

Digital Literature Analysis for Empirical Philosophy of Science

Oliver M. Lean, Luca Rivelli, and Charles H. Pence

Empirical philosophers of science aim to base their philosophical theories on observations of scientific practice. But since there is far too much science to observe it all, how can we form and test hypotheses about science that are sufficiently rigorous and broad in scope, while avoiding the pitfalls of bias and subjectivity in our methods? Part of the answer, we claim, lies in the computational tools of the digital humanities, which allow us to analyse large volumes of scientific literature. Here we advocate for the use of these methods by addressing a number of large-scale, justificatory concerns—specifically, about the epistemic value of journal articles as evidence for what happens elsewhere in science, and about the ability of digital humanities tools to extract this evidence. Far from ignoring the gap between scientific literature and the rest of scientific practice, effective use of digital humanities tools requires critical reflection about these relationships.

1. Introduction

According to a number of critics, traditional philosophy has given us a largely inadequate view of science—one that is in many ways distorted, idealized, and generally out of touch with its varied and complex nature. This gives rise to a general worry that philosophical conclusions based on such a picture of science are likely to be questionable themselves. Among many possible ways of remedying this, philosophers have increasingly made use of empirical methods such as surveys and ethnographic studies (Nersessian [1995]; Leonelli [2016]; for a helpful survey, see Machery [2016]). The hope is that our philosophical claims about science may be more legitimate if they are based, at least in part, on such empirical methodology.¹

¹ While we lack the space to clearly describe these connections here, this move toward empirical methodology is also certainly related to both the ‘practice turn’ in philosophy of science (Soler et al. [2014]) and the rise of experimental philosophy (Machery [2016]), as other ways of responding to the same kind of challenge.

Electronically published October 30, 2023.

The British Journal for the Philosophy of Science, volume 74, number 4, December 2023.

© The British Society for the Philosophy of Science. All rights reserved. Published by The University of Chicago Press for The British Society for the Philosophy of Science. <https://doi.org/10.1086/715049>

Some philosophers have recently applied this empirical spirit to the way they approach scientific literature. While philosophers have always used scientific writings in their work for a wide variety of purposes, it is all too easy to go about this in a way that biases one's conclusions from the start. Indeed, one purported source of error in traditional philosophy is that it has based its generalizations about science on unrepresentative examples, cherry-picked from the literature and lifted out of their broader context. Given the massive and increasing volume of scientific publications, however, it has been hard to see an alternative that could yield results that are both general and reliable. How can we analyse enough of the literature for our claims about it to have broad validity, and without biasing our conclusions through our search methods?

One way to do so is to analyse scientific literature using tools from the digital humanities. Among other methods, the digital humanities involves the use of computational techniques to analyse large collections (corpora) of documents in order to detect features that would evade unaided examination. These techniques can be deployed in a variety of ways: to detect aspects of language use or style in a corpus (Burrows [2002]; Hoover [2010]), to construct networks of relationships and interactions between authors (Börner et al. [2004]; Abbasi et al. [2011]), or to track large-scale differences between different disciplines (Dias et al. [2018]) or conceptual change over time (Chavalarias and Cointet [2013]), to name just a few. There is great promise in utilizing these tools to satisfy philosophy of science's increasing interest in empirical rigour, objectivity and scope.

However, there is an apparent tension created by the application of digital humanities methods to empirical philosophy. It is generally acknowledged that there are major disparities between what is presented in the scientific literature and what actually happens in the process of scientific investigation and discovery. More specifically, it's been argued that the version of events depicted in research articles—of what was done, how, why, in what order, and so on—is a polished, idealized, and generally inaccurate version of what actually happens in the laboratory or in the field (Medawar [1963]; Nickles [1988]; Rouse [1990]; Schickore [2008]). It is partly for this very reason that many philosophers have moved away from a reliance on readings of the literature and gone 'behind the scenes', to observe what is left out of the written record. Because of this view of publications, many philosophers may regard developments in digital literature analysis with suspicion: the disconnect between the content of scientific publications and other scientific practices—which for brevity we simply call the gap—may largely undermine the value of scientific literature as evidence for empirical philosophers. If this distance is too large, then trying to learn anything about science from the literature could be a lost cause.

One way to address this worry would be to point to examples of edifying work using these tools, arguing for them by demonstrating their power in particular cases.²

² We will consider later some broad examples of the kinds of philosophical questions that can be profitably treated using digital tools. For readers who would like to investigate some more completely fleshed out cases of digital analyses in philosophy of science, we can recommend, at least (Overton [2013]; Ramsey

Our aim here, in contrast, is to tackle directly the sources of wholesale scepticism about the value of these methods—scepticism based on the nature of the tools themselves, and on the aforementioned gap between the literature and the rest of scientific practice, or what the sceptic might call ‘actual’ science. Ultimately, the challenges for any given work of digital philosophy of science will depend on the specific question being addressed, the corpus being used, the methods employed, and so on. We intend our discussion here to provide a top-down view of these challenges and why they are surmountable. Our hope is to convince the reader that these worries do not undermine the potential of digital corpus analysis in the broader project of empirical philosophy of science. Instead, they simply highlight the epistemic challenges that must be carefully considered in any project using these tools.

2. From Scientific Literature to Philosophical Claims: Clarifying the Problem

Discussion about what philosophers can learn about science by digital corpus analysis will be more productive with a clear picture of the moving parts in the problem. Our first task, then, is to break the question down into smaller, more philosophically tractable elements. To that end, this section is in effect an extended outline of the article’s structure: it offers a view of our attempt to decompose the problem into distinct components, each of which is dealt with in order in subsequent sections.

Our proposed framework for addressing these issues is represented in figure 1, which shows a series of productive or inferential relationships, each of which can be dealt with semi-independently. The first question concerns the aforementioned relationship between the content of scientific articles and the rest of the practices of science. When it is claimed that journal articles are an unfaithful representation of what science is ‘really like’, what exactly is changed, added or omitted? With notable exceptions we will discuss, little work has been done to gain traction on these questions in ways that offer clear lessons for philosophical methodology. There are various theories about what this gap amounts to, each with different implications for what, if anything, can be learned about science from the literature it produces, whether by digital or traditional means.

While, as we will see, many authors have argued that there is some kind of a ‘gap’ in this first relationship, there are different ways of filling in the details. First, what exactly is the gap a gap between? One answer might be that the gap is between scientific articles and scientific practice: what scientists write does not accurately reflect what they do. But publication practices just are a kind of scientific practice, and writing is a kind of doing. In a sense, then, there is no ‘gap’ between journal articles and

and Pence [2016]; Murdock et al. [2017], [2018]; Pence and Ramsey [2018]; Malaterre et al. [2019]; Mizrahi [2020a], [2020b]; Wray [2020]). We thank an anonymous reviewer for encouraging us to highlight these results.




