Computational History of Knowledge: Challenges and Opportunities

Manfred D. Laubichler, Arizona State University and Santa Fe Institute Jane Maienschein, Arizona State University and Marine Biological Laboratory Jürgen Renn, Max Planck Institute for the History of Science

Abstract: So far, the twenty-first century has been defined by an ever-increasing availability of digital data and substantial advances in computational methods. Taken together, these developments have already affected all aspects of our lives, including the ways research in the sciences and the humanities is conducted. This computational turn is often viewed with unease. But as this essay argues, it also offers exciting new perspectives for the history of knowledge. Rather than fighting these trends, the essay suggests, by embracing new possibilities and actively participating in the development of new computational methodologies the history of knowledge can act as a bridge between the world of the humanities, with its tradition of close reading and detailed understanding of individual cases, and the world of big data and computational analysis. We can gain novel perspectives on the evolution of knowledge that are both detailed and broad.

THE DIGITAL CONTEXT

Our current "digital age" of big data, artificial intelligence, and automation triggers strong responses, often evoking both skepticism and fear.¹ This is true within the history and philosophy of knowledge as well. Digital innovations offer important research challenges and

Manfred D. Laubichler is President's Professor of Theoretical Biology and History of Biology and Director of the Global Biosocial Complexity Initiative at Arizona State University and a Professor at the Santa Fe Institute. His main research question is how something genuinely novel can emerge in evolution and history. School of Life Sciences, Arizona State University, Tempe, Arizona 85287-4501, USA; Manfred.Laubichler@asu.edu.

Jane Maienschein is University Professor, Regents Professor, and President's Professor at Arizona State University, where she directs the Center for Biology and Society. She is also a Fellow at the Marine Biological Laboratory in Woods Hole, Massachusetts. Her research explores the history and philosophy of developmental biology; her books include *Embryos under the Microscope: Diverging Meanings of Life* (Harvard, 2014). At the MBL, she codirects a project on understanding regeneration across the scales of life. School of Life Sciences, Arizona State University, Tempe, Arizona 85287-4501, USA; maienschein@asu.edu.

Jürgen Renn is Director at the Max Planck Institute for the History of Science in Berlin. His research focuses on structural changes in systems of knowledge; the aim is to develop a theoretical understanding of knowledge evolution, taking into account its epistemic, social, and material dimensions. Max Planck Institute for the History of Science, Boltzmannstraße 22, 14195 Berlin, Germany; renn@mpiwg-berlin.mpg.de.

¹ For some recent discussions on this topic see Sean Gerrish, *How Smart Machines Think* (Cambridge, Mass.: MIT Press, 2018); Bruce Schneier, *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World* (New York: Norton, 2015); Eric Siegel, *Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die* (Hoboken, N.J.: Wiley, 2013); and Jacob Silverman, *Terms of Service: Social Media and the Price of Constant Connection* (New York: Harper, 2015).

Isis, volume 110, number 3. © 2019 by The History of Science Society. All rights reserved. 0021-1753/2019/0110-0005\$10.00.

opportunities. We can discover new alliances, transformations of many research fields and the emergence of new ones, new career structures, novel ways of funding and organizing science and technology, new power structures (those who own the data control the science), and diverse public responses to these developments.² In many ways, recent developments in science and technology are a bonanza for all dimensions of science studies (history, philosophy, social studies, policy). Yet as we accept and adopt new methodologies and technologies, such as artificial intelligence (AI), deep learning algorithms, large datasets, complex systems approaches, and so on, the practice and the conceptual and epistemological foundations of the history of knowledge also undergo transformation in ways we cannot fully appreciate at the current moment and that evoke fear and skepticism for some.

The goal of this Isis Focus section and in this essay is to point to opportunities and invite participation.³ We do that by highlighting elements that, taken together, represent first steps toward what can be called computational history of knowledge. In one sense, computational history of knowledge follows a larger trend in digital humanities, which began as a way to make sources more readily available through large-scale digitization, to harvest so-called digitally born data, including social media data, and to provide sophisticated visualizations and annotations for these data.⁴ Novel ways of analyzing data were soon added to the mix. These include statistical methods, computational linguistics, and various forms of network and complex systems approaches. We can interpret this development as a normal case of methodological advancement. Digital data enable new types of analysis and representations; and, depending on the degree of their computational sophistication, historians can take advantage of these methods. But the questions and overall conceptual structure of the work remained grounded within standard historiography. Digital sources and computational methods have thus supported traditional historical work. Examples of such projects are digital editions, bibliographies, interactive databases, and genuinely digital projects, such as the Embryo Project Encyclopedia at Arizona State University.5

