TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science 00 (2024) 1–27

ISSN: 1756-8765 online
DOI: 10.1111/tops.12777

This article is part of the topic "Fellows of the Cognitive Science Society," Andrea Bender (Topic Editor).

How Do Scientists Think? Contributions Toward a Cognitive Science of Science

Nancy J. Nersessian[a,b]

[a]*College of Computing and School of Public Policy, Georgia Institute of Technology*
[b]*Learning Sciences, Harvard Graduate School of Education, Harvard University*

Received 11 June 2024; received in revised form 17 November 2024; accepted 22 November 2024

Abstract

Scientific thinking is one of the most creative expressions of human cognition. This paper discusses my research contributions to the cognitive science of science. I have advanced the position that data on the cognitive practices of scientists drawn from extensive research into archival records of historical science or collected in extended ethnographic studies of contemporary science can provide valuable insight into the nature of scientific cognition and its relation to cognition in ordinary contexts. I focus on contributions of my research on analogy, model-based reasoning, and conceptual change and on how scientists enhance their natural cognitive capacities by creating modeling environments that integrate cognitive, social, material, and cultural resources. I provide an outline of my trajectory from a physicist to a philosopher of science to a hybrid cognitive scientist in my quest to understand the nature of scientific thinking.

Keywords: Scientific cognition; Conceptual change; Analogy; Mental simulation; Model-based reasoning; Distributed cognition; Ethnography; Qualitative methods

Correspondence should be sent to Nancy J. Nersessian, College of Computing and School of Public Policy, Georgia Institute of Technology, Atlanta, GA 30332, USA. E-mail: nancy.nersessian@gmail.com

The whole of science is nothing more than a refinement of everyday thinking.
Albert Einstein, Physics and Reality

1. Introduction

Scientific thinking is one of the most sophisticated expressions of the creative potential of human cognition. It demonstrates what it is possible to accomplish when people attempt to solve complex problems in materially, socially, and culturally rich environments with critical and reflective use of their cognitive capabilities. Yet, how scientists think has thus far played a limited role in research in cognitive science. By and large, our understanding of human cognition has not been informed by thinking of the complexity, sophistication, and reflectiveness seen in science. The study of human cognition has been largely confined to and shaped by controlled experimental research on the cognitive development of children. These studies investigate such phenomena as core cognition, conceptual change, executive functions, and social cognition. Such investigations are designed to provide data on how children reason, represent, imagine, understand, solve problems, learn, and so forth. If we accept, with Einstein, the hypothesis that scientific thinking lies on a continuum with mundane thinking (see also Langley, Simon, Bradshaw, & Zytkow, 1987; Nersessian, 1984b, 2008a), we can, indeed, learn some things about the cognitive basis of scientific thinking from such studies, as well as from the experimental research on mundane problem-solving (mostly by undergraduates). But we can also learn things about the potentialities of ordinary cognitive capabilities by studying the nature of the "refinement" at the scientific end of the continuum. The study of how scientists think provides a novel window on the mind—on what is possible at the highly creative end of the human continuum.

Scientific thinking is best studied in the context of problem-solving, which implicates numerous cognitive processes such as how scientists reason, represent, remember, decide, imagine, plan, understand, and learn. As with cognitive development in children, the domain of scientific thinking affords an opportunity to study cognitive processes that involve major changes in representation and understanding. Further, it provides significant data on metacognition, as evidenced in the articulation and reflective refinement of methods, reasoning strategies, and representational issues by scientists. Scientific investigation requires complex and integrated cognitive processes, that have most often been studied in isolation from one another in research on mundane thinking. It is also highly dependent on scientists creating and using artifacts through which to extend their possibilities for thinking about the world, and often requires coordinated collaboration among members of a research team. On the frontiers of science, especially, it addresses ill-defined problems and takes place in unstructured task environments that are created as the research moves along. In sum, investigations of how scientists think as they solve complex problems can yield valuable insights into the full potential of the human cognitive system.

In what follows, as a contribution to the Fellows Topic, I focus on my near 50-year research program that has aimed to contribute to a cognitive science of science. Philosophers have been participants in the field of cognitive science since its inception. Most cognitive scientists, when they think about philosophy's participation, point to the philosophy of mind and its important contributions to understanding the foundations of mind, as well as to clarifying central concepts of cognitive science. There have, however, also been significant contributions made by philosophers of science, who do not reflect on the foundations of mind or of cognitive science; rather, they aim to understand the nature of scientific cognition. William Bechtel, Lindley Darden, Ronald Giere, David Gooding, and Paul Thagard, for instance, have been philosophical contemporaries in this endeavor. Scientific thinking has also been studied in other fields within cognitive science, viz., cognitive psychology and Artificial Intelligence (AI). The experimental and AI studies, while valuable, have low ecological validity in that they leave out the material, social, and cultural resources that enable sophisticated scientific thinking (Klahr & Simon, 1999). Importantly, they also leave out the thinking that goes into formulating and reformulating scientific problems, into determining or developing methods that might address them, and into devising the requisite infrastructure for problem-solving.

A richer, more accurate account of scientific thinking requires the study of the actual problem-solving practices of scientists and the cognitive processes that are their bases. However, both the contents and the contexts of real-world scientific thinking pose challenges for cognitive science research. Study of scientific thinking requires methodological innovation, since it is not possible to use standard methods of controlled experimental design to gather and analyze these data. Specifically, it requires adapting the qualitative methods of social sciences, such as historical analysis or ethnographic methods, which have their own standards of rigor, to investigate real-world scientific thinking. Furthermore, something of the content and methods of the science needs to be understood, which presents more of a challenge for researchers than, for instance, to grasp the conceptual and methodological aspects of the science problems that children or undergraduates attempt to solve. But these challenges are not insurmountable, as the recent wave of research by psychologists in other areas toward the development of a general psychology of science demonstrates (see, e.g., Feist & Gorman, 1998; O'Doherty, Osbeck, Schraube, & Yen, 2019; Osbeck, Nersessian, Malone, & Newstetter, 2011). Philosophers of science often have the advantage of having a background in the physical or biological sciences. But, even without that, it is possible to develop conceptual knowledge sufficient to understand what scientists are thinking about, without being able to do the science. This form of expertise is what the sociologists Collins and Evans have called "interactional expertise" (Collins & Evans, 2015; Collins, Evans, & Weinel, 2016).

My own cognitive science research is situated in the philosophy of science. Although I was trained as a physicist and a philosopher, I developed into a hybrid cognitive scientist because the aim of my research to understand scientific thinking required that I draw on concepts, methods, and theories or frameworks from philosophy of science, history of science, cognitive psychology, AI, cognitive anthropology, and learning sciences, and reflexively, engage with these fields by feeding back into them analyses and hypotheses about cognition in the context of science. I have advanced the position that to understand the nature of scientific thinking requires an integrative analysis of authentic problem-solving by scientists.