Figure 1. Structure of three related questions about the use of scientific literature in the empirical philosophy of science. The first arrow, the relationship between scientific publication and the rest of scientific practice, is possibly threatened by what we have called ‘the gap’. The second and third arrows concern how we derive generalizations about the literature from its content using digital tools, and how we then fashion those generalizations into philosophical claims. The dashed arrow, finally, indicates the place of empirical philosophy of science that does not rely on journal articles (analysis of experiments, data, ethnographies, and so on). Accounts of all three of the solid arrows should be considered, we claim, in order to fully justify philosophers’ use of digital analyses of scientific articles.

scientific practice: they are a part of scientific practice just like other resources of interest to philosophers, such as experiments, scientific data, or laboratory manuals. Even so, this in itself doesn’t settle the epistemic problem of what can be learned from corpora of journal articles. We still need to understand the relationship between publication, as one thing that scientists do, and the other kinds of practices that make up science. Explicating that relationship requires us to focus in on those publication practices: How exactly does scientific practice give rise to journal articles? What is the role of publication practices in the doing of science? For which audiences and goals (both epistemic and social) are journal articles prepared? In section 3, we consider answers to these questions and their implications for what can be learned about science from corpora of journal articles.

Suppose now that we have an answer to these first questions—that we understand the relationship between the journal literature and science as a whole, and hence what we can learn about science from what is going on in its literature. Even so, we are left with the question of how to find out what is going on in that literature in the first place. Again, given its sheer size, it is hard to see how to generate and test generalizations about the literature with sufficient rigour and scope. Of course, philosophers have long experience with methods of close reading and the construction of case studies from the scientific literature, though they are increasingly recognized to face a variety of limitations and problems (Mizrahi [2020a]). With the advent of the digitization of the journal literature, the tools of the digital humanities offer philosophers a different and potentially powerful perspective. However, there are different theoretical views of how the directly accessible features of the literature—its syntax—relates to the semantic content in which philosophers are generally interested. Importantly, different theoretical views about language and meaning and their relationship to corpora hold different implications for how to go about corpus-based analysis of science. We outline some of these views and discuss their implications for philosophical research in section 4.

Finally, section 5 summarizes the discussion by tying together these various threads. Ultimately, our goal as philosophers is to establish how to approach the scientific

literature in the service of philosophical interests. Again, even with an answer to the first two questions—of how to make reliable generalizations about features of the literature, and of how those features relate to the rest of science—there remains the question of what we as philosophers can do with these outputs. We will not commit to a general account of what constitutes a philosophical (as opposed to ‘merely’ a scientific) question, nor do we believe that much is gained by such an account: disciplinary boundaries are blurred and porous, as they should be. Instead, we simply consider features of science that, as a descriptive matter, philosophers tend to be interested in: questions about the meanings of scientific terms, experimental methods, norms of justification, and so on.

In summary, the task before us is to discuss ways to think about the three solid arrows across the top of figure 1. We will begin by carefully considering the relationship between journal publication and other scientific practices, paying particular attention to various accounts of how the gap might arise or fail to arise. We then consider how we draw generalizations from that literature, in our case roughly equivalent to validating our use of the methods of the digital humanities in the philosophy of science. Finally, we will briefly consider in the conclusion what it is to construct ‘good’ philosophy of science on the basis of these general claims about the literature.

3. Characterizing the Gap between Literature and Science

We noted in the last section that if we view scientific publications as a part of scientific practice, we can understand them in equivalent terms to other things that empirical philosophers study, namely, as data or sources of evidence whose significance to philosophers should be based on theoretical understanding of their production. With that in mind, the next step is to explore alternative views of the role of publication practices in science more broadly, each of which has different implications for the relationship between the scientific literature and other philosophically interesting aspects of science.

When considering the epistemic value of scientific publications, we should begin with the fact—almost too obvious to mention—that scientists themselves make use of publications to keep in touch with the landscape of their field and inform their own investigative work. So unless they are in some way deeply misguided in doing so, we can take for granted that the literature does carry content that is informative to scientific practitioners. Any view of the gap that makes it impossible for scientists to have this kind of relationship to their own literature should be treated with suspicion. This does not make the digital philosopher’s task trivial, however, because it doesn’t imply that the content of articles is transparent. Journal articles are typically written for a specialist audience with a considerable background of shared contextual knowledge, and scientists learn how to write and read publications through a combination of deliberate training and learned skill acquired from experience. It would be all too easy for someone lacking this training and knowledge to misinterpret them. What’s more,

we cannot simply take scientists at their word about how they write and read scientific articles and why. First, much of what one needs to know may be implicit or unconscious, and, second, we mustn't assume that the whole process works in the way it is commonly understood to work. These considerations underlie the need for the third-person perspective of science studies, which aims to make explicit and critically analyse practices that are not necessarily understood, or even known, by those engaging in them. In short, even if the prospect of analysing scientific texts for philosophical purposes isn't futile, it is also not simple.

To that end, this section considers several accounts of how and why scientific publications are produced and, consequently, of the relationship between their content and other aspects of science in which philosophers may be interested. It is only with answers to these questions that we can say what we as empirical philosophers can learn about science from its literature, and how to go about doing so. As we'll see, although different views of this relationship have importantly different implications for how philosophers should interpret their content, all of them attribute a valuable role to scientific literature in some sense.

The ways of understanding the gap that we discuss are far from exhaustive. Despite appearances, they are also not necessarily mutually exclusive: scientific writings are multi-faceted, and some of the views concern influences that only affect some of those facets. And even where they seem to address the same aspect, more than one of the influences they describe may be at work, or they may apply in varying degrees to different disciplines or types of literature. Our sketch here should be seen as a toolkit of possible perspectives for philosophers to take in their work.

Two asides are worthy of note before we continue. First, we should say explicitly that one could reject all of the arguments that push for the existence of any gap at all, and argue instead that scientific articles do, in fact, faithfully represent the content of the rest of scientific practice. In this case, the first of the three arrows in figure 1 doesn't need justification, and the reader can skip to the next section. That said, we believe that many of the arguments in favour of the existence of various disconnects between publications and other parts of science are at least compelling enough to be worth evaluating, and are perhaps even compelling enough to merit cautious response justifying our use of the scientific literature.

Second, we will assume in the following that the scientific articles at issue are produced in what we might call optimal working conditions—that is, that scientists are publishing all and only the articles that they want to be publishing, when they want to be publishing them. It is clear that even in these (impossible) conditions, there will still be a worry about the connection between those articles and the other realities of science. Our discussion neglects, however, what we might call 'ecosystem-level' effects of scientific publishing, such as the file-drawer effect (the fact that negative, uninteresting, or unwanted results are not published), popularity of particular 'hot' topics in high-impact journals, and so forth (see, for example, Ioannidis [2005]; Smaldino and McElreath [2016]). While we believe that these are clearly important

problems that, if pervasive, will make the connection between published articles and other facets of practice more tenuous, we lack the space to pursue them here.

We begin with views that deal specifically with scientific journal articles' portrayal of scientific reasoning and justification. This matter has been discussed at length by Schickore ([2008])—to our knowledge the most detailed treatment to date of any dimension of the gap that considers its implications for philosophers. Schickore's focus is primarily on the gap's implications for epistemology, that is, on where philosophers should look to understand the logic of scientific reasoning and the loci of scientific or theoretical knowledge. We recount here four broad interpretations of scientific justification in journal articles, which we will call the logical, sociological, fraud, and narrative theories.

The logical view is represented in various forms by Reichenbach ([1938]), Suppe ([1998]) and others. These authors assert, in varying ways, that journal articles contain what a discipline considers to be legitimate forms of justification for empirical claims. Reichenbach famously distinguished the context of discovery from the context of justification. Applying that distinction in the domain of scientific publication, one might think of articles as being written in the 'justificatory' mode, as vehicles by which scientists intend to justify their findings and conclusions to each other. The context of discovery—the process by which those conclusions were actually reached through theorizing and experiment—may not be represented in the finished product of research articles. Here, though, the gap offers us an important signal: it marks out the way in which scientists have prepared their findings for presentation, and is evidence of the logical standards by which scientists believe they will be judged. In that vein, Suppe argues that philosophers have so far failed in their task of explicating scientific reasoning precisely because our theories of this process do not capture what we find in the scientific literature.

