The Epistemological Limitations of Google's Knowledge Graph

In 2012, Google introduced the Knowledge Graph, a computer system that instead of providing search results, provides information—what Google calls "knowledge." Now, when people go to Google and search for "Leonardo da Vinci," they see a little box pop up next to the traditional search results, outlining da Vinci's vocation, his birth date, his siblings, his artworks, and so on. While a good deal of work has been done on the broader ethical implications of Google in general ("Search Engines and Ethics"), less has been done on the epistemological implications on the Knowledge Graph in particular—how does the Knowledge Graph affect the landscape of what can be known, and how does it change cultural assumptions about the nature of knowledge? Alexander Monea, in a 2016 paper, made strong headway into the topic. Examining the Knowledge Graph, he argued that "the fundamental data structure of the 'triple,' in essence a subject-predicate-object statement, constitutes a problem immanent to the database itself" (452). Using the "perspective of media theory, philosophy of difference, and epistemology" (452), he demonstrated that the "structure of the 'triple'" (452) sets certain limits on the types of knowledge that graph databases can represent. Taking his work as a starting point, and looking towards speech act theory and Mary Poovey's *A History of the Modern Fact* for theoretical support, I argue that the Knowledge Graph is predicated on specific epistemological assumptions unique to the past five-hundred years or so, assumptions that take it on faith that knowledge consists of facts that can "exist in the world like pebbles, waiting to be picked up" (Poovey 1). Furthermore, I contend that the Knowledge Graph, through Michel Foucault's notion of the "will to truth," limits what we can know by promoting a narrow definition of "the true discourse" (Foucault 54).

In his *Nichomachean Ethics*, Aristotle claimed that knowledge could be divided into two general kinds: the scientific, involving "things whose originative causes are invariable" like science, and the calculative, involving the consideration of "variable things" like how to act when faced with a difficult moral choice (*NE,* Book VI, Chapter I). Within the scientific-calculative distinction, he went on to divide knowledge into five more specific types: art, the knowledge of how to make things; science, the knowledge of "things that are universal and necessary;" practical wisdom, the knowledge of "what is good and expedient" and of "what is to be done;" philosophical wisdom, the knowledge of "the things that are highest by nature" (e.g. God); and intuitive reason, the knowledge of how to go about scientific inquiry (Book VI, Chapters 4-6). Science and philosophical wisdom, according to Aristotle, fall under the category of scientific knowledge, while art, practical wisdom, and intuitive reason fall under the category of calculative knowledge. In my view, the distinction between scientific and calculative knowledge can be more fully understood when paralleled with with the constative-performative distinction made by speech act theory.

When J.L. Austin first developed speech act theory, many philosophers believed that "the sole business, the sole interesting business, of any utterance—that is, of anything we say—[was] to be true or at least false" (Austin 233). They were only interested in utterances like "the cat is on the mat," or "God exists," or "water boils at 212 degrees fahrenheit," and saw other utterances as nonsense. Austin, in developing his theory of speech-acts, points out that in fact people often say things which look "like statement[s] and grammatically, I suppose, would be classed as statement[s], which [are] not nonsensical, and yet [are] not true or false" (Austin 235). He claims that "if a person makes an utterance of this sort we should say that he is doing something rather than merely saying something" (Austin 235). He discusses a few examples:

> Suppose…that in the course of a marriage ceremony I say, as people will, 'I do'—(sc. take this woman to be my lawful wedded wife). Or again, suppose that I tread on your toe and say 'I apologize'. Or again, suppose that I have the bottle of champagne in my hand and say 'I name this ship the Queen Elizabeth'. Or suppose I say 'I bet you sixpence it will rain tomorrow'. In all these cases it would be absurd to regard the thing that I say as a report of the performance of the action which is undoubtedly done—the action of betting, or christening, or apologizing. We should say rather that, in saying what I do, I actually perform that action. When I say 'I name this ship the Queen Elizabeth' I do not describe the christening ceremony, I actually perform the christening; and when I say 'I do' (sc. take this woman to be my lawful wedded wife), I am not reporting on a marriage, I am indulging in it. (Austin 235)