Another, more forward-looking, way to view the emergence of computational humanities is to localize this trend within broader developments related to data, information, computation, digitization, and automation that are currently transforming societies and economies at a global scale.⁶ These approaches include transformation of the ways that knowledge is produced, assessed, and transmitted.⁷ Applying computational methods to understand the history of knowledge is more than just doing traditional history with new tools. It brings novel conceptions

² Viktor Mayer-Schönberger and Kenneth Cukier, *Big Data:* A *Revolution That Will Transform How We Live, Work, and Think* (Boston: Houghton Mifflin Harcourt, 2013); Eric Schmidt and Jared Cohen, *The New Digital Age: Reshaping the Future of People, Nations, and Business* (New York: Knopf, 2013); David Weinberger, *Too Big to Know: Rethinking Knowledge Now That the Facts Aren't the Facts, Experts Are Everywhere, and the Smartest Person in the Room Is the Room* (New York: Basic, 2011); and Tim Wu, *The Master Switch: The Rise and Fall of Information Empires* (New York: Knopf, 2010).

³ For an earlier essay on this topic published in *Isis* see Manfred D. Laubichler, Jane Maienschein, and Jürgen Renn, "Computational Perspectives in the History of Science: To the Memory of Peter Damerow," *Isis*, 2013, 104:119–130.

⁴ Christine L. Borgman, Scholarship in the Digital Age: Information, Infrastructure, and the Internet (Cambridge, Mass.: MIT Press, 2007); Matthew K. Gold, Debates in the Digital Humanities (Minneapolis: Univ. Minnesota Press, 2012); Franco Moretti, Graphs, Maps, Trees: Abstract Models for a Literary History (London: Verso, 2005); Moretti, Distant Reading (London: Verso, 2013); and Susan Schreibman, Ray Siemens, and John Unsworth, A Companion to Digital Humanities (Malden, Mass.: Blackwell, 2004).

⁵ https://embryo.asu.edu.

⁶ Andrew McAfee and Erik Brynjolfsson, *Machine, Platform, Crowd: Harnessing Our Digital Future* (New York: Norton, 2017); Schmidt and Cohen, *New Digital Age* (cit. n. 2); and Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: Public Affairs, 2018).

⁷ Siegel, Predictive Analytics (cit. n. 1).

of what the nature of knowledge is and how it changes, how it is produced, and how it is connected to other societal domains. From this perspective, the concept of computational history becomes simultaneously far more interesting and also controversial.

What do we mean by computational humanities and what is its role in creating an emerging novel conception of knowledge? A full answer would take more than a single essay in a Focus section to develop and defend, and our understanding of the impact will continue to grow as we develop and apply the computational approach. We can gain insight from recent work in history and philosophy of science that recognizes the historical and contextual nature of knowledge. Scholarship in historical epistemology has demonstrated ways that the main organizing principles of knowledge, such as ideas about rationality, evidence, experimentation, observation, and confirmation, reflect the prevailing cultural conditions of past centuries.⁸ It is therefore logical to assume that the nature of knowledge will similarly be affected by current technological, methodological, and social transformations. Furthermore, these effects will apply not only to the natural sciences and technology but also to our understanding of knowledge will thus also be affected by currently emerging trends, including those of an increasingly computational view of reality. In this sense computational humanities will lead to new understandings of knowledge and its roles in both science and society.

Furthermore, if knowledge itself, as well as its contexts and their varied relations, is subject to historical change, then any conception of knowledge must be evolutionary, in the sense that understanding the co-evolutionary dynamics between knowledge and context must be an integral part of it.9 By "co-evolutionary dynamics" we are referring to some recent developments within evolutionary theory that emphasize the combined role of regulatory networks and niche constructions as well as the dynamic feedbacks between them. Within this framework, evolution is not used as a metaphor or a stand-in for history; rather, it directly addresses the mechanisms transforming knowledge. This realization raises profound questions about knowledge and how we can understand the changing past and present, even as our tools for understanding are also changing. The methods and data structures of computational history of knowledge provide a way to capture these dynamics. It is in that sense that we consider computational history of knowledge to be more than just a methodological development within digital humanities; it also brings with it a new epistemology and ontology of knowledge. It allows us to look both at the internal logic of knowledge and at the external contexts in which knowledge exists—something that in theory, if not in practice, has been at the core of the history of science throughout its own history. It also raises deep questions about how we can know the past without also appreciating the epistemological context for that past. Dynamic evolutionary approaches can help us address the evolution of knowledge and also the evolution of its underlying epistemic commitments.