The sociological view, as we call it, is represented in the work of authors such as Knorr-Cetina ([1981]) or Latour and Woolgar ([1987]), and emphasizes the social role of journal articles in scientific communities. According to Knorr-Cetina, scientific findings are not just produced but largely tested and credentialed in the lab, in ways that are messy, contextual, pragmatic, and opportunistic. The reconstruction performed in publication, in contrast, is intended to play a kind of recruiting role: the aim there is to gain converts to the research programme, fend off criticisms from hostile competing factions, and so on. On this view, journal articles contain neither the actual facts of discovery nor their actual justification, but simply attempts to gain followers. This is done in a particular social context involving the kind and degree of scepticism or hostility that the claims are likely to be met with. On this view, what we see in journal articles is to a large extent a kind of power struggle, and so the language within is to be interpreted at least partly in sociopolitical rather than logical terms.

The fraud theory represents the most pessimistic view of journal articles and the reasoning presented in them. This view was notoriously taken by Medawar ([1963]) in his 'Is the Scientific Paper a Fraud?'. Papers, Medawar says, are written as though

scientific inferences are operations of induction—deriving generalizations from passively and dispassionately acquired observational data. Yet this bears no resemblance to how scientists actually reason, which he understands as a Popperian hypothetico-deductive method of testing preconceived hypotheses. Medawar sees no useful purpose to this distortion, and proposes that scientists should instead write in ways that more truly reflect the reasoning they actually accept.³

What we will call the narrative view focuses on another aspect of the role of journal articles: the complexities surrounding the perceived significance of scientific findings. Variations on this view are represented in (Nickles [1988]; Rouse [1990]). Rouse argues that scientific practice is essentially narrative-based: like all purposeful activity, doing science requires a sense of how the past has led to the present and how one's actions might extend into the future. In other words, what a scientist or group of scientists does is always seen as a response to the situation they are currently in, including what they know and don't know, their capabilities, and so on. The narrative is neither linear nor universal in a particular field; the view into both the past and future is constantly changing with new developments, and a field is characterized by competing narratives (for example, of the 'orthodoxy' and challengers to it). In much the same way that narratives can serve a role in scientific explanation, situating events within a larger sequence and showing how those events related to a broader space of possibilities (Currie [2019]), the doing of science relies upon these shared narratives of scientific progress in which the individual happenings within various labs are combined and given sense as part of a larger enterprise.

This narrative view of science affords a very particular role for literature—namely, that it serves to feed the scientific community's sense of narrative that is necessary for science to exist at all. A journal article presents in its introduction a 'state of play' consisting of previous work necessary for understanding its findings and their significance for future research. Crucially, this means that what is considered significant isn't fixed—results previously thought important can later be seen as irrelevant, and vice versa (a fact discussed particularly by Nickles [1988]).

Each of these views holds different implications for what philosophers can infer about the rest of science from some observed feature of scientific literature. According to the logical view, for example, the argumentation we find in journals in fact captures what the scientists themselves consider valid reasons—perhaps even the strongest reasons—for accepting their conclusions, even if it fails to capture how the authors themselves reached them. If one's philosophical interest is in a scientific community's norms of justification, then of course an analysis of journal articles would be an acceptable means of finding this out. While there remain questions on this view about which

³ We will not consider the fraud view extensively in the remainder of the article; while we lack the space to engage with it more fully here, it seems that the investment of time (for example, in writing or reviewing) and engagement (for example, in journal clubs or discussion of articles on social media) provides a strong *prima facie* argument that scientists do derive genuine knowledge about the scientific process from scientific articles, of at least some sort.

publications to use, and which parts (Schickore [2008]), it remains the case that the relevant evidence is to be found somewhere among them.

However, accepting the logical view does cast doubt on the value of journal articles for addressing other philosophical questions, including many that interest philosophers concerned with scientific practice. Since the practice turn, many philosophers now take direct critical interest in the strategies scientists use to investigate and manipulate their objects of study (Hacking [1982]; Wimsatt [2007]; Waters [2008]; Baxter [2019]). Many philosophers of science now view science as skilful problem-solving—the overcoming of various practical, material, and social challenges in order to produce knowledge. For the particular purpose of understanding scientific practices as solutions to these challenges, the literature is indeed likely to be of limited use, since much of this problem-solving occurs elsewhere and isn't reported in the final presentation.

In contrast to the logical view, the sociological view suggests that we cannot, or not easily, discover what scientists accept as good scientific reasoning from analysing articles. The encroachment of sociopolitical factors cautions us not to take any claim in an article at face value, since it may be at least partly a rhetorical move rather than a dispassionate presentation of reasons. Such considerations have implications for other areas of interest to philosophers, such as the role and importance of certain kinds of language. For example, Kay ([2000]) argues that talk of 'information' in molecular biology took hold because of its rhetorical power rather than its importance to actual scientific research, most of which led to dead ends when guided by informational principles.

Claims of this sort are in effect claims about the relationship between journals and other parts of research programmes, asserting that the former are something like the 'public face' of the latter. If true, they suggest that we cannot safely infer the importance of a concept to scientists' daily work from its prevalence in the literature. Of course, we can acknowledge the role of irrational or non-rational forces of social psychology at work in literature exchanges, while also supposing that scientists are at least sometimes or partly swayed by reason over rhetoric. In any case, this makes the task of discovering these motivations in scientific articles more complex.

Despite the complications this holds for the scientific epistemologist, the sociological view of publication practices also reveals a positive role for journal articles. For if social forces manifest themselves in scientific writing, we should be able to detect signals of these forces in the literature. For example, it may be that the amount and strength of rhetorical language an article uses reflects the degree of hostility that the authors expect towards their findings. Learning to detect these signals would, of course, also require cross-referencing with actual sociological observations. In any case, even the sociological view suggests that the language used in articles offers clues about something of value to philosophers interested in scientific practice; namely, the sociological context of the science and its development over time.

What can we say about the use of scientific literature for philosophical purposes from the narrative perspective? Like the sociological view, conceiving of a scientific

article as meaningful only in relation to a narrative certainly complicates the problem, particularly when it comes to judging whether an individual work or body of work is significant. This is something philosophers may be interested to know: for many purposes, the philosophical significance of scientific work will depend largely on its scientific significance. If this significance itself is only relative to a narrative, we must understand the narratives and, crucially, their evolution over time in order to make proper use of that work in philosophical research. This diachronic view is especially needed if, with Nickles ([1988], pp. 37–38), we agree that changes in a narrative can lead to different ways of using and understanding concepts over time, even if they continue to be named by the same words. In any case, given the important role that publications play in constituting and reconstituting the narratives at play in scientific practice, they remain vital sources for interpreting those narratives and their change over time.

Finally, we should mention the well-developed literature critically discussing the role of citations in publication practices. The idealized view of citations is that they are supposed to indicate what previous works influenced the intellectual or practical development of the present work. But this ideal is not necessarily reflected in reality. MacRoberts and MacRoberts ([1989]), for example, highlight several shortcomings of real-world citation practice: influences can be omitted, citations are biased in various ways, especially in favour of the author's own work, citation norms vary between disciplines and across time, and so on.