"Integrative" here has a dual meaning. It signifies the use of an integrative method—drawing from the fields listed above—to create an account of scientific thinking as a system phenomenon that integrates cognitive, material, social, and cultural dimensions (hereafter "cognitive-cultural") of problem-solving. Such an account requires the researcher to gather data on the real-world problem-solving practices of scientists, which can only be obtained through detailed case analysis of the records, archival and public, of past science ("cognitive-historical" analysis) or through in situ studies that use ethnographic methods of data collection and qualitative analysis ("cognitive-ethnographic" analysis), each with their affordances and limitations. Importantly, for the cognitive scientist, the data provided by the case studies are not collected to provide an analysis, for instance, in support of a particular historical claim (although such an outcome is possible), rather they provide a basis for determining features of cognition in the context of a kind of human activity: scientific problem-solving. The way to understand the relation between the specific cases and the general hypotheses is not as that of inductive generalization. Generality is produced in a bootstrapping method commonly used in the sciences. In such bootstrapping processes, hypotheses are made within a context of current understanding, and then "are refined, made more specific, modified, or rejected in light of more constraining data… Surviving hypotheses are then tested against other data and other hypotheses to determine the extent of their validity" (Nersessian, 1991, 683). Bootstrapping is an incremental and iterative, open-ended processes of working back and forth between data and hypotheses/theory until a satisfactory accommodation is achieved.

In the following sections, I discuss important findings from each kind of analysis that I have carried out in research on scientific thinking in physics and in bioengineering sciences. After first sketching my origins as a cognitive scientist in Section 2, I briefly discuss selected findings from analyses directed toward two of the problems central to my research program: the nature of model-based reasoning and its role in conceptual change in Section 3[1] and the nature of cognitive-cultural integration in scientific problem-solving in Section 4.[2] My first problem is pertinent, especially, to the long-standing interests of cognitive science in analogical reasoning and in conceptual change. The second problem is pertinent to the more recent cognitive science interests in what might be called cognition in context: embodied, situated, and distributed cognition and, relatedly, cultural cognition. There are many insights about how scientists think that I could discuss, but by focusing on these main themes, I hope a wide range of cognitive scientists will see the value and importance of promoting research on scientific cognition for a richer understanding of cognition. More extensive analyses can be found in *Creating Scientific Concepts* (Nersessian, 2008a) and *Interdisciplinarity in the Making: Models and Methods in Frontier Science* (Nersessian, 2022), and in publications by other contributors to the cognitive science of science.

2. My origins as a cognitive scientist

I characterize the interdisciplinary path my research had taken in the quest to understand scientific thinking as "following the problem," coupled with leveraging some serendipity in that quest. As the philosopher Karl Popper noted, "We are not students of the same subject matter but students of problems. And problems may cut right across the border of any subject matter or discipline." (Popper, 1962, 67) Interdisciplinary research is a problem-driven enterprise. I would add that formulating the problem and determining the resources needed to advance the problem-solving are a major part of the process. My problem, as I would formulate it now is

> *Science is one of the most significant creative pursuits of humankind. How can we understand and account for the epistemic accomplishments of science given that scientists are limited beings and the natural world is vastly complex?*

This problem is itself complex and my research has focused on two interrelated components: representational innovation and methodological innovation. In unfolding the various dimensions of those, I have found that advancing the problem requires at least the combined resources of philosophy, psychology, history, and anthropology. How I came to be this kind of cognitive scientist is a long story and I highlight just a few transition points along the way.

I have always loved and been intrigued by science. I started out to be a theoretical physicist. I was passionate about math and science for as long as I can remember. I have a vivid recollection of my fifth-grade teacher telling us about Albert Einstein. What I recall is that he used Einstein to introduce the idea that mathematics could be used to understand the universe. What he said must have been inspirational, since I know from that point onward, I wanted to understand what Einstein had done. I wanted to be a theoretical physicist. In college, I was quite a good physics student, but I was disappointed that it was difficult to get my professors to discuss with me the conceptual issues I was interested in, notably about Einstein's work on general relativity.[3] I stumbled into philosophy of science through my choice to get an AB rather than a BS. The only problem was that the degree required the dreaded Introduction to Philosophy, which was notoriously difficult and which my fellow science students found not to their liking. I chose a course that fit an open time slot. It turned out to be an introduction like none other on offer. It was taught by Milič Čapek, a noted philosopher of space and time, and instead of the usual Plato, Aristotle, and Descartes, we read Poincaré, Einstein, and Reichenbach. I found nirvana. I realized there were also physicists in philosophy departments. Since my physics professors offered me no encouragement to pursue a career in theoretical physics, and philosophy offered another avenue to pursue my interests, I set off to study the foundations of the general theory of relativity in a philosophy department.[4]

The next bit of serendipity came when I became frustrated with graduate philosophy of science courses which seemed to have much to do with language, but little to contribute to understanding the nature of the science I was working on. My advisor, Howard Stein, gave me the advice that unfolded the rest of my saga: read the scientists who have worked on the frontiers and have made fundamental transformations in our understanding of nature—in the process they articulate and address deep epistemological, conceptual, and methodological issues. As a physics student, I had only read textbooks or recent publications, but reading the historical scientists, especially their more speculative archival materials: drafts, letters, diaries, and so forth, opened a whole new dimension of what philosophers now call scientific practice to me. Although the works I read by scientists, from its inception onward, were peppered with analogies, visual representations, and thought experiments, it was the works directly relevant to the origins of field theory, those of Michael Faraday and James Clerk Maxwell, that occasioned my "aha" moment about the critical importance of these. Their works were seminal contributions to the development of the field concept, and so directly pertinent to the origins of gravitational field theory. A scientific formulation of a field concept required scientists to determine the nature and production of the forces taking place in the space surrounding and between matter. The notion that forces are transmitted continuously, with a time delay, through space and time, rather than acting instantaneously at a distance was a completely novel concept in the context of 200 years of Newtonian physics and its development was essential to the creation of the non-Newtonian theories of electromagnetism, relativity, and quantum mechanics.

As I worked through Faraday's reasoning about his experimental findings, especially in his Diary (Faraday, 1932), I was struck by all the sketches that filled the margins, and puzzled about the role they played in his thinking, as well as the thought experiments he presented in addition to his laboratory experimental research. In working through the various phases of Maxwell's reasoning, he, too, made extensive use of Faraday's visual representations, while constructing his own, and made extensive use of analogies in the derivation of the field equations, along with providing commentary on how he was using analogy to develop novel mathematical representations. All these heuristics were present in the work of other field theorists, including Einstein. At the time (early 1970s), these heuristics were widely dismissed by both philosophers and historians of science as ancillary—mere aids—to scientific reasoning, understood as logical manipulations of propositional representations. However, the archival data, coupled with published records, provided evidence in favor of my hypothesis that these so-called "aids," themselves, constitute a genuine form of reasoning integral to the scientific problem-solving. This insight led to my quest to develop an explanation of how they function to create or deepen understanding, how they lead to novel insights, and, in some instances, lead to conceptual change. It was at this point I realized I was now more interested in understanding the nature of how scientists think than in what the science they developed said about the nature of the world. But an explanation had to wait, on the one hand, until I could work out more fully both the details of a highly complex process of concept formation and change over nearly a 100-year period (Nersessian, 1984b), and on the other, how analogy, visualization, and thought experimentation (generalized to mental simulation) could work together, in processes of what I called "model-based reasoning," to promote concept formation and change (Nersessian, 1988, 1992, 1999, 2002a, 2008a). The latter led to my interaction with the newly developing field of cognitive science. I will discuss this research in the next section, but before that, I set the problem situation in epistemology and philosophy of science, with which cognitive scientists might not be familiar.

