_{plr}$ , and (c) the  $W_2$  from just before the onset of a slip event to shortly after the start of the next slip. The insets show a sample of the stick regime for each of the measures.

each time series is a result of the dynamics of interacting particles in the granular system, condensed into a sequence of scalar quantities. So, we must use a different approach to construct the state for these measures.

To specify  $G_n$  and  $F_n$  for the considered measures we use an empirical approach, as suggested in [27]. The structure of the state is based on observable *dynamic components*: visible patterns of behavior during the stick regime. To elucidate this

approach, we will first begin with the synthetic data set depicted in Fig. 5.2. Unlike the toy example of the previous chapter, we will not assume any underlying physics. Instead, we will construct a state that describes the visible dynamic components: oscillations superimposed over a slowly increasing level.




Figure 5.2: Sample of synthetically constructed observations from a dynamic linear model of oscillations superimposed over a linearly increasing trend

Before going into the details of constructing the state, we wish to emphasize



the important role of stochasticity in our models. As mentioned, the dynamics of the granular system are quite complex. So, while we typically observe fairly rigid behavior in the measures during the stick regime, the underlying complexities quickly result in a departure from deterministic versions of the models that we will be using. The stochasticity captured in the model error  $W_n$  and observation error  $V_n$  lend the DLM a degree of flexibility that a deterministic model does not possess. As a result, the models adjust to small deviations, while large changes in behavior, e.g., the slip regime, are visible in the predictive error. This is especially important in our case, as the empirically derived models are not inherently connected to the physics of the underlying granular system.

### 5.1 State construction for a synthetic example

The process of empirically constructing the state consists of two basic steps: first, create a deterministic model for the observed signal, then identify the appropriate state variables of the deterministic model. By appropriate, we mean that the chosen variables have linear forecast and observation functions. In Fig. 5.2, the synthetic data appears to roughly consist of a sinusoid superimposed over a line. So, a basic deterministic model is

$$y_n = \beta_0 + \beta_1 n + a \cos(\omega n) + b \sin(\omega n)$$

The first two terms describe the underlying line while the last two terms specify a sinusoid of amplitude  $\sqrt{a^2 + b^2}$  and frequency  $\omega$ . Equivalently, the period is  $P = \frac{2\pi}{\omega}$ .

As a small note, observe that while  $y_n$  is evaluated at integer values  $n \in \mathbb{Z}$ , there is no requirement that the period  $P$  be an integer. That is, with such a deterministic model, we are considering  $y_n$  to be the restriction of a real-valued function  $y(t)$  to

the integers. This is the sense in which we say that  $y_n$  is periodic with a non-integer period. We allow for this because the period of the considered measures does not appear to be strictly integer-valued. For more details on both integer-valued and general periodicity, see [27].

To turn this deterministic model into a DLM, we must capture how the individual variables evolve from frame  $n$  to  $n+1$  (the forecast  $G_n$ ), and then how they combine to reconstruct the value  $y_n$  (the observation  $F_n$ ). Moreover, the forecast and observation functions must be linear. To make this more clear, let us specify the state in this example. It will consist of two dynamic components: one corresponding to the line, and the other to the sinusoid. This will require a state vector with four variables, two for each dynamic component.

We have already seen the model structure for the dynamic component of the line in the toy example in the previous chapter. If  $\theta_n^0$  is the value of the line at  $n$ , and  $\theta_n^1$  is its slope, the forecast function is

$$\begin{aligned}\theta_n^0 &= \theta_{n-1}^0 + \theta_{n-1}^1 \\ \theta_n^1 &= \theta_{n-1}^1.\end{aligned}$$

The dynamic component of the sinusoid is slightly more complicated. First, to be consistent with the DLM framework (linear evolution and observation), the frequency  $\omega$  must be known a priori (or estimated from the data). We then define the *harmonic* value  $S_n$  of the sinusoid to be its height

$$S_n = a \cos(\omega n) + b \sin(\omega n).$$

For the forecast function, we wish to calculate the value of  $S_n$  based on the value of

$S_{n-1}$ . Without further information, this is not possible. With this in mind, we define the *conjugate harmonic* to be

$$S_n^* = -a \sin(\omega n) + b \cos(\omega n).$$

Under this construction, the time evolution of both  $S_n$  and its conjugate  $S_n^*$  can be expressed in terms of each other. That is, even though  $a$  and  $b$  are not known, if we have the values  $S_{n-1}$  and  $S_{n-1}^*$  we can find  $S_n$  and  $S_n^*$  via the following equations:

$$\begin{aligned} S_n &= S_{n-1} \cos(\omega) + S_{n-1}^* \sin(\omega) \\ S_n^* &= S_{n-1} [-\sin(\omega)] + S_{n-1}^* \cos(\omega) \end{aligned} \tag{5.1}$$

The full derivation of this is shown in [27] pp. 105-106, following naturally from trigonometric sum identities. Alternatively, from a geometric point of view, the pair  $(S_n, S_n^*)$  can be thought of as a point on the circle of radius  $\sqrt{a^2 + b^2}$  (equal to  $\sqrt{S_n^2 + (S_n^*)^2}$ , which is constant for varying  $n$ ) that rotates at frequency  $\omega$ .

With  $\omega$  fixed and known this update is linear. In matrix form, Eq. 5.1 becomes

$$\begin{bmatrix} S_n \\ S_n^* \end{bmatrix} = \begin{bmatrix} \cos(\omega) & \sin(\omega) \\ -\sin(\omega) & \cos(\omega) \end{bmatrix} \begin{bmatrix} S_{n-1} \\ S_{n-1}^* \end{bmatrix}$$

To add the sinusoidal dynamic component to the state, we set  $\theta_n^2 = S_n$  and  $\theta_n^3 = S_n^*$ . The forecast, in matrix form, is

$$\boldsymbol{\theta}_n = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos(\omega) & \sin(\omega) \\ 0 & 0 & -\sin(\omega) & \cos(\omega) \end{bmatrix} \boldsymbol{\theta}_{n-1}$$

Lastly, to reconstruct  $y_n$ , the observation function sums the value of the line  $\theta_n^0$  and the value of the sinusoid  $\theta_n^2$ . In matrix form, this is

$$y_n = \begin{bmatrix} 1 & 0 & 1 & 0 \end{bmatrix} \boldsymbol{\theta}_n$$

With this we have largely completed the construction of the state for the synthetic example, capturing the observed trend and overlaid oscillations. The only task remaining is specifying the value of  $\omega$ . For the actual measures, which similarly exhibit oscillations (see Fig. 5.1), we do this through Fourier analysis of a small sample of the signal.

## 5.2 Modeling higher frequency oscillations

This example serves to illustrate how a signal with a linear trend and superimposed oscillation can be decomposed into two dynamic components. However, while the measures we examine sometimes contain nearly pure oscillations as in Fig. 5.2, many appear to have higher frequencies as well, see Fig. 5.1. To model this, we introduce additional dynamic components for higher frequency oscillations. Specifically, we choose frequencies that are multiples of the base frequency.

We refer to the dynamic components as Fourier modes, as this model is based on the Fourier decomposition of a periodic function. We briefly review this decomposition here for the convenience of the reader. First, consider a general well-behaved

periodic function  $y_n$  with known frequency  $\omega$ , referred to as the *fundamental frequency*. Equivalently, the function has a known period  $P = \frac{2\pi}{\omega}$ .

Since we model any local level around which the fluctuations occur by a separate dynamic component, suppose for illustration that  $y_n$  is zero-centered. To understand the idea of a zero-centered function, recall that we model  $y_n$  as the restriction of a continuous function  $y(t)$  to the integers. A zero-centered periodic function is then characterized by

$$\int_0^P y(t) dt = 0$$

The Fourier decomposition states that a zero-centered periodic function can be approximated by the sum of sinusoids whose frequencies are multiples of the fundamental frequency. Similarly to the single sinusoid  $S_n$  in the synthetic example, the  $j$ -th such sinusoid will be denoted by  $S_{n,j}$ , and we call it the  $j$ -th harmonic. More precisely, at frequency  $\omega_j = j\omega$  the  $j$ -th harmonic is given by

$$S_{n,j} = a_j \cos(\omega_j t) + b_j \sin(\omega_j t)$$

Then the Fourier approximation states

$$y_n \approx \sum_{j=1}^J S_{n,j} \tag{5.2}$$

Proof of this approximation can be found in [31].

The construction of the  $j$ -th Fourier mode proceeds identically to the single-frequency construction for the synthetic example. First the  $j$ -th *conjugate harmonic* is defined by

$$S_{n,j}^* = -a_j \sin(\omega_j n) + b_j \cos(\omega_j n)$$

then the forecast is given by

$$\begin{aligned} S_{n,j} &= S_{n-1,j} \cos(\omega_j) + S_{n-1,j}^* \sin(\omega_j) \\ S_{n,j}^* &= S_{n-1,j} [-\sin(\omega_j)] + S_{n-1,j}^* \cos(\omega_j) \end{aligned} \tag{5.3}$$

for each  $j = 1, \dots, J$ .