These criticisms, while well known, have far-reaching implications, including for philosophers aiming to digitally analyse scientific literature. Suppose we are interested in a particular scientific finding or idea and its growth, development, and influence in or across scientific fields. The above suggests that this cannot be straightforwardly inferred from the prominence of the relevant publications in citation networks: that article B cites article A does not necessarily imply that article A provided significant practical or intellectual context for article B's development. Yet, again, accepting these criticisms needn't lead to a wholesale rejection of any attempt to quantify influence by citations. Instead, it calls for philosophers to be cautious, and elaborates on what they should be cautious about.⁴

This section has outlined several ways of conceiving of the role of scientific literature as an element of scientific practices as a whole, and the ways these various perspectives bear on what philosophers can learn from the literature and how. The key lesson is that these are markedly different views of what journal articles are written for, how they are read, what contributes to their uptake, and so on. It bears repeating that these alternatives are not mutually exclusive—each of them may simply capture distinct but equally important aspects of science and the role of publication practices within it. In any case, all these views reveal complexities to publication practices that

⁴ We will discuss some of the ways in which citation networks could be used despite the presence of these critiques in section 4.2.

we as philosophers mustn't ignore. To oversimplify, a scientific article should never be taken at face value. We should always understand a research article in pragmatic terms, as an object whose meaning depends on the role it plays within a larger social activity, which may vary dramatically from one case to another.

None of this means, of course, that articles themselves are useless to philosophers of science. Instead, our conclusion is simply that their interpretation requires an understanding of the larger social activity in which they occur. Crucially, there are alternative ways of conceptualizing this social activity, and so one's interpretation of the literature in a particular case should always be tied to a conceptualization of the nature of publication in general. This is no more or less true of digital methods of literature analysis in particular, as we now discuss.

4. Getting from Corpora to Generalizations about Science

The previous section considered the ways in which we might understand the first arrow in figure 1, that is, the relationship between corpora and other scientific practices. We now turn our focus to the second arrow of that diagram. Assuming that we have a particular body of scientific text, and an understanding of what role those publications played for the community that produced them, how can we extract general claims about what goes on in science? These generalizations, in turn (as we will discuss in the conclusion), often serve as either source material for, or crucial validations of, our philosophical claims about science itself.

On the one hand, computer-aided, corpus-driven tools are potentially an incredibly rich source of claims about science—able to avoid many classic problems of bias and cherry-picking that may render the use of particular case studies problematic (though these same worries remain for corpus analysis, albeit in different guises; see Allen and Murdock [2022]). On the other hand, such methods might result in a spate of unjustified and facile conclusions, drawn unreflectively from the output of misapplied computer-based tools (a worry that has often been raised concerning the digital humanities; see, for example, Pechenick et al. [2015]). We thus need to provide a justification of how these new methods can produce a high-level understanding of the process and practice of science. For our purposes, we will consider two different approaches that could offer us insight into the content of scientific publications, as representative examples: the use of textual analysis and the construction of citation networks.

4.1. Textual analysis

Some digital humanities methods offer us the promise of evaluating, at a large scale and across vast corpora of articles, the very meaning or conceptual structure of the articles at issue. If we could extract information about the use of a concept, for instance, from a large corpus, we could follow its usage across journals and across time, formulating and testing a variety of hypotheses about conceptual dynamics and theory

change. Even purely normative philosophical claims about the content or function of a concept need to be evaluated in the context of knowing how such concepts are actually used by scientists (as also defended by authors like Machery [2016]; as an example, see Overton [2013]). Broadly, such methods are described as different sorts of textual analysis. Let's begin by considering a simple example, co-occurrence analysis. By measuring which words are found near others throughout the corpus (for example, in the same article, in the same paragraph, or in the same sentence), we can create a kind of weighted, undirected network that connects the terms within the corpus.

How does this network allow us to understand the meaning of texts? If we look at the structure of the network—in particular, the way in which it clusters into 'communities', or subsets of terms within which terms are more densely connected to one another than to terms outside the community (Girvan and Newman [2002])—we find surprising possibilities for the automated extraction of novel knowledge from already existing data. To mention just one impressive case from the sciences, Wilkinson and Huberman ([2004]), focusing on a network of co-occurrence of the names of genes in the abstracts of a corpus of biomedical literature, were able to show that detection of communities in that network could be used to automatically discover previously unknown functional modules of genes related to colon cancer. While these functional modules were 'latently' documented (so to speak) in the extant literature, they had escaped the attention of human researchers, due to the enormous number of publications that would have had to be considered in order to find them.

For a more complex methodology, consider the use of topic modelling (well described for a philosophical audience in Malaterre et al. [2019]). Topic models break down documents into clusters of 'topics', here a technical term, defined as a set of probability distributions across all words in the corpus. Such topics, then, can be interpreted as topics in the vernacular sense, subjects that each document is 'about'. Charting the changes in topic distribution in the corpus over time can give us an idea of evolving conceptual emphasis or focus, while examining the probability distributions that make up the topics themselves might help us understand the ways in which authors talk about different subjects in different corpora. A similar method underlies recent work on the reconstruction of the movement of concepts and the formation and dissolution of scientific disciplines (Chavalarias and Cointet [2013]).

4.2. Citation networks

Another of the oldest and most successful constructs applied in the digital humanities, citation networks are networks of articles related by the citation occurring between an article and the ones that directly cite it. These networks can also be transformed into networks of citations between other kinds of entities, such as authors, journals or research institutions. Citation networks can also be constructed starting from indirect relations, such as co-citation (two entities are related when they are cited together by another entity; Small [1973]) or bibliographic coupling (two entities are related when

they both cite a third entity; Kessler [1963]). Intuitively, while co-citation networks detect commonality of topic between related entities, as they connect entities that are found in the same bibliographies and thus serve as common sources for later works, bibliographic coupling could reveal latent or indirect commonality of interests between entities that do not cite one another but rely upon the same kinds of sources, a particularly interesting relation between authors, who might be unaware of these commonalities.

Several kinds of analysis can be performed on basic citation networks, the most common one being community detection, where the goal is to detect areas of mutual citation (groups of authors or other higher-level entities that all regularly cite one another) reflecting, again, shared interests. Once these communities are found, a coarse-grained version of the original network, in which each node corresponds not to a single article (or author, or another entity) but to a single community, can be produced, easing understanding of the network structure at a large scale by letting the observer see the field as a whole and avoid the information overload of a direct observation of the complex original network.

When applied to citation networks built from relationships other than direct citation, community detection could help reveal hidden patterns. In co-citation networks, communities should in general reveal topics grouping together entities that often ‘travel together’ in later literature, while, in bibliographic coupling networks, the same technique could reveal communities of researchers with common interests, possibly in turn indicating ongoing latent processes of unification or splitting of research fields.

The topology of a citation network can shed light both on the general process by which scientific literature grows, and on the ways in which scientific information flows through the network during its dissemination. Citation networks tend in general to grow by preferential attachment: the more an article is cited, the more it gets cited subsequently. This mechanism gives rise to what is called a network with scale-free topology (Barabási and Albert [1999])—where a limited percentage of nodes form poles of attraction (the hubs) cited by many other nodes, with most other nodes considerably less cited. From an information-flow point of view, hyper-cited hubs are the primary sources of information for subsequent articles.