He calls these kinds of utterances "performative utterances," and opposes them to what he calls "constative utterances," which seek simply to describe a certain state of affairs and can be classed as true or false (Austin 235). Aristotle's two types of knowledge—scientific and calculative—divide nicely into constative knowledge and performative knowledge, as expressions of scientific knowledge (e.g. "two things that are equal to the same are equal to each other" or "all objects fall at the same rate") tend to be enunciated through constative utterances, while expressions of calculative knowledge (e.g. "shape the clay with your hands as the pot spins around" or "don't get caught with your hand in the cookie jar") tend to be enunciated through performative utterances.

In *A History of the Modern Fact,* Mary Poovey outlines a very particular type of constative knowledge, one that I would like to highlight in my discussion of the

Knowledge Graph. This type of knowledge, which she names "the modern fact," consists of "noninterpretive…descriptions of particulars"—facts—connected through "systematic claims that were somehow derived from those particularized descriptions" (xii). The stat sheet for a baseball game provides a clear example. Each player's RBIs, home runs, errors, etc… are recorded as simple numbers—"noninterpretive…descriptions of particulars"—and then become meaningful through systematic claims about how they measure players' abilities. A player with a .400 batting average, for example, seems better than one with a .250 batting average, because of the systematic claim that batting average, a description of a particular, indicates something important about a player. One of the most salient epistemological limitations of the Knowledge Graph, I contend, is that it seems to define knowledge not in the broad terms Aristotle sets forth, or in the more schematic but still useful terms speech act theory offers, but in the narrow sense that Poovey's book delineates. The Knowledge Graph's knowledge, as I will show, consists precisely of particulars connected through systematic claims.

Before examining the epistemological assumptions present in the Knowledge Graph, I think it is necessary to first understand the way in which it works. For business reasons, Google does not make the code for the Knowledge Graph public, but even so, much can be learned from papers Google employees have published related to the subject, and from graph database design in general. To give a high-level overview, the Knowledge Graph stores its "knowledge" in "triples," each containing a subject, a predicate, and an object (Dong 2). For example, that the professional hockey player Barry Richter was born in Madison, Wisconsin would be stored as "<Barry Richter, born

in, Madison>" (Murphy). One big problem the Knowledge Graph has to solve is the one of getting true information into the database; to enter the eighteen billion triples currently in the database by hand would not be possible. To solve the problem, it appears that the Knowledge Graph—put simply and perhaps overly schematically—takes a two-step approach. First, it crawls existing knowledge databases (e.g. Wikipedia, CIA World Factbook, Freebase) and web pages to build up a list of what are known as "entities"—things like "Barry Richter" or "Madison." Entities in hand, it then figures out, again by crawling existing knowledge databases and web pages, how those entities relate to each other. In Poovey's language, it first searches the web for "particulars," then tries to figure out the "systematic claims" that connect them (xii). To give a simple example, the Knowledge Graph might in an initial search come across the entities "Barry Richter," "Madison," and "UW-Madison," and put them in the system. Later, it might come across the sentence "In the fall of 1989, Richter accepted a scholarship to the University of Wisconsin, where he played for four years and earned numerous individual accolades" (Murphy). In analyzing the sentence, it would infer a number of things about Richter and the University of Wisconsin, including the triple "<Barry Richter, studiedAt, UW-Madison>." To avoid taking in false information, the Knowledge Graph would not add the new triple right away. First, it would check to see what it already "knew" about Richter; it might find that it already "knew" the triples "<Barry Richter, born in, Madison>" and "<Barry Richter, lived in, Madison>." In light of these triples, the Knowledge Graph would infer that "<Barry Richter, studiedAt, UW-Madison>" was probably true, and officially add it to the graph (Murphy).