In the zero-centered case, the state  $\boldsymbol{\theta}$  is a vector of length  $2J$  consisting of  $[S_{n,1}, S_{n,1}^*, \dots, S_{n,J}, S_{n,J}^*]'$ . The observation function reconstructs the value  $y_n$  by the Fourier sum in Eq. 5.2. In matrix form,  $G_n$  consists of the block-diagonal concatenation of  $J$   $2 \times 2$  matrices, one for each of the Fourier modes, while  $F = \begin{bmatrix} 1 & 0 & \dots & 1 & 0 \end{bmatrix}$ .

Much of the value of the general periodic model given above lies in its parameterization. Only two values need to be specified: the fundamental frequency  $\omega$  and the number of Fourier modes  $J$ . Then, as in the synthetic example, any underlying trend is captured by a separate dynamic component.

### 5.3 Model specification for considered measures

In this section, we specify the three DLMS that model the evolution of the  $v_x$ ,  $f_{plr}$ , and  $W_2$  measures during the stick regime, respectively. For each of the three measures, we begin by constructing the state. As in the synthetic example in Sec. 5.1, the state structure for each measure is decomposed into an underlying trend and oscillation(s). Further, recall from Sec. 5.2 that the oscillations are further decomposed into Fourier modes, for which we must simply specify the fundamental frequency  $\omega$  and the number of Fourier modes  $J$ .

After constructing the state, we specify the model error covariance matrix and observation error variance. Recall from Sec. 4.3.1 that we specify the model error covariance matrix using the discount factor methodology, by setting the parameter

$\delta$ , the discount factor. Further, using the conjugate-variance theory of Sec. 4.3.2, we model the observation error variance as a time-invariant random variable updated by the data at each frame  $n$  by Bayes' formula. In the models for each of the three measures, the required initial distribution is always chosen to be the inverse gamma distribution  $\mathcal{IG}(2, 10^{-3})$ , a vague prior.

Lastly, for each model, we specify the initial prior for the state that is distributed multivariate normal. As noted in Ch. 4, with a vague initial prior the precise value of the mean is relatively unimportant, as the Kalman filter quickly adjusts to the observed data. So, for each measure, we select a prior where the state mean is on the same scale as the typical values of the measure, and the initial state covariance is very large. Notably, when using a conjugate-variance model, the large state covariance is already specified by choosing a vague prior for the observation error variance. Additionally, in all the analysis, we discard the first one hundred frames, to definitively guarantee the filter has adjusted to the data.

Altogether, defining each model consists of specifying: (1) the underlying trend, (2) the fundamental frequency  $\omega$ , (3) the number of Fourier modes  $J$ , and (4) the discount factor  $\delta$ . The underlying trend in each measure is based on a visual examination of the signal over a few samples of stick frames, a process discussed in more detail later in this section. The frequency  $\omega$  is determined from Fourier analysis of a sample of the data, the first ten percent. Lastly,  $J$  and  $\delta$  were determined from model diagnostics of a few samples of stick frames. Our extensive numerical investigation confirmed that all results discussed in Ch. 6 are robust under rather large variations of  $J$  and  $\delta$  and even small variations of  $\omega$ .

The model diagnostics governing the choice of  $J$  and  $\delta$  are based on the normalized predictive error (NME) of the model defined in Eq. 4.24. If the model is a good fit for the data, then the errors are uncorrelated and they are distributed according to

the standard normal distribution [27]. We use the Ljung-Box statistic to test [24] for possible autocorrelation and the Shapiro-Wilk test [30] to see if the normalized errors are distributed according to the standard normal distribution. We chose the models that yield the best diagnostics during the stick regimes shown in Figure 5.3 (and later in Figs. 5.4 and 5.5).

We start by setting up the model for the wall velocity,  $v_x$ . Fig. 5.3 shows that during the stick regime,  $v_x$  oscillates around an almost constant value. Note that oscillations exhibit multiple frequencies. Their dominant frequency  $\omega = 0.2476$  was determined from the first 10% of the data. These oscillations are modeled with two harmonic components with frequencies  $\omega_i = i * \omega$ , for  $i = 1, 2$ .

The value around which the oscillations occur is represented by the last component  $\theta^5$  of the state vector  $\boldsymbol{\theta}$ . The value of  $\theta_n^5$  changes very slowly, and in the deterministic forecast we consider  $\theta_n^5$  to be constant. In the deterministic forecast, the Fourier modes evolve according to Eq. 5.1. For the convenience of the reader, we state the forecast function  $G$  as

$$(G\boldsymbol{\theta}_{n-1})^{2i-1} = \theta_{n-1}^{2i} \sin \omega_i + \theta_{n-1}^{2i-1} \cos \omega_i, \quad (5.4a)$$

$$(G\boldsymbol{\theta}_{n-1})^{2i} = \theta_{n-1}^{2i} \cos \omega_i - \theta_{n-1}^{2i-1} \sin \omega_i, \quad (5.4b)$$

$$(G\boldsymbol{\theta}_{n-1})^5 = \theta_{n-1}^5, \quad (5.4c)$$

for  $i = 1, 2$ . Next, recall from Sec. 5.2 that the observation function,  $F$ , sums the harmonics of the signal and the underlying trend. In this case,  $F$  is given by

$$F\boldsymbol{\theta}_n = \theta_n^1 + \theta_n^3 + \theta_n^5. \quad (5.5)$$

To specify the model error covariance, we use the discount factor  $\delta = 0.6$ . Lastly, the



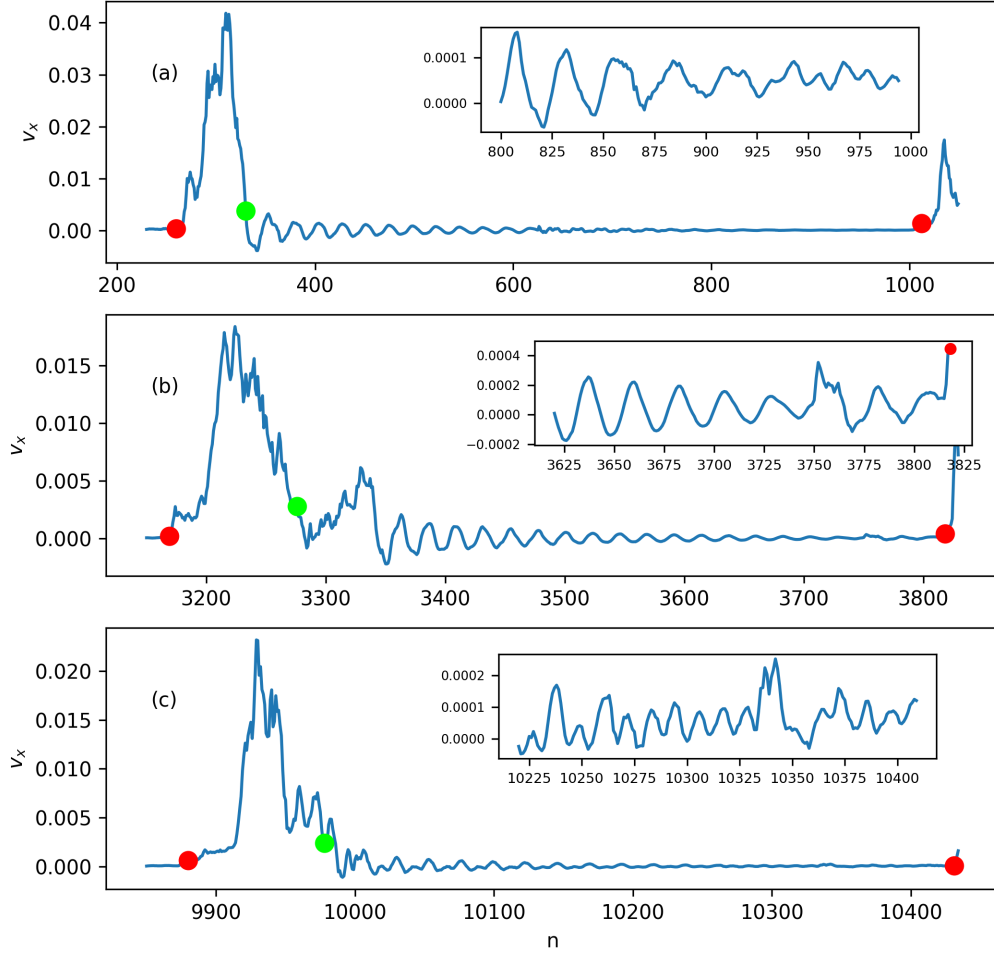


Figure 5.3: Samples of wall velocity  $v_x$  for individual stick-slip cycles. The green dots mark the beginning of the stick regime while the red dots mark the end, as determined from the offline method. Each inset displays a sub-sample of the stick regime. The stick regime from our simulations is characterized not just by low wall movement, but by low-amplitude periodic behavior. This behavior varies from purely oscillatory in (b) to the more erratic periodicity observed in (c). In contrast, the behavior of  $v_x$  during slip is far more erratic, only roughly characterized by its magnitude.

initial prior distribution for the state is a normal distribution with mean  $\mathbf{m}_0$  where  $m_0^j = 0$ , for  $j = 1, \dots, 4$ ,  $m_0^5 = v_x(0)$ , and a scaled covariance matrix  $\bar{C}_0 = I$ . As

noted, this prior is vague as a result of selecting a vague prior for the observation error variance.