⁸ Lorraine Daston, ed., Science in the Archives: Pasts, Presents, Futures (Chicago: Univ. Chicago Press, 2017); Daston and Peter Galison, Objectivity (New York: Zone, 2017); Daston and Elizabeth Lunbeck, Histories of Scientific Observation (Chicago: Univ. Chicago Press, 2011); Hans-Jörg Rheinberger, Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube (Stanford, Calif.: Stanford Univ. Press, 1997); Rheinberger, An Epistemology of the Concrete: Twentieth-Century Histories of Life (Durham, N.C.: Duke Univ. Press, 2010); Rheinberger, On Historicizing Epistemology: An Essay (Stanford, Calif.: Stanford Univ. Press, 2010); Rheinberger, eds., Kuhn's "Structure of Scientific Revolutions" at Fifty: Reflections on a Science Classic (Chicago: Univ. Chicago Press, 2016).

⁹ Jürgen Renn and Manfred D. Laubichler, "Extended Evolution and the History of Knowledge," in *Integrated History and Philosophy of Science: Problems, Perspectives, and Case Studies*, ed. Friedrich Stadler (Dordrecht: Springer, 2017), pp. 109–125; and Renn, *The Evolution of Knowledge: Rethinking Science for the Anthropocene* (Princeton, N.J.: Princeton Univ. Press, forthcoming).

COMPUTATIONAL HISTORY AND PHILOSOPHY OF KNOWLEDGE

Within that broader context, we discuss a number of challenges and opportunities that provide at least glimpses into what we consider to be the main epistemological dimensions of computational history.¹⁰ Computational history and philosophy of knowledge studies the creation, development, and dissemination of knowledge and the structure of knowledge representations through the lens of multidimensional and complex networks at different scales derived from large and diverse datasets. It takes a complex systems approach to the co-evolutionary dynamics of knowledge and problematizes the epistemological foundations of the processes of knowledge creation.

Computational history and philosophy of knowledge faces three significant types of challenges and opportunities.

Data

The first challenge is to provide adequate structured data. This requires both infrastructure and conceptual data models that represent essential features of knowledge structure and evolution (see also Julia Damerow and Dirk Wintergrün, "The Hitchhiker's Guide to Data in the History of Science," in this Focus section). Data can be heavily curated, such as the datasets in big repositories of science and those generated by digital humanities projects, or they can be largely unstructured, requiring more sophisticated tools and efforts to make them usable. The challenges related to data are epistemological: how to agree as to what relevant data are, how to obtain data, how to ascertain and evaluate data, how to structure data, and what data models to use. They are also social and institutional: what incentives there are to work on data curation, how to share data effectively in an open and nonproprietary way, and who provides the infrastructure for data management, storage, and long-term curation.

These are formidable challenges. But the benefits of widely connected structured data are huge. Early on in the development of the internet, the promise of the semantic web and, later, of the epistemic web as an evolving network of data and a new representation of knowledge was an important driving force that encouraged the scientific and scholarly community to contribute to this new world.¹¹ Today developments such as linked open data and its associated data and knowledge models continue these efforts. But we also have to admit that openness has largely given way to commercially motivated fragmentation: the new economy is largely built on the exploitation of data, as has recently been detailed by Shoshana Zuboff in her study of surveillance capitalism.¹² Such trends—seeking to exploit data to address a wide variety of different kinds of questions—also apply to large efforts of social engineering through the pervasive use of algorithms. And here it does not matter if these efforts are driven by state actors such as those we see in China, by markets, or by individual and group behaviors. The increased reliance on quantitative performance metrics at all levels of education and within the scientific process is not only remaking the social structures of science; it is also shaping what counts as knowledge, evidence, and acceptable domains of research.