My training as a physicist made me skeptical as a philosophy student about the value of using only the tools of abstract philosophical analysis to address problems about the nature of the epistemic practices of scientists. The proposal to "naturalize" epistemology by W.V.O. Quine (Quine, 1969) provided some license to recruit resources from history and the sciences to address epistemic issues in conjunction with philosophical analysis. A naturalist stance in philosophy of science holds, basically, that (1) a philosophical account should be informed by the best available scientific understanding of humans that the biological, psychological, and social sciences offer; (2) it should be informed by data on the actual investigative practices, as they are created, used, and justified by scientists; and (3) it should make use of appropriate empirical methods to determine these practices. My research began in a context where the problem of conceptual change was a major issue in philosophy of science. The "historicist" philosophers, notably, Thomas Kuhn (Kuhn, 1962) and Paul Feyerabend (Feyerabend, 1962, 1970), were leading the charge against positivist conceptions of science that were based on logical analysis of the language of science and of "rationally reconstructed" science. They argued that to understand conceptual change in science, philosophers needed to draw from the records of the history of science and from theories of human psychology. Their approach, while moving in the right direction, seemed limited to me in at least two ways (Nersessian, 1979, 1984a, 1984b). First, their claims about the seriousness of "incommensurability" (roughly, the inability of scientists committed to the theory before, and those to the theory after a "revolution" to understand one another) in conceptual change appeared to be based on a decidedly *un*historical approach. To substantiate their claims, they only looked at the endpoints, that is, at the initial and final products of a major conceptual change in science, rather than the processes through which the change came about. For instance, they noted the noncomparability of the concept of "mass" as constant in Newtonian theory and the concept of "mass" as varying with velocity in the special theory of relativity. Kuhn advanced the notion that the process of change consists of the collection of a series of anomalies, and then a new theory appears. Second, Kuhn and Feyerabend used the notion of a "Gestalt switch," drawn from the Gestalt psychologists' theories of human problem-solving (Köhler, 1929; Wertheimer, 1959), to support their claims about the sudden, nonrational nature of the change. This notion, however, is not compatible with fine-grained historical investigation. Investigation of archival and published records provides data, in the form of the reasoning, arguments, and justifications put forth by scientists, that establishes the reasoned nature of conceptual change. To get to relativity theory, for instance, scientists, trained in the Newtonian tradition, had first to develop electromagnetic field theory, and a fine-structure analysis shows this to be a reasoned, complex problem-solving process, spanning the work of many scientists.

In this context, I saw the problem of *how* novel concepts are formed as a fundamental, but neglected, dimension of conceptual change. To address this problem required not only a different kind of historical analysis but also a different kind of psychological theory. In looking for research that might be helpful, I came across *Plans and the Structure of Behavior* (Miller, Galanter, & Pribram, 1960), a formative contribution to cognitive science, which led me to track literature in that field as its research was developing along with my own. It was in this context I dubbed the approach I was using a "cognitive-historical method" (Nersessian, 1987), which other researchers interested in scientific thinking had begun to use as well (see, e.g., philosophers [Darden, 1991; Gooding, 1990; Thagard, 1992] and cognitive psychologists [Gorman & Carlson, 1990; Tweney, 1985]).

3. Model-based reasoning and conceptual change

I have used cognitive-historical (Nersessian, 2008a) and cognitive-ethnographic (Nersessian, 2022; Nersessian & MacLeod, 2022) methods to carry out investigations into the reasoning and representational practices of scientists, especially those leading to innovation. On the one hand, to understand scientific thinking requires a fine-grained examination of problem-solving practices as evidenced in data from historical records or from detailed study of in situ practices of contemporary scientists. The historical dimension of cognitive-historical analysis examines the extant records, published and archival, of the research of one or more scientists over the pertinent period. For instance, in my analysis of the formation of the field concept, I examined the records of scientists who contributed to this representational problem from Faraday through to Einstein (Nersessian, 1984b).[5] The method locates scientists within the problem situations (intellectual, social, cultural, material) of their local communities and wider cultural contexts to understand their accomplishments, to the extent possible, as a cognitive-cultural product of a problem-solving system distributed in time and across individuals and artifacts. However, more detailed, contemporaneous data are required to establish how such cognitive-cultural integration takes place during the processes of problem-solving. To this end, a cognitive-ethnographic method enables, first, collection of as comprehensive a data set of in situ problem-solving as time, human, and financial resources allow and, second, offers the potential to examine how scientists think with artifacts and together with others in ongoing problem-solving processes. For instance, to investigate innovative modeling practices under development in bioengineering sciences, I was fortunate to be able to build a research group that could collect and analyze extensive data sets from four research labs.

The cognitive dimension of both methods begins with a determination of what and how human cognitive capacities might underlie, facilitate, and constrain scientists' investigative practices, given current cognitive science understanding. It starts from the premise that despite clear differences, the cognitive practices scientists have created and developed for solving problems are rooted in cognitive practices humans employ all the time in more mundane forms of problem-solving to meet the challenges of everyday life and work. The pertinent research on mundane cognition is used to provide insight into both how the scientific practices accomplish problem-solving and how they diverge from the mundane findings. Assumptions, methods, and results from both kinds of research are subjected to critical scrutiny, with corrective insights and wider implications moving in both directions. In this way, the findings and theoretical analyses that derive from the investigation of real-world scientific practices not only provide insight into the highly creative end of human thinking, but also reflect on the more mundane end of the spectrum. For instance, within other fields in cognitive science, findings from my research on model-based reasoning in science have been used in research on conceptual change in cognitive development and learning, as well as in the National Academy of Science's Next Generation Science Standards, which promotes a model-based approach to science learning.

One way in which the study of scientific thinking makes an especially valuable contribution is that its complexity calls for accounts that integrate and unify cognitive processes that customarily have been treated as separate research areas in cognitive science (see Nersessian, 2008a, chapter 4). For instance, although analogical, visual, and simulative reasoning and representational processes require specific cognitive accounts, there is ample evidence in both the historical and contemporary cases that in scientific problem-solving they work in combination (model-based reasoning), and thus also require an integrated account. A simple example serves to motivate the need for a unified account.