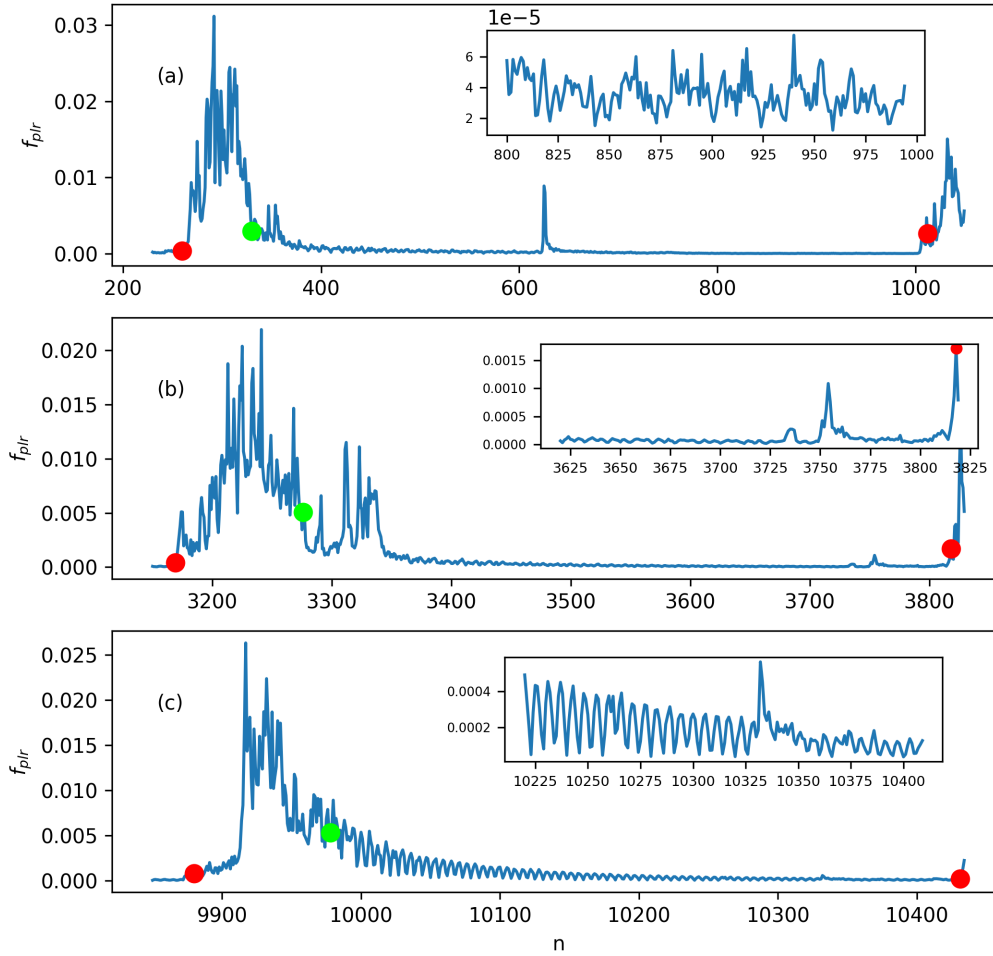


Figure 5.4: Samples of the  $f_{plr}$  measure. Periodic behavior is often present in the stick regime.

Now we turn our attention to the left-right percolation force,  $f_{plr}$ , shown in Fig. 5.4. The dominant frequency of oscillations for this measure is  $\omega = 0.5401$ . Again we use two harmonic components. As in the case of  $v_x$ , the evolution of the variables

$\theta^1, \dots, \theta^4$ , representing the harmonic components, is described by Eqs. 5.4(a)-(b).

The most important difference between the behavior of  $v_x$  and  $f_{plr}$  is that the value around which  $f_{plr}$  oscillates is not changing slowly. This value decreases in an approximately linear manner. To incorporate this decrease into the model, we use a linear trend (just as in the synthetic example in Sec. 5.1). We specify two state variables:  $\theta_n^6$  which models the slowly changing slope, and  $\theta_n^5$  which captures the decreasing value around which the signal oscillates. These two components evolve according to the following equations

$$(G\boldsymbol{\theta}_{n-1})^5 = \theta_{n-1}^5 + \theta_{n-1}^6 \quad (5.6a)$$

$$(G\boldsymbol{\theta}_{n-1})^6 = \theta_{n-1}^6. \quad (5.6b)$$

The observation function,  $F$ , again sums the variables capturing the oscillations and the variable  $\theta_n^5$  representing the value around which the oscillations occur, so  $(F\boldsymbol{\theta}_n) = \theta_n^1 + \theta_n^3 + \theta_n^5$ . To specify the model error covariance we use the discount factor  $\delta = 0.93$ . As in the case of  $v_x$ , the initial distribution for the state vector  $\boldsymbol{\theta}_0$  is a normal distribution centered at  $\mathbf{m}_0$  where  $m_0^j = 0$ , for  $j = 1, \dots, 4$ ,  $m_0^5 = f_{plr}(0)$ ,  $m_0^6 = 0$ , and the scaled covariance matrix  $\bar{C}_0 = I$ .

Fig. 5.5 shows samples of the last considered measure, the zero-level Wasserstein distance  $W_2$ . This measure exhibits oscillations with the dominant frequency  $\omega = 0.5387$ , and using two harmonic components is once again sufficient. The most important difference between the behavior of  $f_{plr}$  and  $W_2$  is that the value around which  $W_2$  oscillates decays in a roughly exponential manner, see Figure 5.5. Thus, it can be modeled by a function  $e^{\mu t}$  where the value of  $\mu$  is allowed to change slowly through the error term. To incorporate this exponential decrease in the model, we use two state components:  $\theta_n^6$  which models the slowly changing value  $\mu$ , and  $\theta_n^5$  which

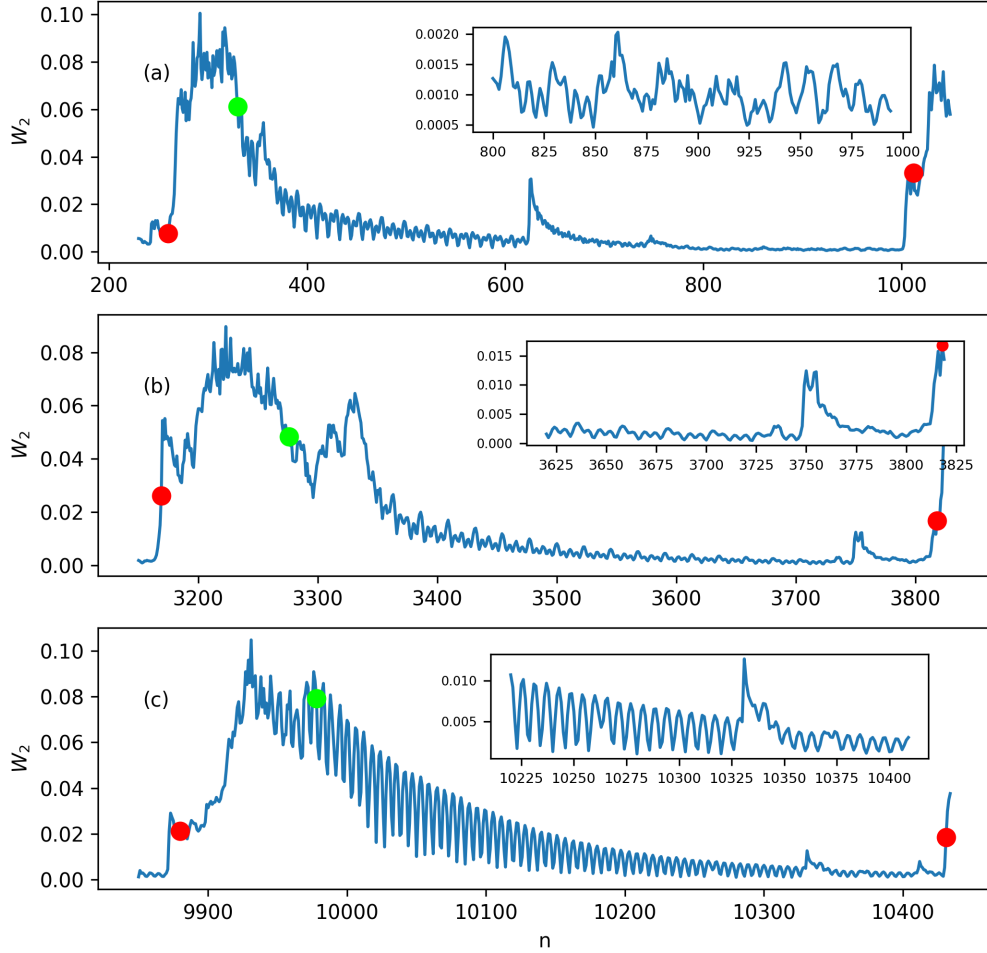


Figure 5.5: Samples of the  $W_2$  measure. The stick behavior exhibits oscillations at multiple frequencies.

captures the behavior of  $e^{\mu t}$ . The following equations ensure the desired evolution of these components

$$(G\theta_{n-1})^5 = \theta_{n-1}^5 e^{\theta_{n-1}^6}, \quad (5.7a)$$

$$(G\theta_{n-1})^6 = \theta_{n-1}^6. \quad (5.7b)$$

The discount factor is  $\delta = 0.67$  and the initial distribution for the state vector  $\boldsymbol{\theta}_0$  is a normal distribution centered at  $\mathbf{m}_0$  where  $m_0^j = 0$ , for  $j = 1, \dots, 4$ ,  $m_0^5 = W_2(0)$ ,  $m_0^6 = 0$  and the scaled covariance matrix  $\bar{C}_0 = I$ .