Again, we are faced with the tensions mentioned earlier, the mix of apprehension and optimism. Clearly, we need to understand these trends and developments. But can we do so with

¹⁰ Throughout this essay we refer to computational history and philosophy of knowledge. However, we mainly discuss issues related to the history of science. For us, the link to philosophy of science results from our own sympathy with an emerging trend in HPS that emphasizes the integration of history and philosophy of science. See, e.g., the growing movement devoted to Integrated HPS: http://integratedhps.org/en/.

¹¹ Bo Leuf, The Semantic Web: Crafting Infrastructure for Agency (Chichester, West Sussex: Wiley, 2006); and Ian H. Witten et al., Data Mining: Practical Machine Learning Tools and Techniques (Burlington, Mass.: Morgan Kaufmann, 2011).

¹² Zuboff, Age of Surveillance Capitalism (cit. n. 6).

our traditional conceptual repertoire? Are our concepts and methods even remotely adequate for this new reality? Beyond these important larger questions, we see opportunities to connect different data repositories and thereby create a system where we can reuse the data that individual researchers have painstakingly discovered and curated, where we can collectively provide annotation and also open review. Such possibilities allow for innovations that we have only begun to imagine and to develop.

Patterns

One focus of big data–driven computational history of knowledge is the detection and analysis of patterns at different scales, especially larger scales that go beyond the capacities of what individual researchers or research groups could accomplish by traditional means. As the essays in this Focus section by Kenneth D. Aiello and Michael Simeone and by Deryc T. Painter, Bryan C. Daniels, and Jürgen Jost show, the analysis of such patterns involves a number of network representations, especially linking different types of networks (social, semantic, epistemic) at different scales, as well as statistical approaches such as topic modeling. A lot of new insights can be gained simply as a consequence of the ability to analyze research questions at larger scales with the aid of computational methods. But the detection of patterns is clearly a domain where new technologies such as machine learning, especially in its more advanced forms of deep learning algorithms, can be applied. Machines are very good at detecting patterns.¹³

The question is whether they are able to detect patterns that are relevant and that can contribute to our understanding of the evolution of knowledge. This is, for now, still an open question. Progress is being made, but the questions that historians ask are clearly hard from the point of view of a machine. One approach that we advocate is to see this research challenge within AI as an opportunity for computational history of knowledge. If we collaborate with computer scientists on the further development of pattern detection methods, we not only acquire powerful analytical tools; we can also come to a better understanding of how these algorithms work. In methodologically complex areas such as computational history it is important to know the possibilities, limitations, and biases of these methods in order to assess critically the kinds of results and insights they might be able to provide. But what can computational historians of knowledge offer AI researchers? The answer is obvious. We have well-structured data and, in many cases, thanks to generations of excellent scholars, a very good understanding of transformations of knowledge, as well as the likely causes for such transformations. Our datasets are therefore ideal training sets for deep learning algorithms as we challenge them to detect changes in patterns.

Such collaboration raises deep and foundational epistemological questions. If deep learning can eventually detect historical transformations, does it mean that these algorithms approach human reasoning and are not just getting results by crunching very large numbers? And given that human knowledge generation is inevitably a form of collective decision making, can AI systems become part of this collective? And if so, how? Again, we see both pragmatic and deep conceptual/epistemological reasons why engaging with these new technological developments can be beneficial for the history and philosophy of knowledge communities.

Dynamics

One goal of historical analysis is to explain underlying dynamics in the evolution of knowledge. This involves the analysis of time series of data as well as the development of causal models of

¹³ Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning (Cambridge, Mass.: MIT Press, 2016); and Shai Shalev-Shwartz and Shai Ben-David, Understanding Machine Learning; From Theory to Algorithms (New York: Cambridge Univ. Press, 2014).

scientific change. Traditionally historians have had an ambiguous relationship to explicit causal models, preferring the much richer form of narrative to account for the development and transformation of knowledge.

The challenges and opportunities of computational history of knowledge for understanding the dynamics of knowledge are twofold. Insofar as computational methods can analyze time series of data at a much larger scale than individual researchers can possibly accomplish (e.g., through the analysis of time series of large networks) and can also provide us with measures and metrics that can detect changes in patterns, they offer additional forms of evidence to support historical narratives. Knowing, for instance, how the network of Darwin's correspondence partners changed over time can support various narratives about the development of his ideas. Or analyzing in detail the correlations between what Darwin read, as we know from archival research, and how his thinking changed, as we know through the digitization of his notebooks, can support very detailed narratives about the development of his ideas.¹⁴

But can we go further? Again, this is an open question, but one that we believe is also worth exploring. Detailed datasets enable us to develop generative models about aspects of the scientific process. These might be more structural models that apply insights from complex systems science, such as general ideas about the nature of social networks and their dynamics. For example, researchers, including members of the Laubichler Lab for Computational HPS, have found evidence that scientific fields have what is called a "rich club" structure, where a highly connected core group tends to accumulate even more influence (measured by the degrees of connectivity within the network).¹⁵ And we can then observe that scientific innovation tends to happen more at the periphery of these networks than in the rich club (see below for more details).