In practice, the Knowledge Graph has worked incredibly well. Almost all of the triples that get added to the Knowledge Graph turn out to represent true statements (e.g. Barry Richter studied at UW-Madison). Moreover, with its introduction, websurfers no longer have to spend laborious hours searching for simple pieces of information scattered across the web; if a person wants to know how many children Barry Richter has, she can now get an immediate answer. However, even with the Knowledge Graph's clear success as a way of making information accessible, it suffers from epistemological limitations. Since its introduction in 2012, one scholar already has expounded upon some of those limitations. In his criticism of the system, Alexander Monea points out that "[w]hat exists for Google is strictly…that which can be abstracted from its context into the numerical form of a triple" (455). He goes on to argue, based on the ideas of Gilles Deleuze, that "potentiality," or the potential for the genesis of knowledge, hinges on the understanding of the virtual, which "is always-already…existing alongside the actual, and can be envisioned as a great plane populated with 'nonnumerical multiplicities'" (458). The Knowledge Graph, he argues, cannot represent these "*nonnumerical* multiplicities" of the virtual because it depends on *enumeration* through the unit of the triple (458). Furthermore, because it depends on "statistically enumerable genera [triples], the Knowledge Graph is cut off from difference in itself, the play of the actual and the virtual" (459). In consequence, he says, "the nonrepresentational and affective milieus that so mark human worlds and ways of being in those worlds [i.e. nonnumerical multiplicities] risk becoming increasingly imperceptible as they are overcoded by representational enumerative schemas" like the

Knowledge Graph (460). In his view, the design of the Knowledge Graph makes a whole realm of knowledge unknowable. For instance, a well-read poet can know the sort of feeling they get when they read a good poem, but that feeling, because it cannot be enumerated in the form of the triple, cannot be known through the Knowledge Graph. For the same reason, it would seem that the four out of Aristotle's five types of knowledge which are not enumerable—art, practical wisdom, intuitive reason, and perhaps even philosophical wisdom—cannot be known through the Knowledge Graph.

While Monea points out an important epistemological issue with the Knowledge Graph, Google seems to wield a large enough influence on the way people think to warrant more attention. Betsy Sparrow et al., looking at four recent studies, highlight the general cognitive effects of technology, concluding that "processes of human memory are adapting to the advent of new computing and communication technology" (778). Introna and Nissenbaum, examing search engines more particularly, discuss the specific connection between technology and "social, political, and moral values" (181). "Philosophers of technology," they write, "have recognized the intricate connection between technology and values" (181). The connection they expound upon, I would argue, exists also between the Knowledge Graph and epistemological values. In doing research for this paper, for example, a quintessential act of knowledge production, I looked to Google a number of times. Through an exploration of Mary Poovey's concept of "the modern fact," then, it is my goal to expand the thus far limited literature on the Knowledge Graph and shed light on the ways in which it affects the way people think—and in particular, the way people think about knowledge.

In *A History of the Modern Fact*, Poovey tells the long story, starting in the fifteenth century, of the epistemological unit she calls "the modern fact." The modern fact, she claims, can first be seen operating in the fifteenth-century system of double-entry bookkeeping. With double-entry bookkeeping, merchants would write on one side of a ledger "interpretive" descriptions of goods traded and on the other side what looked like "preinterpretive or even noninterpretive" costs of the goods (xii). With this system, they became able to record, for example, that they had sold 100 pounds worth of sheep and received 100 pounds of the queen's gold in return, and thus know that all was in order. According to Poovey, "the double-entry system seemed to guarantee that the details it recorded were accurate reflections of the goods that had changed hands *because* the system was formally precise" (30). The numerical aspect of the double-entry system—and its juxtaposition with the descriptive aspect ("sold twenty sheep")—made what had previously been seen as interpreted parts of the world (e.g. the value of twenty sheep) seem "objective" and accurate, and made it seem possible, for the first time, to separate *interpretation* from *description*. This separation, this idea that some components of knowledge were "preinterpretive or even noninterpretive" (xii) and that others were "systematic" (xv), is hard to overemphasize, as it represented a large (and not inevitable) shift in epistemological belief. With Google's introduction of the Knowledge Graph, I argue, this shift has become formalized.