Eq. 5.7(b) is non-linear, so the model is not a DLM. However, recall the theory developed in Sec. 4.4; the extended Kalman filter (EKF) can be used to accommodate the non-linearity of the forecast function [12]. The EKF uses the Jacobian matrix of the forecast function, evaluated at the current state mean, to produce a linearized estimate of the predictive distribution of  $\boldsymbol{\theta}_n$ . There are well-known divergence problems with EKF [12] since using the Jacobian matrix can lead to unrealistically small values of the covariance matrix of the state distribution. This can impede the ability of the filter to assimilate the observed data and eventually cause divergence of the model. In our case, this happens when the value of  $\theta_n^5$  becomes too small and is rounded to zero. To retain numerical integrity during our computations, we re-initialize the filter upon encountering filter divergence. We note that this occurred only once in the entire data set.

In summary, for all of the measures the DLMS are combinations of a trend component (constant, linear, exponential decrease) and a periodic component. The fundamental frequency of the periodic component is different for each measure, as is the discount factor for specifying model error. The models are summarized below in Table 5.1.

	$v_x$	$f_{plr}$	$W_2$
Trend	Constant	Linear Decrease	Exponential Decrease
$\omega$	0.2476	0.5401	0.5387
$J$	2	2	2
$\delta$	0.60	0.93	0.67

Table 5.1: Model specifications for  $v_x$ ,  $f_{plr}$ , and  $W_2$ , detailing the type of trend, the fundamental frequency of the periodic component  $\omega$ , the number of harmonic components  $J$ , and the discount factor  $\delta$ .

#### 5.4 Detections from dynamic linear models (DLMs)

By construction, the DLM for a measure models the typical stick regime behavior. Consequently, the predictive error sharply rises upon departure from the standard stick regime, such as during a slip event. To identify such events we specify a maximal acceptable value of the NME,  $T_e$ . For a given DLM and a value  $T_e$ , we say that the model detects the beginning of a slip event at frame  $n_0$  if the NME,  $e_n$ , satisfies  $|e_{n_0-1}| < T_e$  and  $|e_{n_0}| > T_e$ . We could say that this slip event ends when the value of  $|e_n|$  drops below  $T_e$ . However, it can happen incidentally that the value of  $|e_n|$  becomes small during the slip. Therefore, to declare that slip ends at  $n_1$  we require that all the values  $|e_{n_1}|, |e_{n_1-1}|, \dots, |e_{n_1-m}|$  have to be smaller than  $T_e$  for  $m = 25$ . The value of  $m$  is chosen to be roughly one period of the dominant oscillations during the stick phase, see Fig. 5.3.

We refer to the period during  $[n_0, n_1]$  as a *detection*, defined precisely here.

**Definition 5.1.** *Let  $e_0, \dots, e_N$  be the sequence of normalized model errors of a dynamic linear model. For an error threshold  $T_e > 0$  and minimum event duration  $m > 0$ , the period during  $[n_0, n_1]$  is a **detection** at  $T_e$  if*

1.  $|e_{n_0-1}| < T_e$ ,
2.  $|e_{n_0}| > T_e$ , and
3.  $n_1 = \min\{n > n_0 : |e_{n'}| < T_e \forall n' \in [n - m, n - 1]\}$ .

As previously mentioned, while such a detection may indeed correspond to a slip event, it also may not. The term ‘detection’ indicates the broad variety found by the online method for the various measures, each of which corresponds to some marked change in the force network of the granular system.

Additionally, note that this detection approach is *online* in the sense that the identification of  $n_0$  depends only on data (i.e., the model error values) before it, and the same for  $n_1$ . This is in contrast to the *offline* two-threshold approach to identifying slips detailed in Ch. 2, in which the slip start is determined using data after the fact.

**Remark 5.2.** *For a DLM that accurately models the observed time series, the normalized model error NME will be distributed standard normal. Theoretically, this allows for  $T_e$  to have a standard interpretation across models as a credible interval  $(-T_e, T_e)$  for the standard normal distribution at the corresponding significance level [20]. In practice, the accuracy to which the DLM models the observed time series can affect the magnitude of the NME, occluding the specificity of that interpretation. Nevertheless, the standardization allows us to compare models on a similar scale, regardless of the scale of the underlying time series.*

## Chapter 6

### Results

We now show the results of using the dynamical linear models (DLMs) specified in the previous chapter to detect upcoming regime changes. First, recall the three measures introduced in Ch. 3: the top wall's horizontal velocity  $v_x$ , the left-right percolation force  $f_{plr}$ , and the zero-level second Wasserstein distance  $W_2$ . Based on similar research in sheared granular systems [10, 2] discussed in Ch. 2, we expect the  $W_2$  to be the most sensitive to localized changes that occur before slips, the  $f_{plr}$  to reflect the more globalized changes, and the  $v_x$  to act as a baseline.

For reference, we consider the slips as defined in Ch. 2 from large increases in  $v_x$ , where the threshold velocity is based on an observed bimodal distribution often present in sheared granular systems [17, 10]. As in Ch. 2, any detection corresponding to this second category is called a *micro-slip*, as in that literature. Notably, our methodology observes not only these two classes of detections but a third distinct class, visible only in the  $W_2$  measure. We refer to detections in this class as purely *local changes*, detected only by the zero-level Wasserstein distance  $W_2$  model. Detections in this category are characterized by localized disruptions in the force network which do not spread across the whole system.

With further analysis, we find that the percolation force  $f_{plr}$  model identifies roughly the same events as the  $v_x$  model (e.g., slips and micro-slips), but slightly earlier, implying global changes in the force network tend to occur before increased wall movement. Lastly, the  $W_2$  DLM identifies nearly all these slips and micro-slips



even earlier, as well as the additional category of local changes. This supports our premise that topological measures, such as the  $W_2$ , can be used for the prediction of upcoming slips.

## 6.1 Tri-partite classification of detections

Before continuing to the results, we define precisely what we refer to as ‘slip detections.’ We draw a sharp distinction between what we term *reference slips*, the fixed intervals defined in Ch. 2, and *slip detections*, detected by one of the three models at a specific error threshold  $T_e$ , see Def. 5.1. The former is fixed and effectively treated as a ground truth in our analysis, while the latter varies based on the error threshold and modeled measure.

**Definition 6.1.** A *slip detection* is a detection, as defined in Def. 5.1, whose time interval overlaps a reference slip, as defined in Ch. 2.

This distinction between reference slips and slip detections allows us to analyze how far in advance a measure detects reference slips as the difference between the *detection time*, the start of a slip detection, and the *start time*, the beginning of its reference slip. Note that detection time, and even whether a slip is detected at all, varies with the chosen error threshold. Hence, much of our results focus on the effects of varying the error threshold for each measure.

As already noted, all three considered measures have detections that do not fall into the category of slip detections. The second category is composed of detections made by the wall velocity  $v_x$  model that are not slip detections, and are termed *micro-slip detections*. Unlike reference slips, the references for micro-slips are detections, specifically, those made by the  $v_x$  model at a critical error threshold of  $T_e = 0.1$ . The justification of this critical threshold is explained in Sec. 6.2, but here it suffices to

say that is the smallest error threshold for which Def. 5.1 is reasonably defined over the  $v_x$  model.

**Definition 6.2.** *Detections, as defined in Def. 5.1, made by the  $v_x$  model at the critical threshold  $T_e = 0.1$  that are not slip detections are the **reference micro-slips**, while any detection whose time interval overlaps a reference micro-slip is a **micro-slip detection**.*

Fig. 6.1 shows the densities of two metrics on slip size over the detections of the  $v_x$  DLM at  $T_e = 0.1$ : the maximum velocity during and magnitude of wall displacement, denoted by  $d_x$ , over the detections. The densities clearly show that for either metric, the distributions are bimodal. This bimodality confirms the bi-classification noted in Ch. 2. We note that the slip and micro-slip detections are roughly delineated by the valley separating the peaks.




Figure 6.1: Probability density of (a) maximum of wall velocity  $v_x$  and (b)  $d_x$ , the wall displacement in the  $x$  direction, on the  $\log_{10}$  scale during detections made by the wall velocity model at the critical error threshold of  $T_e = 0.1$ . The bimodal nature of these distributions indicates the presence of two classes.