Is this a general model for scientific innovation? We don't know, but we can test such hypotheses against more and more datasets. Another promising approach is the development of agent-based models that combine what we know about the behavior of participants in knowledge generation and dissemination with information about specific contexts (place, time, culture). Such models can, for instance, tell us something about how new ideas get accepted or how likely different forms of collaboration are. The main epistemological challenges related to these kinds of possibilities are connected to the question of whether general patterns and behaviors that drive the evolution of knowledge exist and, if so, how informative these are in explaining actual historical developments. What we want to suggest here is that this is an important question that cannot be decided *a priori*.¹⁶

WHAT HISTORICAL QUESTIONS CAN COMPUTATIONAL HUMANITIES ALLOW US TO ANSWER?

In this final section we illustrate our earlier reflections by presenting some concrete research questions and some results that demonstrate computational history of knowledge in practice.

¹⁴ Jaimie Murdock, Colin Allen, and Simon DeDeo, "Quantitative and Qualitative Approaches to the Development of Darwin's Origin of Species," Current Research in Digital History, 2018, 1, https://doi.org/10.31835/crdh.2018.14.

¹⁵ Julian J. McAuley, Luciano da Fontoura Costa, and Tibério S. Caetano, "The Rich Club Phenomenon across Complex Network Hierarchies," 2007, arXiv:physics/0701290.

¹⁶ We are aware that our belief that the search for general principles governing social and historical processes is a valid scientific approach is not shared by all historians or historians of science. We are not claiming here that such principles can be found for all aspects of social dynamics or the evolution of knowledge. But we are convinced that this is a question that can only be decided empirically. This requires us to try to see what can be discovered through the computational analysis of large datasets before we *a priori* rule out the possibility that such principles exist. For a recent overview of computational social science see Scott E. Page, *The Model Thinker: What You Need to Know to Make Data Work for You* (New York: Basic, 2018).

How Integrated Are Interdisciplinary Research Fields?

For several decades interdisciplinary research has been pushed by funding agencies, science administrators, and generations of well-intentioned scientists. Interdisciplinary research is needed, so the argument goes, because the problems we face-in medicine, environmental sciences, sociology, anthropology; the list can go on - are too complex to be mapped onto one traditional discipline. While the motivation for interdisciplinary research is clear, its actual success is less obvious. For starters, we don't quite know how to measure interdisciplinarity. We also have a difficult time distinguishing between different degrees of interdisciplinarity. Do we mean actual collaborations between scholars from different disciplines, or are we more interested in a combination of different conceptual and methodological approaches, perhaps in a single person's work? And how closely are those two layers linked? Does the successful application of different approaches require collaboration between scholars with different backgrounds? How can we tell whether any interdisciplinary approach is "better" than another, and in what ways? Traditionally, these questions have been addressed in the context of individual case studies, such as with breakthrough discoveries. While such narratives provide detailed insights into some localized scientific cultures, we have no way of answering questions about interdisciplinarity at a larger scale. Yet understanding across individual cases is exactly the kind of information we need if we want to retool the scientific enterprise toward greater degrees of interdisciplinarity.

In our ongoing study of evolutionary medicine (see also Deryc T. Painter, Bryan C. Daniels, and Jürgen Jost, "Network Analysis for the Digital Humanities: Principles, Problems, Extensions," in this Focus section) we asked these questions. We could do this because the group has a complete dataset (continuously growing and approaching twenty thousand items) of all publications in evolutionary medicine over the last four decades. Evolutionary medicine is an interesting case, as it was quite intentionally created as an interdisciplinary field by Randolph M. Nesse and George C. Williams through a conceptual essay and a successful book.¹⁷ Their argument was quite straightforward. Humans are the product of evolution; so are diseases. In order to understand and treat diseases more effectively, medicine needs to incorporate evolutionary perspectives. Conceptually the argument was easy to follow and quite convincing. Yet the actual scientific practice is another matter, and it is difficult to assess from individual reports whether and how evolution was actually making a difference. Looking at the large corpus of all publications claiming to incorporate evolutionary biology into medicine can give us an answer to the question, What difference did it make to bring evolution into medicine?