Specific dynamics of network evolution can be detected by analysing sequences of snapshots of the citation network over time. Diachronic analysis, especially of coarse-grained networks, could reveal large-scale phenomena in the development of corpora of scientific literature, producing fruitful results both from a historical perspective—seeing when such processes have already occurred—and from a philosophical one. As an example of the latter, in a network of communities of authors who cite the same common entities, the real nature and scope of a debate could be clarified by studying the ‘latent’ extension, both in space and in time, of the communities that fuel it.

Scientific articles can, of course, be further classified into kinds (research articles, review articles, letters, and so on), each with a typical range of uses. Diachronic reconstruction of the movement of ideas through citation networks—for example, from

the initial announcement in the form of a letter, to the appearance of the polished results in the research article, to the subsequent inclusion in a review article and then in further research articles—could help model the ever-changing narrative landscape of science stressed by authors like Nickles ([1988]).

When modelling the evolution of citation networks at shorter temporal scales, machine learning techniques (as envisioned in Shibata et al. [2012]) could be applied to predict which articles will most likely be cited at a later time. This is a way of modelling the dynamics of an author in the process of choosing which articles to cite, with the potential to reveal in general how the biases induced by the structure of the extant corpus of literature shape its subsequent growth through their influence on the practice of scientific writing. Thus biases, far from being purely negative, can be viewed as useful data for modelling the production of articles: biases in the literature influence the (mis) perception of the information conveyed to the scientists by the literature itself, providing philosophers with information about how to start tackling the ‘gap’ problem.

As mentioned above, citation networks have received their fair share of criticism (MacRoberts and MacRoberts [1989]; but see Wouters [1998]). In part, these critiques arise as a result of the fact that citations are objects with meanings, just like any other element of a scientific text: they occur in a certain context, and express the possibly biased intentions of an author when referencing other articles. There is thus a distinction between syntax and semantics at work in the analysis of citations, just as there is in the textual analysis of scientific concepts. We must, of course, be precise about just what it is that we hope to analyse. If our object is the citation network as a whole, then we are not directly interpreting the intent of an author in citing a particular article—this network is a product of those aggregated intentional behaviours, as performed by all of the writers. It can then be defended on statistical grounds (for example, van Raan [2004]) that when working with large data sets of reasonably highly cited articles, sets of citations get their random semantic biases averaged out, and in this sense they can be treated in a purely formal way, without regard to their semantic content. This runs parallel to the consideration that while semantics is likely significant for the theory of the production of single articles, it can be reasonably ignored when algorithmically analysing certain features of an entire corpus.

This line of reasoning alone seems to provide a sufficient rationale for advocating that analysis of patterns in large citation networks can still offer us a variety of insights into the structure and evolution of the scientific community. It is even likely that citation networks themselves will be indispensable for studying the disconnect between the varied and possibly biased intentions of authors driving their use of references and the ‘objective’ influence of articles on one another, as the citation network should indirectly retain, in its structural and statistical features, indirect information about the intentions behind the formation of references.⁵

⁵ The most trivial example is the typical scale-free topology of citation networks, which reveals a pervasive bias in producing references.

4.3. Analysing meaning

In all these methods, whether analysing text or citation networks, an automated algorithm, with access only to the syntax of the corpus and no access to the meaning of the terms within, reveals general or high-level structure and regularity in that syntax, which we interpret as related to its meaning. This offers us yet another justificatory challenge: how should we understand the jump that is made here from data, like frequencies of terms within documents or citation relations, to the meaning of texts? Just as inferences from what's said in the literature to other aspects of scientific practice are tied to theories about the gap between them, inferring meaning from text depends on theories about the relationship between syntactical structure and semantics.

This is, of course, by no means the first time that such a question has been broached in the literature. 'How can we get from syntax to semantics?' is a refrain perhaps most commonly found within the literature on corpus linguistics, where the entire *raison d'être* of the field is, in some sense, tied up in an answer to this question. Unfortunately, unlike traditional philosophy of language, philosophical approaches to linguistics are still somewhat in their infancy (though for a philosophy-of-language approach highly informed by interplay with linguistics, see, for instance, Santana [2016]). Our goal in this subsection will simply be to point out two approaches to the syntax-semantics relation in the corpus linguistics community and, similar to our discussion of the gap in section 3, demonstrate that on either reading it is possible to draw interesting generalizations from the scientific literature.

One approach is grounded in the social use of language (we take here as representative Teubert [2005]). On this view, the meanings of fragments of text are entirely defined by past segments of text that themselves are also to be found within the discourse. 'Meaning', Teubert writes, 'does not concern the world outside the discourse' (p. 3). But because text is diachronically connected—because 'what is said today is a reaction to what has been said before, an argument in a simultaneous debate and an anticipation of what we expect to be said tomorrow' (p. 4)—we can still extract meanings from texts by looking at the ways in which players in that discourse paraphrase, self-describe, define, even 'negotiate' what they consider 'meanings' as they arise within a particular discourse tradition.

Such a perspective on semantics has both obvious resonances and tensions with the traditional ways in which philosophers of science have treated the meanings of scientific terms. On the one hand, there is a welcome emphasis on use present here: we shouldn't trust our own reconstructions of a concept, or our potentially Whiggish or naive suppositions about how such a concept connects to the outside world. On the other hand, this account requires a very subtle understanding of the relationship between these concepts and the remainder of scientific practice. If meanings are objects within the discourse itself, then we need to carefully build our account of the ways that our other practical and social uses of language can affect those meanings. It is here that, on a sociolinguistics perspective, we can find room for scientific practice

and the social structure of science to re-enter. Such a view requires a quite subtle account of the nature and function of scientific concepts, and their accessibility to study by linguistic means.

If this approach strikes the reader as having too-far divorced meaning within the corpus from questions of reference and mental content, consider instead the perspective on corpus analysis derived from cognitive approaches to linguistics. Here, by contrast both with the social approach described above and a Chomskyan view on which language processing is handled by a dedicated mental ‘module’, cognitive linguists view syntax, semantics, pragmatics, and the other components of the structure of language as fundamentally mental, and connected to our other processes of cognition (Glynn [2014], p. 8). What, then, is the role of corpora on such a perspective, where meanings, at least in part, are grounded within the heads of the users of a language rather than lying solely within the discourse itself? Corpora of real-world text are essential in order to falsify or verify hypotheses that we derive about semantics. We can operationalize claims about the meaning and use of concepts by investigating when they become common as linguistic units in a particular corpus, that is, when they function in regular and acceptable uses of language (their ‘entrenchment’ as ‘grammatically acceptable’ in a particular corpus, in Glynn’s phrasing).

This sort of picture, again, comes with both advantages and disadvantages. On the one hand, the consistent focus on operationalization of linguistic hypotheses as claims about corpora is welcome in a digital humanities context, within which we can imagine directly implementing ways to test claims about the meanings of scientific concepts within dedicated systems of text analysis (such as Pence [2016]). On the other hand, such freedom comes with a significant price, as developing such operationalizable hypotheses about scientific concepts will by no means be a simple endeavour. We hope that digital humanities methods can be widely accessible and easy to use, and thus developing ‘best practices’ surrounding these kinds of syntax-semantics inferences in a philosophy of science context should, if we adopt this linguistic perspective, be a significant and important task for digital philosophy of science in the future.