The third category we consider is seen exclusively in the  $W_2$  metric. A significant portion (28%) of the detections made by the  $W_2$  model do not fall into the category of slip or micro-slip detections. Fig. 6.2 shows an example of such a detection. The fact that there is no visible change in the  $f_{plr}$  quantifying the global changes of the network implies that the detected change of the force network is localized. Indeed, for this particular detection, there is a large spike in the broken force,  $f_{bc}$  as shown in Fig. 6.2(d), which is a measure of the force network defined as the sum of all (normal) forces at the contacts that are broken between subsequent frames. The spike in  $f_{bc}$  is the result of a single large broken contact, indicating that, at least for this case, the  $W_2$  detection is strongly related to broken contacts and highly localized.

The example of Fig. 6.2 is not unique; nearly all the detections in the third category, which we will define shortly, are not detected by the  $f_{plr}$  model. Hence, we feel justified in calling these *local changes*. Similarly to the reference micro-slips, a set of reference local changes is defined based on detections made by the  $W_2$  model at a critical error threshold. In this case, a threshold of  $T_e = 0.4$  is used, which will be justified in Sec. 6.2.

**Definition 6.3.** *Detections, as defined in Def. 5.1, made by the  $W_2$  model at the critical threshold  $T_e = 0.4$  that are not slip or micro-slip detections are the **reference local changes**, while any detection whose time interval overlaps a reference local change is a **local change detection**.*

We note that similar local changes, including non-affine motion or rotations of the particles, have been discussed in previous works [25, 4].

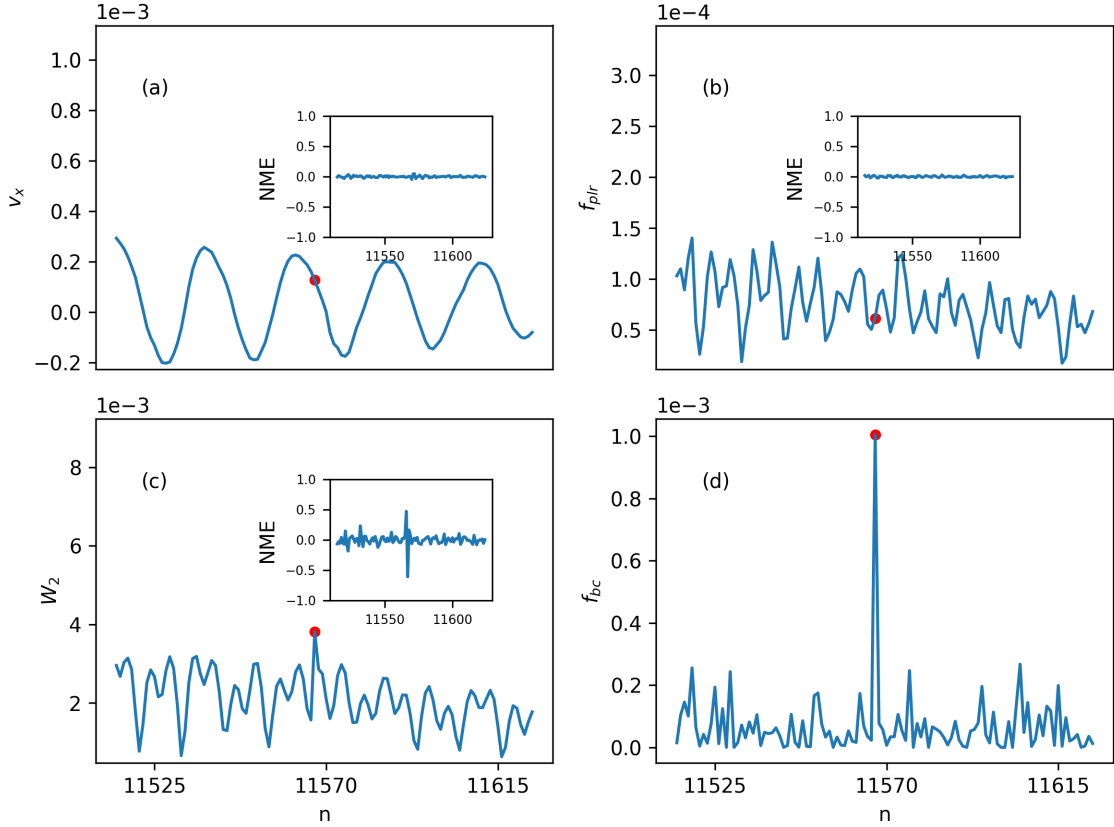


Figure 6.2: Example of a detection from the  $W_2$  model at  $T_e = 0.4$  that is not detected by the  $v_x$  (or  $f_{plr}$ ) model. Panels (a)-(c) show the  $v_x$ ,  $f_{plr}$ , and  $W_2$  with the inset displaying the corresponding NME, while (d) depicts the broken contact force,  $f_{bc}$ . The red point marks the detection time of the  $W_2$  model. The fact that  $v_x$  and  $f_{plr}$  do not detect this event suggests that it corresponds to a local change in the force network. The corresponding spike in  $f_{bc}$  is the result of a single broken contact with a large force value.

## 6.2 Detections for varying error thresholds

Detection time, and even the detections themselves, depends on the error threshold  $T_e$ . Fig. 6.3 is a representative example of how detection depends on the error threshold for a slip (a)-(c), and a micro-slip (d)-(f). We begin by exploring the trends in this example before generalizing to the entire data set.

Starting with Fig. 6.3(a)-(c), we compare the detection times of a slip made by

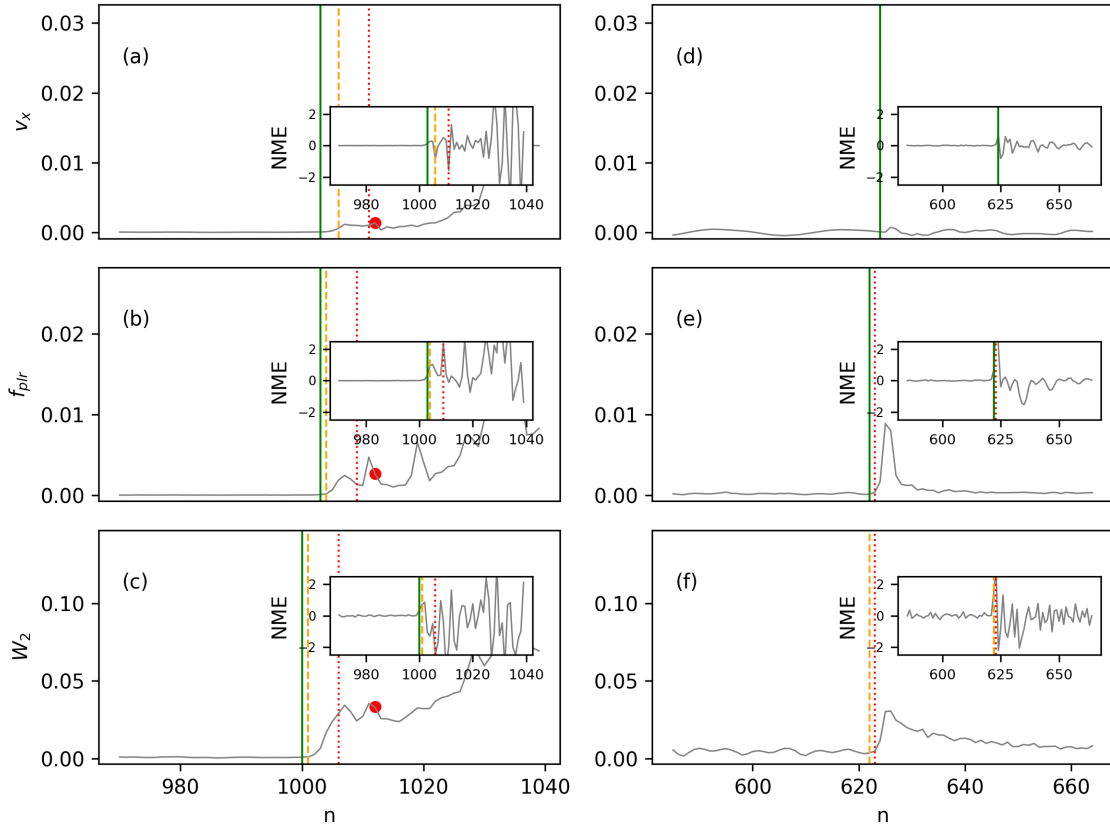


Figure 6.3: Two detection examples: (a)-(c) a slip event detected by the offline method (shown in Figs. 5.3(b), 5.4(b), and 5.5(b)), and (d)-(f) a micro-slip event. The measure  $v_x$  is shown in (a), (d),  $f_{plr}$  in (b), (e) and  $W_2$  in (c), (f). In (a)-(c), the red dots identify the start of the slip event determined by the offline method. The insets show the respective NME for individual DLMs, with the vertical lines marking detection times for individual DLMs using three different values of  $T_e = 0.1, 0.4,$  and  $1.5$  for the solid green line, dashed orange line, and dotted red line, respectively.

the  $v_x$ ,  $f_{plr}$ , and  $W_2$  models, respectively, for three different error thresholds. In this example, decreasing the value of the error threshold  $T_e$  leads to an earlier detection time. We note that for each fixed value of  $T_e$ , the  $W_2$  model detects the slip first, followed by  $f_{plr}$ , and then  $v_x$ .