In analyzing the evolution of evolutionary medicine, we observe that the field steadily grew, that the pattern of its establishment follows that of many other scientific fields (from informal gatherings and interest groups to the establishment of a scientific society and a journal), and that it shows a growth rate standard for a new area of science. But just how interdisciplinary has it been? To address this question, we created different types of networks and quantitative metrics that crucially depend on our large dataset and computational analyses.

In a nutshell, here is what we have found. First, evolutionary medicine as a scientific community attracts researchers from different backgrounds. It includes people from medical fields as well as evolutionary biologists, with the latter dominating slightly. At the level of participation, then, the field is indeed interdisciplinary.

Second, publications in evolutionary medicine appear in a number of different journals. Both medical and evolutionary biology journals are well represented, and indeed evolutionary biologists are increasingly publishing their research in medical journals, while medical

¹⁷ George C. Williams and Randolph M. Nesse, "The Dawn of Darwinian Medicine," *Quarterly Review of Biology*, 1991, 66:1–22; and Nesse and Williams, Why We Get Sick: The New Science of Darwinian Medicine (New York: Times Books, 1994).

509

researchers are sending their work to journals of evolutionary biology. This again supports the claim of interdisciplinarity.

Third, when we analyze the conceptual evolution of evolutionary medicine an interesting pattern emerges. Looking at the whole corpus, we find a conceptual history that reflects some of the emerging narratives offered by practitioners in the field, with a gradual diversification of topics and concepts. But if we perform a more fine-grained analysis based on the different kinds of journals where research is published and the background of the researchers in question, we can identify distinct subdiscourses that follow their own logic and history and have their own communities. Network analysis allows us to quantify these differences (see the essay by Painter *et al.* in this Focus section).¹⁸ This result challenges some of the perceptions that evolutionary medicine is a unified interdisciplinary discourse. It also suggests that we need comparative analyses of other interdisciplinary fields in order to investigate whether these dynamics represent a common pattern. Of course, such analyses can only be done in the context of computational history, ideally based on linked databases.

Fourth, when we look at the patterns of collaboration in our corpus, we find that it is common for evolutionary biologists to collaborate with each other—and likewise for people with a medical background. However, the rate of collaborations between evolutionary biologists and medical researchers is around 1 percent of all publications, a challenging number for a self-professed interdisciplinary field.

These results of computational analyses allow us to frame the debate about interdisciplinarity more effectively by distinguishing different layers as well as providing quantitative evidence. Our analyses complement historical narratives of individual cases; they can detect both successes and failures and therefore provide a broader context for science policy.

Where Do Scientific Innovations Originate?

Another important set of questions in the history of knowledge concerns innovation.¹⁹ Understanding innovation is important not just for historians, but also for STS researchers and philosophers and administrators. What can computational approaches contribute? Here we focus on one specific aspect: Where do scientific innovations originate within the networks of science? To provide answers, we need complete data about a given scientific field in order both to construct social networks of science at different scales (see the essay by Painter *et al.* in this Focus section) and to describe the conceptual evolution of that field. For the latter, we need to analyze the full texts of all publications in the field. We also need to define what we mean by an innovation. Conceptually, we can distinguish between the emergence of a new idea represented by a scientific concept, a change in the meaning of an existing concept, or a reconfiguration of a conceptual network. For the purpose of this essay we focus on the first kind of innovation—the emergence of a new concept.

¹⁸ Further work along these lines will be appearing soon: Deryc Painter, Julia Damerow, and Manfred D. Laubichler, "The Evolution of Evolutionary Medicine," in *The Dynamics of Science: Computational Frontiers in History and Philosophy of Science*, ed. G. Ramsey and A. de Block (Dordrecht: Springer, forthcoming); and Painter, Bryan C. Daniels, and Laubichler, "Innovations Are Disproportionally Likely in the Periphery of a Scientific Network," *Theory in Biosciences* (forthcoming).