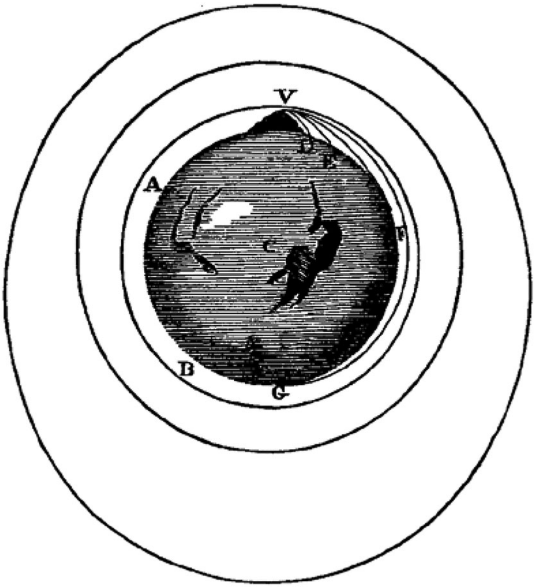

Fig. 1. Newton's rendering of a projectile thrown from a high mountain with increasing velocity. He provided text to direct the reader on how to mentally simulate the paths of the moving projectile (Newton, 1687, vol II, book III, 3).

Newton's famous analogy between the moon and a projectile provides a demonstration of how all three—visualization, analogy, and imaginative simulation—work together. Newton created an analogy between the motion of a cannon ball, shot with successively greater velocity from a mountain rising high above the surface of the earth, and the orbit of the moon. Here, Newton used his prior analysis of projectile motion as a source analogy. Newton provided a diagram (Fig. 1) that represents the analogy in a visual format, along with text, that guides the reader through a thought experiment in which one imaginatively simulates successive paths of the projectile, as its velocity increases, ending with the escape velocity where the projectile, too, would orbit the earth under the effect of centripetal force (gravity). Although the diagram is specific, it is understood to represent generic structures and processes that can be represented mathematically. The inference that the forces which keep the moon in orbit and the projectile directed toward the earth are the same marked a radical departure from prior concepts of earthly and celestial bodies as different in kind and contributed to Newton's creation of the mathematical concept of universal gravitation, which conceives all bodies as made of the same materials as earthly bodies, and, for analysis, casts all bodies as point masses.

The focus of my initial cognitive-historical research was on the formation of the concept of electromagnetic field, especially the contributions of the Faraday–Maxwell pair. It was the work on Faraday and Maxwell that first made me realize the need for a unified account. Briefly, one analysis I conducted showed, how, using the sketches in the margins of his research diary, Faraday reasoned about the deflection of a magnetized needle hung near a current-carrying wire by means of analogy and imaginative simulation through ongoing modifications of these sketches (visual representations) (Nersessian, 1984b, chapter 3). The sketches instantiated a model of the motion, which was refined through iterative interaction with the experiment. In the experiment, the motions appeared to be linear back and forth movements. However, Faraday used his sketching process to arrive at the hypothesis that the needle was in fact rotating around the wire. With considerable further work, this hypothesis contributed to the development of his field concept, as well as to the first electric motor. In the case of Maxwell, I constructed a detailed analysis of how, through iterative model-building processes Maxwell used resources from continuum mechanics and machine mechanics, in conjunction with known experimental findings on electricity and magnetism, to construct a series of hybrid analog models. These were accompanied by diagrammatic representations and written instructions for how to animate them mentally (Nersessian, 2002b, 2008a, chapter 2). They are hybrid in that they integrate information abstracted from analyses of mechanical phenomena and of electromagnetic phenomena. Through this process, Maxwell was able to construct a set of unified mathematical representations of the electromagnetic field concept, and put forward the hypotheses that electromagnetic actions are transmitted continuously through space at the speed of light, and that light, itself, is an electromagnetic phenomenon. As with the case of Newton, Maxwell's electromagnetic field concept radically transformed the accepted understanding of electric and magnetic forces as different in kind, and are thus not interactive, actions-at-a-distance, and unrelated to light.

The data provided by their historical records, especially the archival materials—drafts, letters, diaries, notebooks, artifacts, and so forth—opened an entirely new perspective on scientific reasoning and conceptual change, which, after years of research, can be summarized as follows. In this case, as in numerous others across the sciences, the processes underlying the accessible phenomena pertaining to complex, dynamical systems scientists wish to investigate are often inaccessible, either practically or in principle. Instead, they create models that provide a means of thinking and making inferences about the target phenomena. These models (conceptual, physical, or computational) provide structural, behavioral, or functional analog representations of the phenomena of interest. Scientists reason about the target phenomena by manipulating the models. I have argued that the human capacity to construct and manipulate iconic mental models provides the cognitive basis for this scientific problem-solving practice (Nersessian, 1992, 2002a, 2008a; Nersessian, 2008b). The mental models are schematic, partial representations that have correspondences with the real-world models, and interact in problem-solving processes to create reciprocal modification. Here, I can only highlight one important feature of such model-based reasoning: models are created specifically to serve as analogical sources for target phenomena.

This analogical practice in science is quite unlike any considered in the philosophical and cognitive science literatures. Usually, analogy is cast as a process of making sense of what we do not understand (target) in terms of what we do (source). In the scientific case, often little is understood about either the source (model) or target (real-world phenomena) at the outset. In analogical problem-solving, as customarily understood, the reasoner retrieves a previously solved problem that, with some modification, provides a source analogy, determines a mapping between source and target, transfers features from source to target, and evaluates inferences with respect to the target domain. What we customarily understand as analogy does, of course, occur in science. However, when scientists grapple with complex, novel phenomena, it is often the case that no ready-to-hand analogical source exists to provide a model. Instead, scientists need to "build" the source analogy through iterative and incremental processes of design, construction, manipulation (simulation/experimentation), and redesign in interaction with the goals and constraints of the target problem. Such model-building is a bootstrapping process that furthers the articulation of the problem as well as a potential solution. Each iteration of the model provides data and these insights to be worked into the next iteration. In this fashion, the modeling process serves to continually improve the creation of the source analogy until a satisfactory mapping is established.

There are several sources of data on scientific problem-solving, including historical, think-aloud protocol, and ethnographic that provide ample evidence of this important *representation-building* aspect of analogy (see, e.g., Nersessian, 1995; Nersessian, 2008a, 2022; see also Clement, 2008; Clement, 2022). For instance, our research on modeling methods in various fields in bioengineering sciences analyzed how, as one respondent expressed, researchers "build models to predict – or you hope will predict – what happens in real life." These analog models take the form of either living in vitro simulation models, which are composed of selected tissues or cells and nonliving engineered materials, or computational simulation models that are built with data from biosciences and engineering methods for modeling human-made systems, adapted to the requirements of biological systems. Importantly, our analyses showed how in conducting experiments (simulations) with the various source models, novel structures and behaviors can emerge, which, when represented can lead to novel concepts, and processes of conceptual change. Of course, there is no recipe for how to build a source model or any guarantee that model-building will lead to viable candidate problem solutions, but there are numerous cases in which these processes have led to profound and verified scientific achievements.