Next, Fig. 6.3(d)-(f) showcases micro-slip detections and illustrates some difficulties that may arise at different error thresholds. First, observe how the  $v_x$  model does not detect this event for the larger thresholds  $T_e = 1.5$  and  $T_e = 0.4$ , since its NME rises only slightly above 0.1. On the other hand, using the smaller threshold  $T_e = 0.1$  with the  $W_2$  model does not result in a detection inside the depicted time window. Rather, at this threshold the entire depicted time window is inside a detection that starts far earlier. So, if  $T_e$  is too large, then slip events may be detected late or even missed, while choosing  $T_e$  too small can cause large stretches to be captured in a single detection, merging detections that should be distinct.

Fig. 6.4(a) shows the number of detections as a function of  $T_e$ , for the three measures. Observe that each of the measures shows a prominent peak, located at  $T_e = 0.1$  for the  $v_x$  and  $f_{plr}$ , and  $T_e = 0.4$  for the  $W_2$ . If we decrease  $T_e$  below this threshold (as specified for each of the measures), then distinct detections start to merge and their number decreases. Hence, we consider these to be *critical thresholds* (mentioned in Sec. 6.1 in Def. 6.1, 6.2, and 6.3), as they are the lowest threshold for which distinct detections do not merge. This is made precise in the following definition.

**Definition 6.4.** *The **base detections** for each of the three considered models are the detections made by a model at its **critical error threshold**, denoted by  $T_e^c$ :  $T_e^c = 0.1$  for the  $v_x$  and  $f_{plr}$  models and  $T_e^c = 0.4$  for the  $W_2$  model.*

For the remainder of our analysis, we will only consider the values of  $T_e \geq T_e^c$ .




Figure 6.4: Panel (a) shows the total number of detections from the  $v_x$ ,  $f_{plr}$ , and  $W_2$  models as a function of  $T_e$ . The curves corresponding to  $v_x$  and  $f_{plr}$  show a well-pronounced maximum at  $T_e = 0.1$ , while the curve corresponding to  $W_2$  shows a maximum at  $T_e = 0.4$ . The dashed line indicates the number of reference slips in the data set. Panel (b) shows the number of detections after removing aftershocks, i.e., the number of base detections (as defined by Def. 6.4) that are still detectable at  $T_e$ .

Note that with this terminology, the reference micro-slips in Def. 6.2 are simply the  $v_x$  base detections that are not slip detections, while the reference local changes in Def. 6.3 are the  $W_2$  base detections that not slip or micro-slip detections.

Observe how in Fig. 6.4(a) for the  $v_x$  and  $f_{plr}$  there is a pronounced plateau followed by a gentle rise. Further investigation revealed that some slips are followed by almost immediate aftershocks, see Fig. 5.3(b) for an example. Detection of these aftershocks causes this rise because, for small values of  $T_e$ , the main event and its aftershock are identified as one detection; however, as  $T_e$  increases, both the main event and the aftershock are eventually identified separately. To avoid the identification of aftershocks as separate events and to accurately capture the number of distinct detections, we consider an adjusted detection count, shown in Fig. 6.4(b). The aftershocks are removed by considering ‘different’ detections at some  $T_e \geq T_e^c$

to be the same if they occur during a single base detection. The adjusted detection count steadily decreases for all measures and error thresholds  $T_e \geq T_e^c$ .

### 6.3 Insight provided by different measures

Now we turn our attention to the relation between the reference slips and the detections from each of the measures. To better understand how the composition of the detections varies over the error threshold  $T_e$ , we use two ratios: the true positive rate and the false positive rate. For a given measure, the true positive rate,  $r_{tp}(T_e)$ , is the ratio of the number of reference slips detected at the threshold to the total number of reference slips. The false positive rate,  $r_{fp}(T_e)$ , is the ratio of the non-slip detections, to the total number of detections made at that threshold.

Fig. 6.5 depicts the ratios,  $r_{tp}(T_e)$  and  $r_{fp}(T_e)$ . Observe in Fig. 6.5(a) that for each measure there is a range of error thresholds for which  $r_{tp} = 1$ . That is, each of the three models is capable of detecting all reference slips, within these respective ranges. However, in Fig. 6.5(b) note that for each measure there is a considerable number of false positives at thresholds near the critical threshold  $T_e^c$ . The large false positive rates are caused by the micro-slip detections (and local change detections, in the case of the  $W_2$  model).

So, increasing the error threshold  $T_e$  reduces the proportion of non-slip detections but also results in some reference slips not being detected (and, as we will show in Sec. 6.4, later detection times). At this point, we cannot reliably distinguish between slip detections and the other two types until after the fact. Further research will be necessary to address this issue; we propose some potential avenues in Ch. 7.

As with the detection counts in Fig. 6.4, we calculate an adjusted false positive rate that does not distinguish aftershocks and is thus more representative of the






Figure 6.5: (a) True positive rate,  $r_{tp}(T_e)$ , and (b) false positive rates,  $r_{fp}(T_e)$ . The dashed line marks the false positive rate and the solid line marks the adjusted false positive rate obtained by removing the aftershocks.

proportion of non-slip detections. The adjusted false positive rate at  $T_e$  is the ratio of base non-slip detections still present at  $T_e$ , and the total number of base detections still present at  $T_e$ . Fig. 6.5(b) shows that the adjusted  $r_{fp}(T_e)$  is always decreasing and decreases significantly faster than the non-adjusted  $r_{fp}(T_e)$ .

To further investigate the relation between wall movement and force network rearrangements, we introduce a measure defined over the base detections, the *maximal error threshold*  $T_e^m$ . The maximal error threshold of a base detection is defined as the largest value of  $T_e$  for which it is still detectable. The maximal error threshold is a measure of model error during a detection, indicative of the magnitude of the deviation from typical stick regime behavior for that model. So, we use  $T_e^m$  for the  $f_{plr}$  as a proxy measure of global change and  $T_e^m$  for the  $W_2$  as a proxy measure of local change.

Fig. 6.6 shows the relation between the values of  $T_e^m$  for different measures, with




Figure 6.6: Scatter plot of the  $T_e^m$  values for base detections made by both considered measures (a)  $v_x$  and  $f_{plr}$ , (b) W2B0 and  $v_x$ , and (c) W2B0 and  $f_{plr}$ . The blue color indicates slip detections while the green color indicates micro-slip detections.

each point corresponding to a base detection made by both of the measures specified on the axes labels. For each pair of measures, there is a general correlation between the  $T_e^m$  values, indicating a rough agreement between the ‘magnitude’ of detections, at least as quantified by the maximal error threshold.

However, in the comparison between the  $v_x$  and  $f_{plr}$  base detections shown in Fig. 6.6(a), there is a notable difference in slope between the slip and micro-slip detections. The maximal error thresholds of the  $v_x$  slip detections increase more rapidly with the  $f_{plr}$  maximal error thresholds, compared to the micro-slip detections. Further, an  $f_{plr}$   $T_e^m$  value larger than  $\approx 3$ , almost certainly indicates a slip detection (this can also be seen in the near-zero adjusted  $r_{fp}$  in Fig. 6.5(b)), while an  $f_{plr}$  detection with  $T_e^m < 1$  is a micro-slip. That is, there is a relatively well-defined critical value for the ‘magnitude’ of the global change of the force network, and surpassing this value leads to a slip event.

In contrast, Fig. 6.6(b)-(c) indicates that the magnitude of local change, as measured by the  $W_2$   $T_e^m$ , is far less connected to the distinction between slips and micro-slips. While very large (small) values of  $W_2$   $T_e^m$  always correspond to a slip (micro-slip)

detection, there is a wide range of maximal error thresholds obtained by the  $W_2$  model for which we observe both slip and micro-slip detections. That is, the connection between wall activity and local changes in the force network, viewed through persistent homology, is not clear, since even large local changes may not lead to a slip.

Next, we explore the relation between maximal error thresholds of base detections and the more traditional measure of wall displacement,  $d_x$ , over a detection. Fig. 6.7(a) shows that the values of  $T_e^m$  obtained by the  $v_x$  measure are linearly proportional to the wall displacement. We also observe a relatively clear distinction between micro-slip and slip detections, since Fig. 6.7(b) shows a clear change of slope around the boundary between the slip and micro-slip detections. This change of slope further indicates that for the slips the wall activity increases more rapidly with the growing global change in the force network than for the micro-slip detections. Fig. 6.7(c) relates the size of the local change of the force network and the wall displacement. As in Fig. 6.6(b)-(c), there is a large range of values  $T_e^m$  obtained by  $W_2$  for which both slip and micro-slip detections occur. A local change of the force network with  $T_e^m$  in this range has uncertain consequences. It might dissipate without causing a global change in the force network or trigger a global change accompanied by a micro-slip or slip. However, sufficiently small local changes never trigger a slip while sufficiently large ones always do.