¹⁹ Innovation is a perennial topic in the history of science and technology. It also drives business investment, governmental action, and science policy. Consequently, both the field of science and technology studies and historians of science and technology have had a lot to say about this issue. Many of the recent debates have focused on the patterns of scientific innovation, following Kuhn's model of discontinuous scientific change. This has led, for example, to an emphasis on the details of knowledge transfer, the role of cultural and political contexts, and the role of technology in driving scientific change. There is a vast body of literature here, but we want to point selectively to Thomas S. Kuhn, *The Structure of Scientific Revolutions* (Chicago: Univ. Chicago Press, 1962); Richards and Daston, eds., *Kuhn's "Structure of Scientific Revolutions" at Fifty* (cit. n. 8); and John Ziman, ed., *Technological Innovation as an Evolutionary Process* (Cambridge: Cambridge Univ. Press, 2000).

By analyzing a complete corpus, we can easily pinpoint the first occurrence of a concept. But for a new term to count as an innovation it also needs to persist through time, something we can also measure. Here we are adopting Joseph Schumpeter's distinction between invention (first occurrence) and innovation (persistence and/or success). Again, using our study of evolutionary medicine as an example, we can detect a number of new concepts that by our definition qualify as innovations. This then raises the question, Where in the network of the field did these innovations originate?

As mentioned earlier, scientific fields generally display a very prominent structure of social networks—the rich club structure, where already highly connected members tend to become more connected through time (the rich getting richer). This pattern can be found for coauthorship and for co-citation. Evolutionary medicine is no exception. When we mapped the first occurrence of a conceptual innovation onto the network of evolutionary medicine (we used a co-citation network), we found that innovations tend to occur not within the rich club but at the periphery. This result is consistent with lots of anecdotal evidence, but it can be quantified only within a network-based representation of the whole field. We also found that innovations, defined as persisting inventions, diffuse rather rapidly into the conceptual structure used by the rich club, which can be used as one metric to define success. One example of such an innovative discovery is the role of nitric oxide during altitude acclimatization, a discovery first made by people at the periphery of evolutionary medicine that has rapidly become a core example of a complex evolutionary adaptation.²⁰

Computational analysis allows us to contribute to discussions about innovation with some very clear metrics and quantitative results that provide evidence for models that emphasize the role of the periphery. For future research, we can then also further characterize the periphery as well as the specific papers that contain conceptual innovations and ask if these have a higher or lower degree of conceptual diversity or bring together different kinds of ideas and sources, which might suggest some causal models of the innovation dynamics.

How Do the Social Mechanisms of Science Evolve?

Our final example is connected to the social organization of science, a prime application for computational analysis. There is a huge body of literature on social networks that provides the context for such analyses. We want to illustrate this domain of computational history of knowledge with a specific result from our analysis of the history of the microbiome concept.²¹

Microbiome research is one of the fastest growing areas of biomedicine. We analyzed the evolution of this concept on the basis of a complete collection of all microbiome papers published up to the present. In building the collection, we noticed an interesting pattern. As soon as the open-access journal *PLOS One* was established in 2008, it was the number one journal for microbiome research, as measured by number of publications. It was also the most conceptually diverse journal and introduced the highest number of conceptual innovations as we defined them above. The overall quality of the publications was also high, and we could not distinguish it from more established journals by a number of metrics, including citations. So, our first interpretation was that we had found a nice correlation between the availability of a new publication venue and a newly developing scientific research field.

²⁰ Allison J. Janocha et al., "Nitric Oxide during Altitude Acclimatization," New England Journal of Medicine, 2011, 365:1942– 1944, https://doi.10.1056/NEJMc1107887.

²¹ Kenneth D. Aiello, "Systematic Analysis of the Factors Contributing to the Variation and Change of the Microbiome" (Ph.D. diss., Arizona State Univ., 2018).

FOCUS

But once we also analyzed the abstracts of funded National Institutes of Health (NIH) projects—NIH is the major source of funding for microbiome research—we noticed additional correlations that led us to a new organizational mechanism in biomedical research. *PLOS One* publications have the highest correlation with funded NIH abstracts. This might at first seem surprising, as *PLOS One* is often regarded as a place for conceptual innovation and experimentation, two dimensions that many critics feel are not obviously associated with NIH panels. Then we looked at the funding numbers and found that, overall, microbiome funding has been growing more or less exponentially, which raises additional questions.