The Knowledge Graph's entities (e.g. Barry Richter) and predicates (e.g. studiedAt) correspond precisely to the "observed particulars" and "systematic claims" characteristic of the modern fact (Poovey xv). "In contrast to ancient facts," Poovey writes, “which referred to metaphysical essences, modern facts are assumed to reflect things that actually exist” (29). Entities, which include things like books and movies (Singhal) and according to Google “describe real-world entities like people, places, and things,” are indeed by definition “assumed to reflect things that actually exist” (Poovey 29)[1]. In addition, just as Poovey claims that modern facts, though they pretend to be “preinterpretive” (xii), are really “*both* observed particulars *and* evidence of some theory,” so too are Knowledge Graph entities really *both* observed particulars *and* evidence of some theory (8). Here, the word theory is being used in the sense of seeing the world a certain way. The observed particular that 100 pounds of sheep were sold, for example, hinges on a way of seeing the world that presupposes that sheep can be quantified into a monetary value. In the Knowledge Graph, the triple “<Barry Richter, studiedAt, UW-Madison>” pretends that the entities “Barry Richter” and “UW-Madison” are just preinterpretive observed particulars. However, it is in fact the case that the entities are employed in the Knowledge Graph as evidence of at least three theories: first, the theory that a walking, talking, two-legged organism can have a stable human identity (like “Barry Richter”); second, the theory that a gathering of buildings, programs, and people can constitute a university (like “UW-Madison”); and third, the more foundational theory that the two entities—Barry Richter and UW-Madison—can be disentangled from each other and then related in a meaningful way—that observed particulars are “different in kind from the analytic accounts that accompanied them” (Poovey xv).

[1] It is interesting to note that this parallels Monea’s claim that the Knowledge Graph cannot represent “the nonrepresentational and affective milieus that so mark human worlds” (460).

The division between the entities and the predicates so characteristic of the modern fact can be seen quite clearly through a visualization of the Knowledge Graph's underlying graph data structure. As mentioned earlier, the Knowledge Graph stores knowledge in <subject, predicate, object>, or rather <entity, predicate, entity>, triples. These triples, taken together, constitute a graph that looks something like Figure 1, or more realisically, Figure 2. In both visualizations, the entities are shown as circular

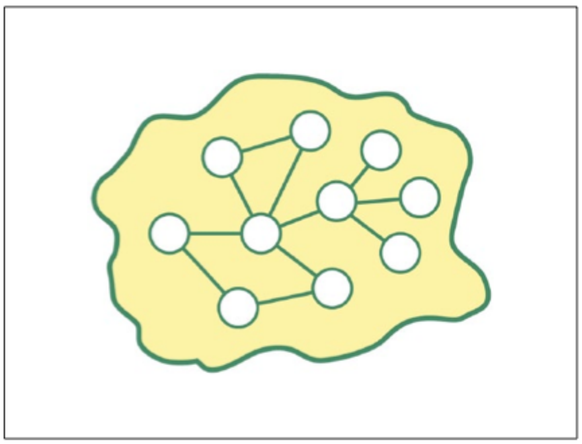

Figure 1: A Simple Graph (Monea 455)

Figure 2: A Complex Graph (Monea 455)

nodes and the predicates are shown as lines between them. Seen visually, it becomes clear that the Knowledge Graph treats entities as "observed particulars" able to be separated from both each other and the predicates, or "systematic claims," which connect them (Poovey xv). It is easy indeed to imagine one of the predicate-lines on the graph disappearing, leaving an entity to float off into space by itself, an event which in terms of the underlying data structure indeed requires simply changing a one to a zero[2]. And not just the entities' disconnectedness, but also the systematic claims which connect them, can be seen in the visualization. With a small bit of imagination, the predicate-lines can be visualized as forming branch-like substructures lurking within the greater tree of the graph; Barry Richter was not just born in Madison; he also lived in Madison and went to school at UW-Madison. Thus, the entities and predicates of the Knowledge Graph do work like the observed particulars and systematic claims of the modern fact.