In sum, while there are again multiple approaches to making the jump from the syntactical aspects of journal articles to their meaning (just as there were multiple ways of theorizing the connection between corpora and other aspects of scientific practice), on each of these approaches are we able to make semantic generalizations arising from careful study of the journal literature. The kinds of meta-philosophical work needed to ensure that those claims are justified will differ depending on the other kinds of commitments we have in play, but, in our view, such meta-level work seems possible.

5. Putting It All Together

Let’s step back and take stock. It has been our aim here to offer a general justification for the use of inferences from large corpora of scientific articles to claims in

the philosophy of science. This is, to reiterate, a different pursuit from that of offering examples of these inferences in the literature, though we have gestured at some such cases in the last section. We began by arguing in section 2 that the general justificatory problem we want to tackle here may be decomposed into a number of separate, related questions. One piece of the puzzle, as discussed in section 3, is to consider how the literature is related to other sorts of scientific practice, which requires considering what publication practices are like and what they are for. The next step discussed in section 4—that of producing generalizations about what is in the literature—involves a justification of the ability of digital humanities tools to extract meanings from scientific corpora. It is important to underline that these are two separate but related questions. Importantly, because there are multiple coherent understandings available both of the relationship between publication and other scientific practices, and of the justification of generalizations drawn from corpora, we cannot take either of these questions for granted. We owe ourselves answers to each that can work in concert to produce a coherent picture of what kinds of philosophical inferences the literature can be taken to support. We conclude here by examining the very last solid arrow in figure 1, in pursuit of a few glimpses of this broader picture: how can we get from empirical generalizations about the contents of the literature to philosophical arguments or conclusions?

What exactly counts as a philosophical question is, of course, a question much too broad and contentious for us to settle here. For our purposes, we take it to be at least partly a matter of generality. A philosophical theory of science—a theory of ‘how science works’—is more than a laundry list of empirical particulars: while we needn’t insist on anything like universality, we aim to abstract and generalize to at least some degree over particular empirical facts about scientific practice. In any case, we take it as uncontroversial that philosophical claims about science should be based on well-supported empirical generalizations about what science is really like. For those who might not be convinced, we want to close by hinting at a few ways in which such generalizations could serve as an important tool for philosophical work. In doing so, however, we do not intend to minimize the importance of ‘cross-validating’ the kind of work we describe here with other methods. For that, it seems that one would need, for instance, to observe the practical realities of scientific research that the literature doesn’t show. Even then, however, the comparison between the face of the research presented by the literature and what lies behind it is philosophically interesting in itself. Analyses of the literature, including those made possible with digital humanities tools, would lend themselves to a principled comparison of that kind.

How could generalizations about practice be useful for philosophy of science? The question has already received some discussion in the literature (Machery [2016]; Mizrahi [2020b]); we can point to a few important possibilities. First, we might think of them as playing the role of hypothesis generators, or exploratory tools. Consider, for example, the crucial role played by historical reflection on the Fresnel-Maxwell debate in the development of ontic structural realism (Worrall [1989]; French [2011]), which was used as a way to demonstrate the possibility—via its actual use by historical

scientists—of a novel way of thinking about the status of unobservable entities as grounded in facts about their relational structure. Such a hypothesis, while of course relying on a host of other concerns raised in the literature and the varying configurations of the debate over scientific realism, was in no small part spurred by the realization that another mode of thinking about science could be extracted by generalizing from this particularly fruitful instance of historical theorizing. Generalizations from the literature, we think, have often played such a catalytic role.

Second, we often rely upon such generalizations as ways in which we might test our philosophical hypotheses. As was seen in the tight connection between Medawar's fraud view of the gap and Popperian falsificationism as, he thought, the 'true' theory of scientific knowledge generation, the kinds of claims made by philosophers and historians of science about issues like theory change or scientific progress often come with a variety of very important implications, whether descriptive ones for our interpretation of the process of science, or normative recommendations for practising scientists. Our ability to test these philosophical conjectures thus relies largely on our ability to extract well supported evidence from the varied products of the scientific process itself. In short, on either of these views—or any of the numerous others that one might propose—the derivation of empirical generalizations from the scientific literature should be a key element in our studies of science.

On the approach we presented here, how could the questions we have raised work together in defence of philosophical claims? We will close the article by presenting several such 'packages' of responses to the various questions we have raised, and elaborating some of the kinds of digital analyses that would be licensed on these understandings of the nature, function, and meaning of the scientific literature. First, consider what would result if one takes the social view of the gap, and adds to it a social-linguistics understanding of the connection between syntax and semantics. Claims about conceptual analysis are indeed accessible on these premises via the literature, but those concepts as they appear in the literature are a sort of socially negotiated entity, constantly in flux, subject to definition and redefinition over the course of a particular scientific episode. Digital analyses grounded in this way would need to focus on diachronic methods for reconstructing scientific debates. For example, the 'phylomemy' approach of Chavalarias and Cointet ([2013]; though this is not the only methodology available, see Tahmasebi et al. [2021]), which extracts clusters of concepts from text and examines the ways in which these concepts change over time, would be particularly suitable. But because these concepts are serving a social role in the literature, we would need to support that with a clear understanding of the scientific communities at play, for which citation-network methods might offer a very useful indicator (Braam et al. [1991]). Textual analysis tools could help us identify places where the definitions of such concepts were renegotiated or altered—and we would be able to see by comparing that analysis with relevant social changes whether and how those conceptual changes were related to 'moves' in the social game described by authors like Latour. We would therefore be able to approach questions

of conceptual change and its social motivation: How have different groups utilized scientific concepts to social ends? What are some important examples of these renegotiations of meaning? How do social networks influence conceptual change?

Next, turn to the combination of a narrative view of science with a cognitive-linguistics approach to the meaning of texts. We are searching here for signals of the entrenchment of new terms within the corpus, as ways to test our hypotheses about the meaning and function of the concepts they represent. Those concepts, in turn, function as structuring elements in the narrative that scientists build for themselves about the content of scientific articles. Here, simpler methods may well prove fruitful. Topic modelling (Blei [2012]), perhaps extended to dynamic topic models capable of change over time (Blei and Lafferty [2006]), are a well understood technique designed to classify articles, without direct human supervision, according to the topics or concepts that they contain. Such models, when visualized over time (for example, applying the ‘alluvial diagrams’ of Rosvall and Bergstrom [2010]) can provide a powerful way in which to understand the ways in which concepts can enter and exit these narratives at a broad scale, giving us insight that heretofore could only be approximated by tracing the impact of ‘landmark’ journal articles.⁶ These analyses could aid in the consideration of more traditional questions of conceptual analysis. To what use have certain scientific concepts been put in fashioning the narrative of scientific progress? Are there, for instance, delays between the introduction of a concept and its spread in the literature? Do the dynamics of such narrative change differ across different fields of the sciences?

Third (and finally, for our purposes here), imagine taking the logical view of the scientific literature and combining it with the cognitive-linguistics approach. If the logical view holds, we now not only have access to conceptual features in the text—exactly as we did on the previous ‘package’ of views—but also have access to structural features of scientific argumentation, where these are now taken not only to be relevant merely for ‘getting the article published’, but further, to represent the genuine argumentative structures a field finds most compelling. Computational analysis of these argument structures would therefore become a viable enterprise. While this is still a very young field, argument mining, a new branch of natural language processing, could provide an exceptionally powerful tool here (Lawrence and Reed [2020]). Extracting and comparing the ‘community approved’ structures of inference from the literature as a whole would give us an entirely novel way to understand the patterns of scientific reasoning. What argument patterns are commonly used by a given community? To what extent are these patterns discipline-specific? How do they relate to our extant philosophical literatures on scientific reasoning or explanation?