#### 6.4 Predictive power of different measures

So far, we have compared the detections made by each of the three measures. This section examines how far in advance the reference slips can be detected by our online methodology. The representative example of Fig. 6.3(a)-(c) suggests that the detection time of a slip depends on both the considered measure and  $T_e$ . To examine this




Figure 6.7: Scatter plot of the wall displacement during a base detection and its values  $T_e^m$ , for events detected at  $T_e^c$  by (a)  $v_x$ , (b)  $f_{plr}$ , and (c)  $W_2$ . The blue color indicates slip detections while the green color denotes micro-slip detections. In c) the red color indicates local change detections.

quantitatively, we consider the times at which the offline method detects the start of the reference slips as a baseline, referred to as the *start times*. We compare the start times with the *detection times* obtained by different measures and values of  $T_e$ , the first frame of the detection. From this, we focus on two metrics: the *median advance notice* and the *advance ratio*, measuring how far in advance slips are detected and how many slips are detected in advance, respectively. We emphasize that the purpose of this comparison is quantifying the effectiveness of the online methods based on DLMS, and not comparing the performance of online and offline methods. The offline approach uses future data for identifying slip events and is not appropriate for predictive purposes.

For a given slip detection, let  $t_0$  and  $t'_0(T_e)$  be the start time and detection time, respectively. We then define the advance notice of the detection as  $t_a(T_e) = t_0 - t'_0(T_e)$ . If  $t_a(T_e) > 0$ , then the slip is detected in advance. To analyze how far in advance the




Figure 6.8: Comparison of offline and online methods: a) the median advance notice  $\langle t_a \rangle$  taken over slip detections by the given measure at  $T_e$ ; (b) the ratio of slip detections that were detected in advance  $r_a(T_e)$ .

reference slips can be detected, using different measures and values of  $T_e$ , we consider the median advance notice  $\langle t_a \rangle$  defined as follows. For a given measure and  $T_e$ , the value  $\langle t_a \rangle(T_e)$  is the median over the slip detections at that error threshold. The advance ratio,  $r_a(T_e)$  is simply the ratio of the number of reference slips detected in advance and the total number of reference slips. Fig. 6.8 shows these two metrics as functions of  $T_e \geq T_e^c$ .

We immediately observe in Fig. 6.8(a) that the  $W_2$  measure tends to provide the earliest detection times for a wide range of error thresholds. In particular, for  $T_e = T_e^c$ , the median advanced notice is over ten frames, while the best value that can be achieved by  $f_{plr}$  is only around five, and  $v_x$  barely detects slip events in advance at all.

Fig. 6.8(b) shows that  $r_a$ , as expected, decreases with  $T_e$  for each considered measure. Once again, the shape of the  $W_2$  curve is different from the other two. This




Figure 6.9: The median magnitude of the wall velocity at the detection time for different measures and  $T_e$ . For each measure, the solid lines mark the median of  $v_x$ , and the shaded regions mark one standard deviation. The dashed lines show the two thresholds used by the offline method: (black)  $v_L$  used to identify slip events and (red)  $v_S$  used to identify their start.

curve decreases slower and is the only one that approaches unity as  $T_e$  gets close to  $T_e^c$ . That is, this measure is the only one that can predict almost all slips in advance.

Lastly, we investigate the wall velocity at detection times over different measures as further confirmation that the  $W_2$  metric can detect slips before the wall has begun to move. Fig. 6.9 shows the median (and one standard deviation) of  $v_x$  at detection times as a function of  $T_e$ . The median at each fixed  $T_e$  value is taken over all slip events detectable by the given model and  $T_e$ . As expected, all the curves are increasing with  $T_e$ . Consistent with the results shown in Fig. 6.8, we find significantly lower values of  $\langle v_x \rangle$  when considering the  $W_2$  measure. That is, persistent homology indeed detects changes in the force network well before the wall has begun to move appreciably.

## Chapter 7

### Conclusions and future work

In this manuscript, we present a method for detecting the onset of slip that can be used for an incoming stream of data. We consider three different measures on the granular system: the wall velocity  $v_x$ , maximum percolation force of the differential force network,  $f_{plr}$ , and the zero-level second Wasserstein distance between the persistence diagrams quantifying the force networks,  $W_2$ . For each of these measures, we build a dynamical linear model (DLM) capable of accurately predicting its behavior during the stick regime. Naturally, the predicted values of the models become increasingly inaccurate leading up to and during slips. Hence, to detect the upcoming slip we analyze the predictive error: the differences between the predicted and observed values of these models.

Our analysis of  $v_x$  shows a clear classification between slip and micro-slip detections. We find that global changes in the network, as measured by  $f_{plr}$ , are nearly always followed by a subsequent slip or a micro-slip. So, the  $f_{plr}$  model provides a slightly earlier prediction of the upcoming slip events than that of the wall velocity. Next, by using the  $W_2$  measure, we can identify local changes in the force network which can either dissipate or spread over the force network, culminating in a slip or micro-slip. While micro-slips were reported in the previous works considering similar systems, see e.g. [10], and local changes (called ‘local avalanches’) were discussed recently as well [4], we are not aware of such events being used for predictive purposes. From the predictive perspective, we note that the  $W_2$  model shows the best

performance, in terms of sensitivity. On the other hand, the high sensitivity of this measure also leads to a significant percentage of ‘false positives,’ in the form of local changes that do not lead to slip events.

In summary, we observe the following timeline of the changes leading to a slip event. Typical slip events start with a local change in the force network, the magnitude of which can be roughly quantified by the model error of the  $W_2$  DLM. If the size of the local change is sufficiently large, then it always becomes global and triggers a slip. On the other hand, a sufficiently small local change does not trigger a slip. There is, however, a large range of local change sizes for which the outcome is uncertain; the initial local disruption may either dissipate or trigger a global change resulting in a micro-slip or even slip event. While the chronology of a slip described above has been discussed in the literature already, our method allows for precise quantification of this timeline.

### 7.1 Future avenues of research

As noted above, a significant fraction of the local changes of the force network fall into an intermediary range and may or may not result in a slip or micro-slip event. Naturally, it is of great interest to explore whether there are distinguishing characteristics of local changes that could be used to predict if the given change will lead to a slip event. One potential route is the analysis of individual dynamic components of the DLMs, similar to the component-wise analysis of models in [34].

Figure 7.1 isolates the amplitude of the periodic component(s) of the individual DLMs, focusing on the start of the detected events. In separating the slips, micro-slips, and local changes, there do appear to be subtle differences between these categories, even before detection. However, at the time of detection, the distribu-