To recapitulate, then: Poovey, in her book, argues that before the time of double-entry bookkeeping people believed that description could not be separated from interpretation, but that in the past five hundred years, a shift has taken place. People have begun to see knowledge as essentially divisible between "noninterpretive or even preinterpretive" facts and the systematic claims which connect them (Poovey xii). Google's Knowledge Graph formalizes and makes stark this division.

The modern fact operates at another level in the Knowledge Graph. Zooming out a little bit, one can think not of each entity, but of each triple as a modern fact. In this view, the "systematic claims" (Poovey xii) that connect the triples is the Knowledge Graph itself; the Knowledge Graph, through the underlying structure it imposes, and through the specific information-collection algorithms it uses, defines and limits the ways in which the triples connect. Thus, Barry Richter simultaneously exists as a series of

---

[2] The Knowledge Vault, which is used to populate the Knowledge Graph and probably shares its underlying data structure, stores triples in a "sparse E x P x E 3d matrix G, where E is the number of entities, P is the number of predicates, and G(s,p,o) = 1 if there is a link of type p from s to o, and G(s, p, o) = 0 otherwise" (Dong 2). To remove a predicate p, one must simply change G(s,p,o) from one to zero.

disconnected, preinterpretive triples (<Barry Richter, studiedAt, UW-Madison>, <Barry Richter, born in, Madison>, and <Barry Richter, lived in, Madison>), and also, through the structure and algorithms of the Knowledge Graph, as a person who was born in Madison, lives in Madison, and went to school at UW-Madison. In this view, too, the Knowledge Graph seems permeated by the modern fact. And here, even more than in the previous view, it becomes clear that the "facts" delineated by the Knowledge Graph are *not* preinterpretive. Barry Richter's living in Madison, for example, clearly hinges on a cultural conception of residency, which requires a system of territoriality, and law, and so on—all things which seem even more clearly than "Barry Richter" to require interpretation. Yet the Knowledge Graph, because it builds into its DNA the notion of the modern fact, makes all of its triples, its information, seem preinterpretive.

If the Knowledge Graph does assume and formalize into a computer system the notion of the modern fact, it is worth wondering what epistemological ramifications may follow. As Poovey notes, "modes of representation inform what we can know" (xv). Though I cannot hope to give a full account here of the limitations introduced by the modern fact's permeation of the Knowledge Graph, I do think that I can start such an account. First, as Poovey points out, people tend to see the modern fact as a "unit of value- and theory-free representation available for producing systematic knowledge about the social and natural worlds" (xxv). The Knowledge Graph's use of the modern fact might therefore cause people to forget that all knowledge is, in the last analysis, interpreted[3]. Furthermore, in the seventeenth century, the modern fact, because it seemed theory- and value-free, began to exercise social authority. "[Robert] Boyle," for example, "and…other members of the [Royal] Society, propped their own social authority on the claim that facts were theory- and value-free" (Poovey xviii). The discipline of accounting, due to its use of double-entry bookkeeping, eventually "came to exercise social authority" (33). In the same way, Google's Knowledge Graph, because it is permeated by the modern fact and indeed seems to present knowledge as "just the facts," devoid of theories and values, has the potential to exert social authority. Almost as if by providence, I was walking by a pair of students the other day and overheard one, in the heat of an debate, say something like "we need to settle this now—let's Google it…"