There are many interesting features of model-based reasoning in service of building the analogical source. But given space limitations, what follows is a general account, derived from analyses of data from studies across the sciences, of the most salient features:

1. building processes are goal-directed
2. models are built toward instantiating features germane to the epistemic goals, and are evaluated on this basis
3. omitted or negative features can provide resources for further development
4. elements used in building models derive from the target, one or more source domains, and from the model itself
5. various abstractive processes are used to select features to be instantiated and to integrate data, constraints (mathematical, physical, computational), or materials from feature 4 into a model
6. several iterations are usually needed to build a satisfactory model
7. each iteration provides insights that serve to improve the analogy until a satisfactory mapping is established or shown not to be feasible

8. interaction between target domain, source domain(s), and model is ongoing in the building process
9. mappings between the model and the target are established during the building processes, so in most cases, mappings develop over time

Some of these features need a bit of explication in this abbreviated context. With respect to 5, although "abstraction" is commonly used for a separate process alongside "idealization" and other abstractive notions, it is preferable, in general, to use the term as a comprehensive notion comprising various processes, including idealization, approximation, simplification, omission, limiting case, and generic modeling. All these abstractive processes can play a role in building analogical sources. With respect to 9, although for structural analogies, such as in the Maxwell case, mapping relational structure is important for the productivity of the analogy, as well as its evaluation (Gentner, 1983), for behavioral or functional analogies, replication of the behavior or function of interest is it what makes for a good and productive analogy. For example, with an in vitro model of an artery, the selected tissues and cells in the model need to replicate salient real-world features such as gene expression when exposed to blood forces. Finally, 4–8 are significant for my bootstrapping account of how problem-solving through model construction can lead to conceptual change. Iterative incorporation of elements from multiple domains into models creates novel, hybrid representations that can provide the basis for concept formation and conceptual change (see Carey, 2009; De Benedetto & Poth, 2024; Nersessian, 2008a; Nersessian & Chandrasekharan, 2009).

4. Cognitive-cultural integration

To understand how scientists think, it is necessary to consider the contributions of the material, social, and cultural environments that enable such sophisticated cognition. As the philosopher Daniel Dennett has succinctly stated about cognition, generally, "Just as there is little carpentry you can do with your bare hands, there is little thinking you can do with your bare mind" (Dennett, 2000, 17). Having worked for some time at the interface of philosophy of science, history of science, social studies of science, and cognitive science, in the 1990s, I began to think about ways of bridging "the cognitive-cultural divide" (Nersessian, 2005). Accounts of scientific practices have tended to focus on either "cognitive/rational" factors or "social/cultural" factors. On the one hand, philosophical and cognitive science studies focus largely on individual scientists, the rational practices and standards they use to carry out and justify experimental and theoretical research, and the cognitive representations and processes they use in problem-solving. On the other hand, historical and social studies of scientific practice focus instead on interests, motivations, and a range of social, cultural, and material factors in play within communities of scientists. This interpretive divide lies in a mix of complex issues rooted in 20th-century intellectual history. Here, it suffices to agree that any such divide, though at times analytically useful, imposes boundaries on dimensions that are inherently integrated in scientific practice.

My endeavor to develop an analytical framework that facilitates analysis of these dimensions in relation to one another led me to examine "environmental perspectives" in cognitive science—accounts of cognition as an embodied, artifact-using, and situated process. Environmental perspectives seek to understand cognition as inherently cultural, that is, as being shaped in ongoing processes of evolutionary history, of the development of the human child, and of the various sociocultural environments in which learning and work take place. In reading that literature in cognitive psychology, cognitive anthropology, learning sciences, and philosophy of mind, I was struck by a statement by Edwin Hutchins: "Humans create their cognitive powers by creating the environments in which they exercise those powers" (1995, 169). It led me to think that a way to analyze cognitive-cultural integration in scientific problem-solving would be to investigate in situ how scientists create environments to build the cognitive powers they need to realize their epistemic aims.

Since the inception of what we now know as "science," scientists have created artifacts through which to think about, investigate, and represent nature (e.g., Newton's calculus or Galileo's inclined planes or Maxwell's diagrams). They have also worked together with other scientists, either locally or through correspondence to think about research problems. However, as a source of data, historical science is limited in the extent to which it is possible to unearth and examine how the cognitive-cultural resources that scientists create are integrated in formulating, refining, and solving problems. The material dimension is largely confined to textual and diagrammatic representations, or, for an experimentalist such as Faraday, instruments that might be left behind. As for the social and cultural dimensions, one can locate scientists within their historical problem situation, such as Maxwell in the context of Cambridge physics. However, richer, contemporaneous data are required to establish how integration takes place in ongoing problem-solving. With cognitive-ethnographic research, one can develop a rich data archive on scientific problem-solving as it is occurring. Interdisciplinary science provides a particularly good locale in which to study the creation of scientific environments because concepts, methods, materials, and epistemic norms and values from different domains need to be fit together to formulate and address research problems of interest to the participatory domains. My research group has conducted cognitive-ethnographic studies of four bioengineering research labs in different areas. These are pioneering labs that bring together resources from the quite disparate fields of engineering, biological science, and computational science to address problems the researchers claim "no one has attempted before."

4.1. *Extending the framework of distributed cognition to science*

The analytical framework of distributed cognition (hereafter "D-cog") provides a good starting point for a cognitive-cultural analysis for the kinds of problem-solving that take place in science. It incorporates all the dimensions advanced by the range of environmental perspectives, and so scientists are cast as embodied agents who create and work within complex distributed cognitive-cultural systems. The D-cog perspective offers an analytical framework for studying scientific problem-solving practices, in general, in all the contexts in which science is and has been done. The grain size of the system relevant to the analysis—individuals, groups, or groups of groups with their attendant artifacts and sociocultural structures—depends on the focus. I began research within this framework in the philosophy and history of science in the late 1990s. Work in cognitive anthropology and the learning sciences also examines scientific practices from a D-cog perspective and helps to establish the broader fruitfulness of this approach for a cognitive science of science (see, e.g., Alac & Hutchins, 2004; Becvar, Hollan, & Hutchins, 2008; Charbonneau, 2013; Goodwin, 1995; Hall, Stevens, & Torralba, 2002, 2010). However, the D-cog analytical framework has developed largely in contexts quite different from those in which scientists work (see, e.g., Hutchins, 1995; Lave, 1988), and there are ways the framework needs to be broadened to accommodate scientific practices, especially for the philosopher of science. I note three main ways, although more might be identified as research on science develops.

First, much of the research that has contributed to the initial framework has focused on well-defined task environments, with problems and goals specified from the outset and cognitive artifacts ready-to-hand. Scientific environments, by contrast, are in ongoing development along all dimensions. For instance, in the environment of the cockpit, features of the problem-solving situations change in time, but how to land a plane has well-defined problems and goals, and the technological artifacts, the practices surrounding them, and the knowledge the crew brings to bear are relatively stable. Scientific environments, however, have open, ill-formed problems, tasks, and goals. Methods, knowledge, and problem formulations can all change in the course of problem-solving. Additionally, in the university setting, graduate students are conducting research as they are learning to become scientists. Such ongoing development is especially evident in frontier science environments, where not only are the problems and goals ill-defined, but the methods, artifacts, and social organizations needed for problem-solving are under development.