These are only three possible combinations of the (non-exhaustive) set of positions we have described here, of course, and other packages of responses would produce

⁶ These approaches to disentangling narratives could equally well be combined with the analyses of term redefinition and negotiation mentioned in the previous example, if we wished instead to combine the narrative view with a sociolinguistics perspective.

dramatically different uses of these tools. We have also only mentioned a few uses from the myriad that would be available given each such ‘packages’, and a few such questions those methods could be used to explore. We want only to insist that in the absence of such a set of responses to the questions we have raised here, and a careful consideration of what kinds of tools that set of responses would justify, it will remain difficult to determine whether the conclusions drawn in any given study are properly justified (for similar worries, see also Allen and Murdock [2022]).

6. Conclusion

There is a long road ahead for the integration of corpus analysis into the methodological toolkit of empirical philosophy of science. However, we hope to have shown that these problems are not qualitatively different from other empirical methods—at least not in a way that is fatal for the possibility of a role for them in philosophy. Even on the most pessimistic views about the relationship between publication and other parts of scientific practice, we have seen that a number of avenues for digital analysis will remain open, which can shed light on a variety of philosophically interesting parts of the scientific enterprise. In turn, this insight will be further contextualized by an understanding of the local context of the role of publication acquired by other empirical means such as ethnographic studies. Rather than replacing those other methods, they can and should coexist in a relationship of mutual scrutiny and enrichment. In short, the digital humanities and its challenges are simply business as usual for empirical study of science (with significant advantages, we think, thanks to the unprecedented scale of the data involved). Again, the way we go about the empirical part of our philosophical research, in this case as in all cases, must be adapted, consciously and repeatedly, to the types of questions we ask and to the particular environments of the scientific fields about which we ask those questions.

In this sense, we hope to be joining a larger chorus of authors within the science studies community who, inspired by the recent successes of empirical approaches to the study of science, have turned to dedicated studies of the justification of our own philosophical practices. We think such a turn is long overdue, and if our work here can encourage further reflection on these important questions, we will consider it successful.

Acknowledgements

This work was supported by the Fonds de la Recherche Scientifique (grant F.4526.19). Our thanks to two anonymous referees for insightful comments. For comments on previous versions of this article, we thank Mark Alfano and Carlos Santana. This article benefited from audiences at the XPhiSci 2019 conference at Aarhus University, especially Sam Schindler, Edouard Machery, Moti Mizrahi, Deena Skolnick Weisberg, Carlos Santana, Michiru Nagatsu, and Andrew Shtulman; and at

the Center for Philosophy of Science and Societies Seminar at the Université catholique de Louvain, especially Michel Ghins, Gwenael Laurent, Alexandre Guay, and Peter Verdée.

Oliver M. Lean
Institut supérieur de philosophie
Université catholique de Louvain
Louvain-la-Neuve, Belgium
oliver.lean@uclouvain.be

Luca Rivelli
Institut supérieur de philosophie
Université catholique de Louvain
Louvain-la-Neuve, Belgium
luca.rivelli@uclouvain.be

Charles H. Pence
Institut supérieur de philosophie
Université catholique de Louvain
Louvain-la-Neuve, Belgium
charles.pence@uclouvain.be

References

- Abbasi, A., Altmann, J. and Hossain L. [2011]: ‘Identifying the Effects of Co-authorship Networks on the Performance of Scholars: A Correlation and Regression Analysis of Performance Measures and Social Network Analysis Measures’, *Journal of Informetrics*, **5**, pp. 594–607.
- Allen, C. and Murdock, J. [2022]: ‘LDA Topic Modeling: Contexts for the History and Philosophy of Science’, in A. De Block and G. Ramsey (eds), *The Evolution of Science: Computational Frontiers in History and Philosophy of Science*, Pittsburgh, PA: University of Pittsburgh Press, pp. 103–19.
- Barabási, A.-L. and Albert, R. [1999]: ‘Emergence of Scaling in Random Networks’, *Science*, **286**, pp. 509–12.
- Baxter, J. [2019]: ‘How Biological Technology Should Inform the Causal Selection Debate’, *Philosophy, Theory, and Practice in Biology*, **11**, available at <doi.org/10.3998/ptpbio.16039257.0011.002>.
- Blei, D. M. [2012]: ‘Probabilistic Topic Models’, *Communications of the ACM*, **55**, pp. 77–84.
- Blei, D. M. and Lafferty, J. D. [2006]: ‘Dynamic Topic Models’, in W. W. Cohen and A. W. Moore (eds), *Proceedings of the 23rd International Conference on Machine Learning*, New York: ACM, pp. 113–20.
- Börner, K., Maru, J. T. and Goldstone, R. L. [2004]: ‘The Simultaneous Evolution of Author and Paper Networks’, *Proceedings of the National Academy of Sciences USA*, **101**, pp. 5266–73.
- Braam, R. R., Moed, H. F. and van Raan, A. F. J. [1991]: ‘Mapping of Science by Combined Co-citation and Word Analysis, I: Structural Aspects’, *Journal of the American Society for Information Science*, **42**, pp. 233–51.

- Burrows, J. F. [2002]: ‘“Delta”: A Measure of Stylistic Difference and a Guide to Likely Authorship’, *Literary and Linguistic Computing*, **17**, pp. 267–87.
- Chavalarias, D. and Cointet, J.-P. [2013]: ‘Phylomemetic Patterns in Science Evolution: The Rise and Fall of Scientific Fields’, *PLOS ONE*, **8**, available at <doi.org/10.1371/journal.pone.0054847>.
- Currie, A. [2019]: *Scientific Knowledge and the Deep Past: History Matters*, Cambridge: Cambridge University Press.
- Dias, L., Gerlach, M., Scharloth, J. and Altmann, E. G. [2018]: ‘Using Text Analysis to Quantify the Similarity and Evolution of Scientific Disciplines’, *Royal Society Open Science*, **5**, available at <https://doi.org/10.1098/rsos.171545>.
- French, S. [2011]: ‘Shifting to Structures in Physics and Biology: A Prophylactic for Promiscuous Realism’, *Studies in History and Philosophy of Biological and Biomedical Sciences*, **42**, pp. 164–73.
- Girvan, M. and Newman, M. E. J. [2002]: ‘Community Structure in Social and Biological Networks’, *Proceedings of the National Academy of Sciences USA*, **99**, pp. 7821–26.
- Glynn, D. [2014]: ‘Polysemy and Synonymy: Cognitive Theory and Corpus Method’, in D. Glynn and J. A. Robinson (eds), *Corpus Methods for Semantics: Quantitative Studies in Polysemy and Synonymy*, Amsterdam: Benjamins, pp. 7–38.
- Hacking, I. [1982]: ‘Experimentation and Scientific Realism’, *Philosophical Topics*, **13**, pp. 71–87.
- Hoover, D. L. [2010]: ‘Authorial Style’, in D. McIntyre and B. Busse (eds), *Language and Style: Essays in Honour of Mick Short*, Basingstoke: Palgrave, pp. 250–71.
- Ioannidis, J. P. A. [2005]: ‘Why Most Published Research Findings Are False’, *PLOS Medicine*, **2**, available at <doi.org/10.1371/journal.pmed.0020124>.
- Kay, L. E. [2000]: *Who Wrote the Book of Life? A History of the Genetic Code*, Stanford, CA: Stanford University Press.
- Kessler, M. M. [1963]: ‘Bibliographic Coupling between Scientific Papers’, *American Documentation*, **14**, pp. 10–25.
- Knorri-Cetina, K. D. [1981]: *The Manufacture of Knowledge: An Essay on the Constructivist and Contextual Nature of Science*, Oxford: Pergamon.
- Latour, B. and Woolgar, S. [1987]: *Laboratory Life: The Construction of Scientific Facts*, Princeton, NJ: Princeton University Press.
- Lawrence, J. and Reed, C. [2020]: ‘Argument Mining: A Survey’, *Computational Linguistics*, **45**, pp. 765–818.
- Leonelli, S. [2016]: *Data-Centric Biology: A Philosophical Study*, Chicago, IL: University of Chicago Press.
- Machery, E. [2016]: ‘Experimental Philosophy of Science’, in J. Systema and W. Buckwalter (eds), *A Companion to Experimental Philosophy*, New York: Wiley, pp. 473–90.
- MacRoberts, M. H. and MacRoberts, B. R. [1989]: ‘Problems of Citation Analysis: A Critical Review’, *Journal of the American Society for Information Science*, **40**, pp. 342–49.
- Malaterre, C., Chartier, J.-F. and Pulizzotto, D. [2019]: ‘What Is This Thing Called Philosophy of Science? A Computational Topic-Modeling Perspective, 1934–2015’, *HOPOS*, **9**, pp. 215–49.
- Medawar, P. B. [1963]: ‘Is the Scientific Paper a Fraud?’, *Listener*, **70**, pp. 377–78.
- Mizrahi, M. [2020a]: ‘The Case Study Method in Philosophy of Science: An Empirical Study’, *Perspectives on Science*, **28**, pp. 63–88.