The most severe, and perhaps insidious, ramification of the Knowledge Graph's formalization of the modern fact is suggested by the work of Michel Foucault. Foucault, in a 1970 lecture, talks about something he calls the "will to truth" (54). The will to truth, he writes, is "a system of exclusion, a historical, modifiable, and institutionally constraining system" (54). By dividing the "true" from the "false," the will to truth limits discourse to the "true discourse" (54). "For the Greek poets," he claims, the true discourse consisted of that which was "pronounced by men who spoke as of right and according to the required ritual," that "which dispensed justice and gave everyone his share" (54). "Yet already a century later," by the time of Plato, "the highest truth no

[3] N. Katherine Hayles, in a discussion of scientific inquiry, makes a strong case for this. In her view, "observables [cf. observed particulars] really mean observations made by humans located at specific times and places and living in specific cultures" (77).

longer resided in what discourse was or did, but in what it said," a trend that continues to this day (55). In a remarkable parallel, the division between what discourse does and what it says corresponds precisely to the division speech act theory makes between performative and constative discourse. The will to truth, then, functions today in the realm of constative discourse; in particular, it functions to exclude that constative discourse which says "false" things about the world. To give a simple example, the will to truth is seen functioning when one presidential candidate yells at another that their last statement "just isn't true." In so uttering, the candidate's goal is to limit their opponent's discourse through the will to truth. In addition to limiting false constative discourse, I would suggest that the will to truth limits performative discourse as well, since performative discourse not only is not true, but cannot be conceived of in terms of the true-false dichotomy at all. In the case of the twentieth-century philosophers to whom Austin was responding, for example, they claimed that performative discourse was "nonsense" and thus declared it unworthy of study.

Google's Knowledge Graph, because it is seen as an arbiter of truth, embodies and perpetrates the will to truth. Two students get into an argument, they google a contentious point, and they limit their discourse to that which the search results support, that which is "true." Because, as I have laid out, the Knowledge Graph assumes a very narrow conception of knowledge—as observed particulars connected through systematic claims—it limits what can be known, and thus limits what can be "true." Insofar as people look towards Google to separate the true from the false, then, they limit their discourse—and their knowledge with it—to the narrow field permitted by the

Knowledge Graph. Of Aristotle's five types of knowledge, they limit themselves to scientific knowledge representable through the particulars of the modern fact only, and de-emphasize art, practical wisdom, intuitive reason, and philosophical wisdom.

One might argue that the modern fact, along with the general assumption of knowledge as scientific in the Aristotelian sense, only permeates the Knowledge Graph in its current iteration. Future iterations, it could be argued, will be able to transcend the limitations in what the Knowledge Graph can know; new data structures and new algorithms will solve the challenge. Indeed, this seems to be a common enough view. In a talk given by a software engineer working on the Knowledge Graph's extraction team, the engineer remarked that "we are very well aware that these knowledge bases only have factual sort of declarative statements—that Barack was born in Hawaii. And they don't know anything about apples, or common sense, or stuff that kids know. And we want to solve that problem too. We have a team working on common sense knowledge extraction and reasoning, and that's a work in progress" (Murphy). While this seems like a laudable attitude, it is possible that there exist fundamental limitations to what computer-based knowledge systems like the Knowledge Graph can know, limitations stemming from the basic principles upon which computers are built.

Computers, since they have been built from scratch by humans, are understood fairly well—when I used to work on robotics, I would often hear the phrase "the robot is only doing what you told it to do," a true, if frustrating, statement. The mysteriousness often found in nature and in human systems is not found in computer science. To understand computers, computer scientists look towards the mathematical theory of automata. Automata theory, basically, models a computer as “a directed graph” consisting of a “set of states [or nodes] Q,” an “input alphabet A,” a “set of initial states I,” a “set of final states T,” and “a set of transitions E,” which is a subset of Q x A x Q and defines how inputs act on the automata (Sakarovitch 51). An automata—and in turn, a computer—can thus be specified as the graph $<Q, A, I, T, E>$. Put in simpler terms, automata theory explains that computers work by maintaining a certain internal state, waiting for human (or otherwise external) input, and then upon that input, moving to a new state. For example, a word processor—basically an automata, a computer, within a computer—works by maintaining an internal state constituted by the file that is open, the words that are visible, the panels that are shown and/or hidden, the location of the cursor, and so on. Once the human user presses some key, the ‘H’ key for instance, the word processor will then move to a new state, constituted the same way except that the words visible now include an extra ‘H’. At any given moment, every computer in the world could be said to have a specific, (theoretically) knowable, and certain state. Each works by moving from state to state in reaction to external inputs.