Second, as with all problem-solving, there are many cognitive capacities that come into play in science, for instance, memory, representation, reasoning, imagination, abstraction, and executive functions. Much theoretical and empirical work remains to be done on how these function within D-cog systems. For science, the interaction between representation and reasoning is especially important to thinking about the integral roles of artifact representations in problem-solving. Such external representations can be linguistic, mathematical, or visual, and they include gestures, physical models, computational models, and dynamic visualizations. Cognitive processes of mental modeling or mental simulation (which implicate many other mental resources) are especially relevant to thinking about the interaction between mental and artifact representations in inferential processes. One bioengineering researcher aptly characterized the relationship between mental representations and the artifact models they create as "putting a thought into the benchtop to see if it works." There are several literatures in cognitive science that can aid our understanding of their interaction within a D-cog framework (see Nersessian, 2008a, chapter 3). For instance, as with research on qualitative physics and on mental animation (see, e.g., DeKleer & Brown, 1983; Hegarty & Just, 1994; Schwartz & Black, 1996), the notion of internal–external representational coupling is useful for characterizing the relationship between scientists' mental models and artifact models. Customarily, D-cog has characterized cognitive processes, such as memory, as "off-loaded" to the representational artifacts. "Coupling" is a more apt metaphor for representational interaction in scientific thinking, where each part of the system, mental model and artifact model, can improve as a result of inferential processes. “Coupling” can be understood on analogy with mechanics: heterogenous components interacting dynamically in a feedback loop to improve the function of the system, in this case, model-based inference in a D-cog system.

Third, unlike most other practices investigated in the D-cog framework, science has epistemic goals. Scientists need not only to solve problems, but also to justify their methods and outcomes. Although one can study problem-solving in cognitive science for various purposes without attending to its epistemological dimension, philosophical analyses of scientific thinking need to attend to the epistemic norms and values within the D-cog system that guide the principles and the considerations that warrant claims to have produced knowledge or understanding. Studying scientific problem-solving in the situations and contexts where it takes place can provide valuable insights into features of the environment pertinent to how norms, standards, and methods emerge, develop, and function in scientific communities.

4.2. The research lab as a problem space

“The lab” is often associated with the physical space that houses the research-specific technologies, instruments, artifacts, and workbenches, along with the researchers who work within it. For a D-cog system, however, the lab is not simply a physical space existing in the present, but rather a problem space, extended in space and time, constrained by the research program of the director that reconfigures itself continually as the research program moves along and takes new directions in response to what occurs both in the lab and in the wider community of which the research is a part. At any point in time, the lab-as-problem-space contains resources for problem-solving which include people, technologies, methods, concepts, knowledge resources (literature, the internet, data bases, and so forth), financial resources, epistemic norms and values, problems, relationships, and lab history in the form of artifacts and lore.

The overarching question of this part of my research has been how does cognitive-cultural integration take place in research labs on the frontiers of bioengineering science, where that environment is continually under construction as the research moves along. The frontier university labs I have studied are *adaptive* problem spaces in that researchers are continually in the process of merging resources from different fields to create the infrastructure to enhance their natural capacities for problem-solving. The choice of labs in this area was serendipitous. As Director of the Program in Cognitive Science, I was contacted by senior biomedical engineers who were seeking assistance to develop a new educational program suited to the aims and challenges of the 21st-century bioengineering sciences. I saw this as an opportunity to pursue my own research goals, since to help them develop the curriculum, I needed to understand the nature of their research problems and practices. Setting on a cognitive-ethnographic program of research with this dual purpose enabled me to receive sufficient funds from the United States National Science Foundation over a dozen years to fund a highly interdisciplinary research group—unusual for a philosopher—focused on “cognition and learning in interdisciplinary cultures.”

In what turned into a 20-year project, we investigated the epistemic practices of four university labs working on the frontiers of the bioengineering sciences. Bioengineering scientists aim to make fundamental contributions to basic biological research, with an eye to the future creation of novel artifacts and technologies for application, especially in medicine. In the labs we studied in tissue engineering, neural engineering, and integrative/computational systems biology, researchers were focused on understanding complex biological phenomena at the system level. For reasons of ethics or control, these phenomena need to be studied indirectly, though in vitro models or computational (in silico) models that can simulate the target phenomena under experimental conditions. Thus, the research centers on building models to provide analogical sources, as discussed above, that, ultimately, can be used to develop hypotheses about the behaviors of complex biological systems "in the real world." Once a model is evaluated to be a satisfactory source, an extended process which can take years of research, outcomes of experiments with the models are transferred as hypotheses to target phenomena, and evaluated to the extent possible given the state of science. Modeling methods are continuously undergoing improvement as materials, technologies, methods, and so forth advance.

The problem spaces of the four labs differed significantly, which provided the opportunity to focus on the roles of different environmental resources in our analysis of the problem-solving systems in each lab. Importantly, despite the differences, in each lab, during all phases of the building process—design, construction, evaluation, redesign, experimentation—the model is the locus of cognitive-cultural integration. Each lab created an environment, organized specifically for the kind of model it needed to build. The primary model(s) of the lab research are sites of intersection of biological and engineering concepts, methods, and materials, and of epistemic norms and values. They are the artifacts through which researchers think about and improve their understanding of the biological phenomena in repeated interactions with mental models (coupling). They are sites where processes of mentoring, identity formation, and learning take place, and where the history of the lab is learned and appropriated hands-on. So, too, they are the basis for the interaction with members of the wider community through presentations and publications, and efforts to gain funding and institutional support for the research.

Briefly, in each lab, we used the method of cognitive ethnography to study in situ investigative practices. This form of ethnography was dubbed "cognitive" (Hutchins, 1995) in that it focuses on problem-solving activities of individuals or groups as situated in real-world contexts, to determine the cognitive processes implicated in these. In our case, we investigated problem-solving situated in the ongoing development of modeling environments in communities of bioengineering scientists. In each lab, we collected ethnographic observations, conducted open and semi-structured interviews, and collected a range of archival data (grant proposals and papers at various stages of development, power point presentations, emails, wiki contents, writing and drawings on lab whiteboards, and so forth) over the course of 5 years. Notably, by collecting data over a sustained period, we were able to track the formation of problems and goals; to log the various methods, steps, and iterations of model-building; to ascertain specific concepts, theories, methods, and materials in use, and changes to these; to probe decisions and judgments behind the development and alteration of a specific model; to examine how and what kind of inferences an experimental simulation with such models enables; and to note interactions among the researchers relevant to problem-solving processes. Having multiple sources of data enabled us to corroborate findings ("triangulation") and construct robust, trustworthy accounts (Eisner, 2003; Lincoln & Guba, 1985). Finally, in data analysis, we used a variety of mutually complimentary qualitative methods: interpretive coding, thematic analysis, case study analysis, and cognitive-historical analysis. Such analyses aim to move from the specificity of the case to construct a broader interpretive account by using systematic procedures to abstract and coalesce interpretive categories, and, where appropriate, formulate candidate hypotheses to transfer across multiple cases (Geertz, 1983).