Together, these factors suggested a mechanism that can explain our computational results. Any successful NIH grant application needs to point to peer-reviewed publications. Traditional journals are often intrinsically conservative, generally space limited, and slow. *PLOS One* has a unique systems of peer review that is supportive of conceptual innovation, is not space limited, and is fast. So once *PLOS One* publications were accepted by NIH panels as sufficient proof of concept—evidenced by the conceptual correlation between *PLOS One* publications and funded NIH grants—a new positive feedback loop between NIH-funded microbiome research and *PLOS One* as a journal of record was established. And, since *PLOS One* is an open-access journal, this newly established link has a high potential for disruption. In short, this analysis revealed how the social mechanisms of science can evolve at rapid pace and how these new configurations in the organization of science can influence the subsequent dynamics of knowledge evolution. Given the amount of data we needed to analyze and the methods we needed to deploy, this phenomenon could only be discovered in the context of computational history of knowledge.

IN SUMMARY

These examples, as well as the more detailed accounts presented in the other essays in this Focus section, demonstrate that computational history and philosophy of knowledge is not just an idea but also an emerging practice that is generating results. As such, it is at a minimum an addition to the methodological repertoire of the history and philosophy of knowledge communities. But we argue that it is more than simply an additional set of methods. The data structures used and the kinds of results obtained contribute to a conception of knowledge as a complex co-evolving system with properties similar to those of other such systems.²²

But we are also claiming that history will benefit from an evolutionary approach. Some historians will ask why we need an explicitly evolutionary focus for studying the history of knowledge, and there is some reason for skepticism given earlier attempts. Applying evolutionary ideas to the history of knowledge has been tried before, sometimes naively (ideas about inevitable progress or simple adoption of meme-based selection models) and sometimes outright dangerously (Social Darwinism). So why is it important to try nonetheless?

First, evolutionary approaches to cultural phenomena have led to a series of novel insights into human history, from the evolution of language to social learning and cooperation, all of which are important for understanding the evolution of knowledge. Second, parts of evolutionary theory have progressed substantially and have now explicitly addressed one challenge that still plagues the history of science: the split between internalist and externalist perspectives. Evolutionary theory embraces the study of both evolving systems and the larger environment within

²² Manfred D. Laubichler and Jürgen Renn, "Extended Evolution: A Conceptual Framework for Integrating Regulatory Networks and Niche Construction," *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 2015, 324:565–577; Renn and Laubichler, "Extended Evolution and the History of Knowledge" (cit. n. 9); and Renn, *Evolution of Knowledge* (cit. n. 9).

which they exist—at multiple levels of complexity. The current mainstream of cultural history of science, with its focus on rich contextual narratives, has not been able to integrate this perspective with a similarly deep understanding of the internal details of the scientific knowledge system.

Emerging developments in evolutionary theory, such as extended evolution theory, do not simply apply evolutionary models, such as population genetics, to cultural phenomena. Instead, this emerging framework focuses on the dynamics of interactions between nested layers of internal and external regulatory structures or niches. By "niches," evolutionary biologists mean constructed environments in the broadest sense, a theme that is obviously relevant to cultural evolution. In contrast to a more traditional population genetics perspective, extended evolution theory does not view these niches as given (such as fitness landscapes) but, rather, focuses on how they are constructed by the actions of interacting agents and can then also take on their own transformative trajectory.

Another feature that addresses our main point in this Focus section is that history of science and technology can gain significantly from the strong connection of this approach to newly available possibilities afforded by the big data and computational methods described here. We should also emphasize that approaches driven by big data and computational methods are not only necessary for understanding the recent past, the present, and the future but can also be successfully applied to earlier historical periods thanks to pathbreaking developments in digital humanities and widespread digitization and conservation efforts directed toward preserving cultural heritage. And yes, we are aware that this process is only just beginning and will take time to come to full fruition. But it is already clear that the availability of digital resources greatly increases the participation of scholars from developing countries who otherwise would have no access to these materials.

Taken together, the computational and evolutionary approaches enable us to ask new questions and to broaden our analysis of previously researched cases. We see huge opportunities and also exciting challenges. We see much room for optimism and many reasons not to embrace fear. But let us also emphasize that computational analyses complement, not replace, narrative interpretations and traditional historical approaches. They provide different forms of quantitative and qualitative evidence, which in turn will lead to richer narratives of the evolution of knowledge. Such innovations will be important as both the world we study and the ways we study it change along with newly emerging configurations of knowledge in the Anthropocene.