- Mizrahi, M. [2020b]: 'Hypothesis Testing in Scientific Practice: An Empirical Study', *International Studies in the Philosophy of Science*, **33**, available at <<https://doi.org/10.1080/02698595.2020.1788348>>.
- Murdock, J., Allen, C. and DeDeo, S. [2017]: 'Exploration and Exploitation of Victorian Science in Darwin's Reading Notebooks', *Cognition*, **159**, pp. 117–26.
- Murdock, J., Allen, C. and DeDeo, S. [2018]: 'Quantitative and Qualitative Approaches to the Development of Darwin's *Origin of Species*', *Current Research in Digital History*, **1**, available at <<https://doi.org/10.31835/crdh.2018.14>>.
- Nersessian, N. J. [1995]: 'Opening the Black Box: Cognitive Science and History of Science', *Osiris*, **10**, pp. 194–211.
- Nickles, T. [1988]: 'Reconstructing Science: Discovery and Experiment', in D. Batens and J. P. van Bendegem (eds), *Theory and Experiment: Recent Insights and New Perspectives on Their Relation*, Dordrecht: Reidel, pp. 33–53.
- Overton, J. A. [2013]: '"Explain" in Scientific Discourse', *Synthese*, **190**, pp. 1383–405.
- Pechenick, E. A., Danforth, C. M. and Dodds, P. S. [2015]: 'Characterizing the Google Books Corpus: Strong Limits to Inferences of Socio-cultural and Linguistic Evolution', *PLOS ONE*, **10**, available at <doi.org/10.1371/journal.pone.0137041>.
- Pence, C. H. [2016]: 'RLetters: A Web-Based Application for Text Analysis of Journal Articles', *PLOS ONE*, **11**, available at <doi.org/10.1371/journal.pone.0146004>.
- Pence, C. H. and Ramsey, G. [2018]: 'How to Do Digital Philosophy of Science', *Philosophy of Science*, **85**, pp. 930–41.
- Ramsey, G. and Pence, C. H. [2016]: 'EvoText: A New Tool for Analyzing the Biological Sciences', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **57**, pp. 83–87.
- Reichenbach, H. [1938]: *Experience and Prediction: An Analysis of the Foundations and Structure of Knowledge*, Chicago, IL: University of Chicago Press.
- Rosvall, M. and Bergstrom, C. T. [2010]: 'Mapping Change in Large Networks', *PLOS ONE*, **5**, available at <doi.org/10.1371/journal.pone.0008694>.
- Rouse, J. [1990]: 'The Narrative Reconstruction of Science', *Inquiry*, **33**, pp. 179–96.
- Santana, C. [2016]: 'What Is Language?', *Ergo*, **3**, available at <doi.org/10.3998/ergo.12405314.0003.019>.
- Schickore, J. [2008]: 'Doing Science, Writing Science', *Philosophy of Science*, **75**, pp. 323–43.
- Shibata, N., Kajikawa, Y. and Sakata, I. [2012]: 'Link Prediction in Citation Networks', *Journal of the American Society for Information Science and Technology*, **63**, pp. 78–85.
- Smaldino, P. E. and McElreath, R. [2016]: 'The Natural Selection of Bad Science', *Royal Society Open Science*, **3**, available at <doi.org/10.1098/rsos.160384>.
- Small, H. [1973]: 'Co-citation in the Scientific Literature: A New Measure of the Relationship between Two Documents', *Journal of the American Society for Information Science*, **24**, pp. 265–69.
- Soler, L., Zwart, S., Lynch, M. and Israel-Jost, V. [2014]: *Science after the Practice Turn in the Philosophy, History, and Social Studies of Science*, New York: Routledge.
- Suppe, F. [1998]: 'The Structure of a Scientific Paper', *Philosophy of Science*, **65**, pp. 381–405.
- Tahmasebi, N., Borin, L. and Jatowt, A. [2021]: 'Survey of Computational Approaches to Lexical Semantic Change Detection', in N. Tamahsebi, L. Borin, A. Jatowt, Y. Xu and S.

- Hengchen (eds), *Computational Approaches to Semantic Change*, Berlin: Language Science Press, pp. 1–91.
- Teubert, W. [2005]: ‘My Version of Corpus Linguistics’, *International Journal of Corpus Linguistics*, **10**, pp. 1–13.
- van Raan, A. F. J. [2004]: ‘Measuring Science’, in H. F. Moed, W. Glänzel, and U. Schmoch (eds), *Handbook of Quantitative Science and Technology Research*, Dordrecht: Springer, pp. 19–50.
- Waters, C. K. [2008]: ‘Beyond Theoretical Reduction and Layer-Cake Antireduction: How DNA Retooled Genetics and Transformed Biological Practice’, in M. Ruse (ed.), *The Oxford Handbook of Philosophy of Biology*, Oxford: Oxford University Press, pp. 238–62.
- Wilkinson, D. M. and Huberman, B. A. [2004]: ‘A Method for Finding Communities of Related Genes’, *Proceedings of the National Academy of Sciences USA*, **101**, pp. 5241–48.
- Wimsatt, W. C. [2007]: *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*, Cambridge, MA: Harvard University Press.
- Worrall, J. [1989]: ‘Structural Realism: The Best of Both Worlds?’, *Dialectica*, **43**, pp. 99–124.
- Wouters, P. [1998]: ‘The Signs of Science’, *Scientometrics*, **41**, pp. 225–41.
- Wray, K. B. [2020]: ‘How Is a Revolutionary Scientific Paper Cited? The Case of Hess’ “History of Ocean Basins”’, *Scientometrics*, **124**, pp. 1677–83.