My research group has published numerous richly detailed accounts of how such cognitive-cultural integration take place as scientists create modeling environments in ongoing research. Since most cognitive scientists do not study science, I think it is important to provide an exemplar here, although only an overview is possible. As with all ethnographic research, the details of the case are specific to it, but the generalities listed above transfer across our cases in different labs with different kinds of modeling practices.

4.3. Creating a neural engineering modeling environment

The example I use to anchor the notion of building modeling environments to create cognitive powers in a D-cog system took place over 5 years. It is drawn from a detailed case study of a pioneering neuroengineering lab that was seeking to understand learning in living networks of neurons. The full case study provides an examination of how the researchers integrated conceptual, methodological, and material resources from engineering, neuroscience, and computational science to create different kinds of distributed problem-solving environments that enhanced their natural capabilities, for instance, for reasoning, visualization, abstraction, imagination, and memory, to attain their epistemic aims. It is a complex story that focuses on the development and interaction of two analog source models: one, an in vitro "dish" model-system built to investigate learning in living networks of neurons; the other a computational simulation model of the in vitro model-system built to better-understand its dynamics. So, the in silico model provides a second-order analogy. Chapters 2 and 3 of my 2022 book provide a detailed analysis of how these model-systems were built, the D-cog systems they created, and the research they enabled.

Prior research on learning in living neurons had been conducted on single neurons. The lab director argued that since learning in the brain involves dynamic processes of synaptic growth in response to electrical signals, it needed to be studied in living networks. Before setting up his lab, the director had done postdoctoral research in a lab that was building an in vitro model-system to enable experimentation on, and real-time imaging of, living neuron networks. The in vitro dish is a hybrid model-system, composed of embryonic rat cortical neurons (approximately 40,000), dissociated, and plated on a specially designed grid of 64 electrodes, called a "multi-electrode array" (MEA). The neurons generate new connections to become a living network. At the time the director established the neural engineering lab, this was a completely new kind of model-system, and the lab was one of the first to investigate its properties and behavior. They began by developing software to "communicate" with the dish by sending and receiving signals from it ("open-loop physiology") and proceeded to develop computational and robotic "embodiments" through which the dish might learn from "sensory" feedback, as the brain does from the body ("closed-loop physiology"). All these model-systems were designed to function as source analogies to develop an understanding of the dynamics of learning in neural networks in the brain (target system). The researchers hypothesized that although the dish provides a highly abstract analogy to the target, if they could produce and control learning in the embodied dish model-system, this would provide insights into the neural mechanisms in the dish network that might be transferred and evaluated as mechanisms of learning in the brain in further research in neuroscience.

The lab conceptualized "learning" in terms of the Hebbian notion of learning as plasticity (basically, changes in the brain from adding or removing neural connections or adding cells in response to experience), the mathematical formulation known as the Hebbian rule ("neurons that fire together wire together"), and the standard notion of memory, which is the ability to retain and retrieve experiences. The researchers expected to modify the concept of plasticity and the associated equation, since the original concerns two neurons and they were investigating populations of neurons. They operationalized learning as "a lasting change in behavior resulting from experience." They argued that the research would contribute to understanding learning in neuron networks if they were able to demonstrate that they could build dish model-systems in which they could reliably create and control learning. Thus, their overarching problem became to develop a control structure for supervised learning in the in vitro embodied dish model-systems.

In the D-cog framework, the "cognitive powers" created by environments built for problem-solving are analyzed in terms of the affordances of the specific artifacts or human interactions. One of the affordances of the dish model-system are that as an abstraction, working with it enables the researchers to selectively focus on synaptic dynamics in response to electrical stimulation in various experimental setups. Others are that, with the neurons plated to an MEA array, it is possible to develop a means to visualize the dish's electrical signals, which enables perceptual inferences, to record experimental sessions, which enables researchers to revisit/recall the experiments, and to track the history of each dish, which provides a memory of its development. The latter is important, because the dish is a living system and so continually undergoing development over the course of its life, which could be as long as 2 years. In addition, the researchers could do optical imaging of real-time synaptic formation and, thus, see connections as they formed, which supports making perceptual inferences about these dynamics.

As part of the suite of software tools they developed to send, record, and analyze dish electrical signals, the researchers decided to create a visual display of neuronal behavior as electrical activity in a format similar to how engineers represent electrical signals on an oscilloscope, as an eight-by-eight grid that displays the electrical activity, as it is occurring over time, in each individual MEA channel (Fig. 2). This visualization shows the electrical activity of clusters of neurons around each electrode (64 electrodes, approximately 40,000 neurons).

Early on they encountered a problem: the MEAscope representation showed there to be continual spontaneous electrical activity taking place across the dish. They borrowed the concept of burst (spontaneous electrical activity) from single neuron studies, now generalized to a population of neurons to interpret the phenomena. Fig. 2 exhibits a pattern of bursting behavior as spikes in each electrode channel. They initially understood this behavior to

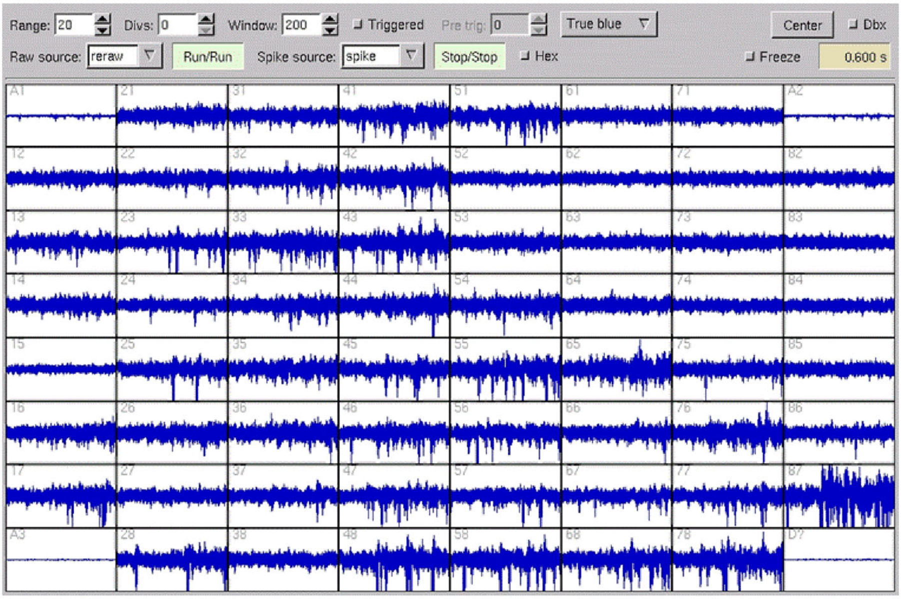

Fig. 2. A screen shot of the MEAscope per channel visualization of the in vivo dish activity showing spontaneous bursting (spikes) in the channels of the dish.

create a significant obstacle to their goal of getting the dish to learn. They used the engineering concept of noise to interpret the bursting behavior: “noise in the data – interference, it’s clouding the effects of learning we want to induce.” Such bursting behavior does not occur in a normally functioning adult brain, so one researcher, D4, focused her research on finding an electrical stimulation pattern (sensory input) that would eliminate bursts. It took a frustrating year, which included numerous failures to quiet the dish, but when she succeeded, a strange thing happened. They still were unable to get the quieted dish to learn when they tried numerous stimulation patterns. So, after 2 years of work, the learning research was at an impasse.