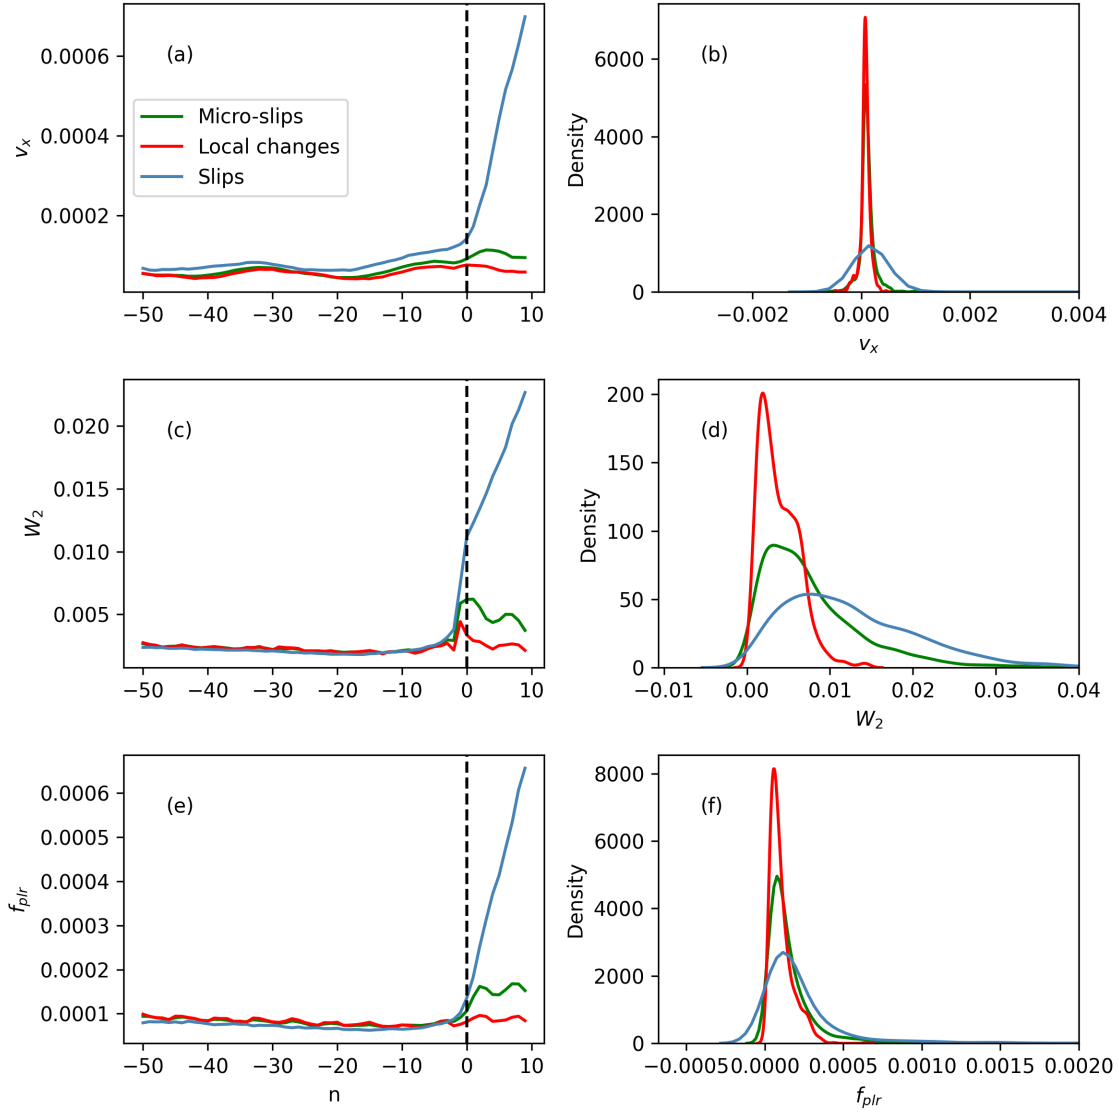


Figure 7.1: Median amplitude of the periodic component(s) of the DLMs for (a)  $v_x$ , (c)  $W_2$ , and (e)  $f_{plr}$  at the frames near detection (marked by  $n = 0$ ), taken over detections. The medians are separated over slips (blue), micro-slips (green), and local changes (red). Panels (b), (d), and (f) show the full distributions of amplitude values at detection time. The detection time is from the  $W_2$  model at its critical threshold  $T_e^c = 0.4$ .

tions largely overlap, which impedes the ability to distinguish the detection category. Nevertheless, these plots suggest this type of analysis could be useful in future work.

Another potential avenue is the examination of other measures. In this work, we

use  $W_2$  to identify local changes. This measure compares the topological properties of the force network encoded by the persistence diagrams. Hence, the precise geometry of the force network is ignored. Further research will be necessary to establish if the geometry of the network, or even more detailed information on the particles, can be used to predict the outcome of local changes. Our results suggest that the DLM framework can be a powerful tool when considering the appropriate measures on the granular system.

## Bibliography

- [1] David Acheson. *Elementary Fluid Dynamics*. 1990.
- [2] E. Aharonov and D. Sparks. “Stick-slip motion in simulated granular layers”. In: *Journal of Geophysical Research: Solid Earth* 109.B9 (2004).
- [3] Lucilla de Arcangelis et al. “Statistical Physics Approach to Earthquake Occurrence and Forecasting”. In: *Physics Reports* (2016).
- [4] Jonathan Barés et al. “Local and global avalanches in a two-dimensional sheared granular medium”. In: *Physical Review E* 96 (Nov. 2017). DOI: 10.1103/PhysRevE.96.052902.
- [5] Andrea Bizzarri, Alberto Petri, and Andrea Baldassarri. “Earthquake dynamics constrained from laboratory experiments: new insights from granular materials”. In: *Annals of Geophysics* 64 (Nov. 2021), SE441. DOI: 10.4401/ag-8613.
- [6] W. F. Brace and J. D. Byerlee. “Stick-slip as a mechanism for earthquakes”. In: *Science* 153.3739 (1966), pp. 990–992.
- [7] P. Bretz, Lou Kondic, and M. Kramar. “Stochastic methods for slip prediction in a sheared granular system”. In: *Physical Review E* 107 (May 2023). DOI: 10.1103/PhysRevE.107.054901.
- [8] David Cohen-Steiner et al. “Lipschitz Functions Have  $L_p$ -Stable Persistence”. In: *Foundations of Computational Mathematics* 10 (Feb. 2010), pp. 127–139. DOI: 10.1007/s10208-010-9060-6.

- [9] Fergal Dalton and David Corcoran. “Self-organized criticality in a sheared granular stick-slip system”. In: *Physical review. E, Statistical, nonlinear, and soft matter physics* 63 (July 2001), p. 061312. DOI: 10.1103/PhysRevE.63.061312.
- [10] Lucilla De Arcangelis et al. “Micromechanics and statistics of slipping events in a granular seismic fault model”. In: *Journal of Physics: Conference Series* 319 (Sept. 2011), p. 012001. DOI: 10.1088/1742-6596/319/1/012001.
- [11] J. A. Dijksman et al. “Characterizing granular networks using topological metrics”. In: *Phys. Rev. E* 97 (2018), p. 042903.
- [12] Geir Evensen. *Data Assimilation: The Ensemble Kalman Filter*. New York, NY: Springer, 2009. ISBN: 978-3-642-03710-8.
- [13] Behrooz Ferdowsi and Allan Rubin. “A Granular Physics-Based View of Fault Friction Experiments”. In: *Journal of Geophysical Research: Solid Earth* 125 (June 2020). DOI: 10.1029/2019JB019016.
- [14] Clark Givens and Rae Shortt. “A Class of Wasserstein Metrics for Probability Distributions”. In: *The Michigan Mathematical Journal* 31 (Jan. 1984). DOI: 10.1307/mmj/1029003026.
- [15] Yu Guo and Jennifer Curtis. “Discrete Element Method Simulations for Complex Granular Flows”. In: *Annual Review of Fluid Mechanics* 47 (Jan. 2015), pp. 21–46. DOI: 10.1146/annurev-fluid-010814-014644.
- [16] H. M. Jaeger, S. R. Nagel, and R. P. Behringer. “The physics of granular materials”. In: *Physics Today* 49 (1996), p. 32.
- [17] P. A. Johnson et al. “Acoustic emission and microslip precursors to stick-slip failure in sheared granular material”. In: *Geophys. Res. Lett.* 40.21 (2013), pp. 5627–5631.

- [18] Paul Johnson and X. Jia. “Nonlinear dynamics, granular media and dynamic earthquake triggering”. In: *Nature* 473 (Oct. 2005).
- [19] Paul Johnson et al. “Effects of acoustic waves on stick-slip in granular media and implications for earthquakes”. In: *Nature* 451 (Feb. 2008), pp. 57–60. DOI: 10.1038/nature06440.
- [20] Karl-Rudolf Koch. *Introduction to Bayesian Statistics*. New York, NY: Springer, 2007. ISBN: 978-3-540-72723-1.
- [21] Lenka Kovalcinova, Arnaud Goulet, and Lou Kondic. “Percolation and jamming transitions in particulate systems with and without cohesion”. In: *Physical Review E* 92 (Feb. 2015). DOI: 10.1103/PhysRevE.92.032204.
- [22] Miro Kramar et al. “On intermittency in sheared granular systems”. In: *Soft Matt.* 18 (2022).
- [23] Miroslav Kramár et al. “Quantifying force networks in particulate systems”. In: *Physica D: Nonlinear Phenomena* 283 (Nov. 2013). DOI: 10.1016/j.physd.2014.05.009.
- [24] G. M. Ljung and G. E. P. Box. “On a measure of lack of fit in time series models”. In: *Biometrika* 65 (1978), pp. 297–303.
- [25] Craig Maloney and Anaël Lemaître. “Universal Breakdown of Elasticity at the Onset of Material Failure”. In: *Physical review letters* 93 (Dec. 2004), p. 195501. DOI: 10.1103/PhysRevLett.93.195501.
- [26] Yuriy Mileyko, Sayan Mukherjee, and John Harer. “Probability measures on the space of persistence diagrams”. In: *Inverse Problems - INVERSE PROBL* 27 (Dec. 2011). DOI: 10.1088/0266-5611/27/12/124007.

- [27] Giovanni Petris, Sonia Petrone, and Patrizia Campagnoli. *Dynamic Linear Models with R*. New York, NY: Springer, 2009. ISBN: 978-0-387-77237-0.
- [28] Massimo Pica Ciamarra et al. “Unjamming Dynamics: The Micromechanics of a Seismic Fault Model”. In: *Physical review letters* 104 (June 2010), p. 238001. DOI: 10.1103/PhysRevLett.104.238001.
- [29] C. Scholz. *The Mechanics of Earthquakes and Faulting*. Jan. 2019. DOI: 10.1017/9781316681473.
- [30] S. S. Shapiro and M. Wilk. “An analysis of variance test for normality (complete samples)”. In: *Biometrika* 52 (1965), pp. 591–611.
- [31] E. Stein and R. Shakarchi. “Fourier Analysis: An Introduction”. In: *Fourier Analysis: An Introduction* (Jan. 2003).
- [32] Lawrence D. Stone et al. *Bayesian Multiple Target Tracking*. Norwood, MA: Artech House, 2014. ISBN: 978-1-60807-553-9.
- [33] Antoinette Tordesillas, Zongzheng Zhou, and Robin Batterham. “A Data-Driven Complex Systems Approach to Early Prediction of Landslides”. In: *Mechanics Research Communications* (2018).
- [34] H. Wang, Q. Qian, and A. Tordesillas. “Modeling Big Spatio-Temporal Geo-Hazards Data for Forecasting by Error-Correction Cointegration and Dimension-Reduction”. In: *Spat. Stat.* 36 (2020).
- [35] Afra Zomorodian and Gunnar Carlsson. *Computing Persistent Homology*. Springer, 2004.