The states of a computer, I contend, correspond to the observed particulars of the modern fact. Though a full exploration of computers as formalizations of the modern fact is beyond the scope of this paper, I will point out that by definition, computer states have to be discrete, and as mathematical objects, numerical. Poovey, in her book, emphasizes multiple times that “numbers…epitomize…the modern fact” (4). Thus, it would at least seem that computers in general, like the Knowledge Graph, are permeated in their foundations by the modern fact.

To conclude, the Knowledge Graph seems to assume a narrow conception of knowledge, as solely scientific, or rather, as solely constative. In the terms Mary Poovey sets forth, the Knowledge Graph seems to assume that knowledge consists of facts that can "exist in the world like pebbles, waiting to be picked up" (1), as "noninterpretive… descriptions of particulars" connected by "systematic claims that [are] somehow derived from those particularized descriptions" (xii). If it is true that Google plays a large role in determining what constitutes knowledge and thereby what constitutes "truth", then it would seem that the Knowledge Graph thus limits knowledge, and discourse in general, to a narrow realm. And while it is laudable that computer scientists are working hard to expand the Knowledge Graph's reach, automata theory suggests that any such attempt will fail to escape the grips of the modern fact. That being said, the most promising avenue for computer science researchers may be that suggested by the work of Douglas Hofstadter. In his research, Hofstadter has tried to understand how human consciousness arises out of what he sees as the formal, or mathematical, system of physical reality. By understanding how consciousness arises in humans, he hopes to take the formal, mathematical, system of the digital computer, and use it to build an artificial consciousness. If he or his successors succeed, it may be possible for computers to represent epiphenomenal knowledge of the various types human can understand, but until that time, it may prove wise to look for knowledge in places outside of the Knowledge Graph.

References

Austin, J.L. How to do things with words. Oxford university press, 1975.

Dong, Xin, et al. "Knowledge vault: A web-scale approach to probabilistic knowledge fusion." Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, 2014.

Foucault, Michel. "The Order of Discourse." Untying the Text: A Post-structuralist Reader. Boston: Routledge & Kegan Paul, 1981. 48-78. Print.

Hayles, N. Katherine. "Constrained constructivism: locating scientific inquiry in the theater of representation." New Orleans Review 18.1 (1991): 76-85.

Introna, Lucas D., and Helen Nissenbaum. "Shaping the Web: Why the politics of search engines matters." The information society 16.3 (2000): 169-185.

Monea, Alexander. "The Graphing of Difference Numerical Mediation and the Case of Google's Knowledge Graph." Cultural Studies and Critical Methodologies 16.5 (2016): 452-461.

Murphy, Kevin, et al. "Knowledge Vault: a web-scale approach to probabilistic knowledge fusion." 20th International Conference on Knowledge Discovery and Data Mining. *VideoLectures.net*. 21 Nov. 2016, videolectures.net/ kdd2014_murphy_knowledge_vault.

Poovey, Mary. A history of the modern fact: Problems of knowledge in the sciences of wealth and society. University of Chicago Press, 1998.

Sakarovitch, Jacques. Elements of automata theory. Cambridge University Press, 2009.

Singhal, Amit. "Introducing the knowledge graph: things, not strings." Official google blog (2012).

Sparrow, Betsy, Jenny Liu, and Daniel M. Wegner. "Google effects on memory: Cognitive consequences of having information at our fingertips." science 333.6043 (2011): 776-778.