During this period, the other researchers were engaged in largely separate, but interrelated activities aimed at developing the embodied model-systems (Fig. 3, before the dashed line). D2 was working on the software module needed to translate signals between the dish and the motor commands to control its computational and robotic embodiments. One of the real-world embodiments, which was to figure prominently in the supervised learning research, was the robotic drawing arm he designed in collaboration with another research group. D11 had been working with him, but early in the burst-quieting period, he decided to branch away from the work with the in vitro model-system and develop a computational model that could simulate the behavior of the in vitro model. This computational dish model is a second-order model (or second-order analogy) built to gain insight into the behavior of the living dish model. Computational modeling had not been part of the practices of the lab because of the director’s experience with the limitations of neural network modeling, but D11 believed the affordances of this kind of model, in particular, that “you can measure everything, every detail of the

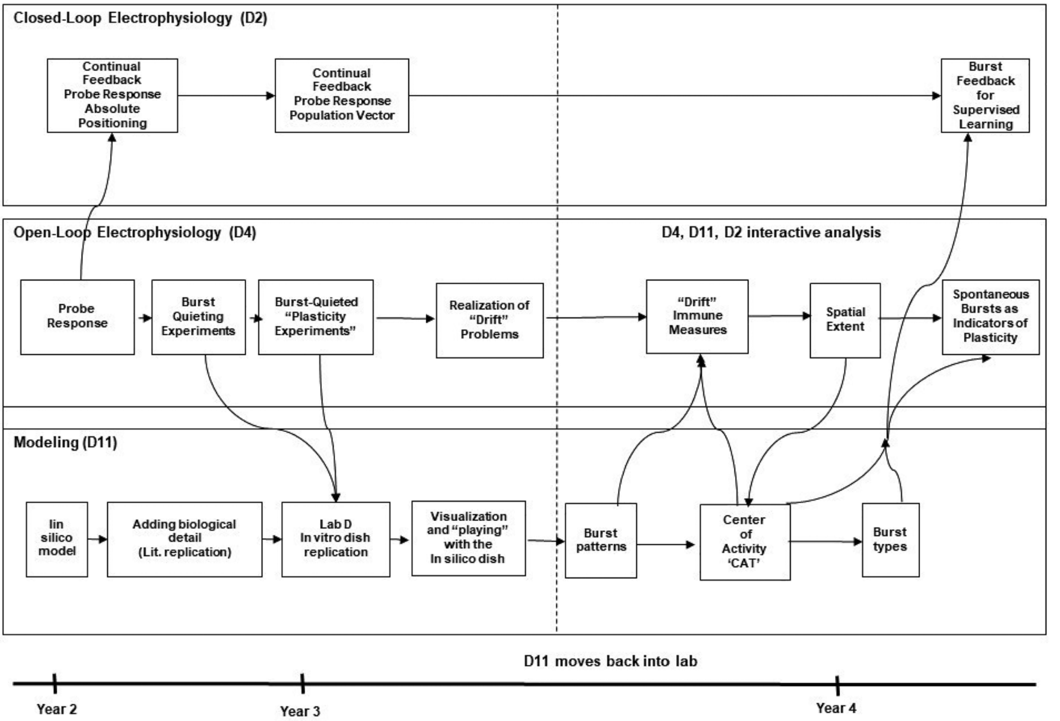

Fig. 3. My representation of Lab D as a distributed problem-solving system. The boxes represent models or model activity and researcher interactions with them. The arrows indicate interactions among the researchers. The period after the dotted line shows how the in silico model coalesced the researchers into a highly interactive system.

network," might provide "some new information about the problem [bursting and control] we could not solve at the time." In particular, he thought that the computational dish would enable him to "see" the activity at the level of the individual neurons, make precise measurements after every experiment, and run significantly more controlled experiments than were possible with the in vitro dish.

Building the computational model took many iterations and the processes are too complex to detail here (see Chandrasekharan & Nersessian, 2015; Nersessian, 2022). I highlight one major affordance of computational simulation models, which, in this case, proved quite instrumental to their effort to investigate bursting: the capacity for dynamic visualization. The modeler has significant discretion about how to visualize the model dynamics, for instance, D11 could have used the same kind of grid format they used for the in vitro dish (Fig. 2); however, he chose to visualize it the way he imagined it: as propagation of synaptic weight changes in a dynamic network of neurons, as shown in Fig. 4.

With this representation, as D11 expressed, you "can visualize fifty thousand synapses…so you can see…after you deliver a certain stimulation, you can see those distributions of synaptic weight change." He also made movies of the visualized behavior in numerous simulations and showed them to the other group members (and to us), so that everyone could see the

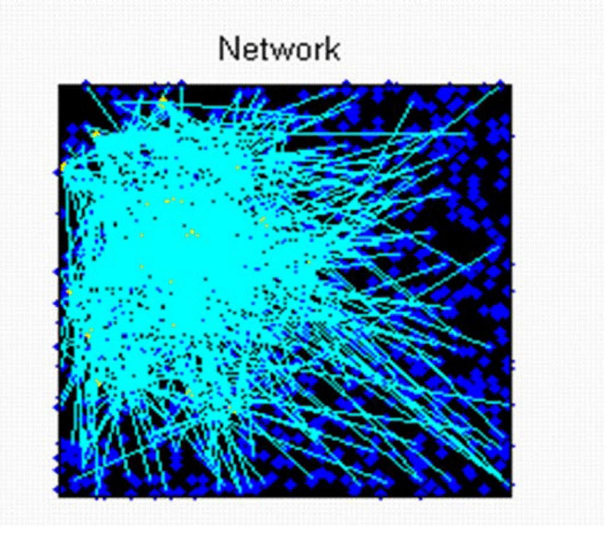

Fig. 4. A screen shot of the computational visualization of the burst activity of the in silico dish displayed on the computer screen as the propagation of synaptic weight changes across the network of neurons. The actual dynamic display would show how the activity moves across the computer screen.

behavior and come to an agreement that he might have discovered a remarkable feature of bursting behavior. D11 noticed that there were repeated spatial patterns within the bursting activity, which he characterized as "similar looking bursts," that propagated across the network. He was able to count, what the display showed to be, a limited number of what he called "burst types." The understanding he reached from these findings was that if the bursts were stable, then it might be possible to exploit bursts as signals, rather than eliminate them as noise in the data. The others agreed when they watched the videos, and they started working together to develop a way to track and mathematically formulate the activity of possible limited and stable bursting across the network.

One implication of this insight was that the group's understanding of bursts changed from noise in the data to signals that might be exploited to control the behavior of the dish, and indeed, with considerable further joint problem-solving through mathematical analyses and experimentation on both the in vitro and in silico models, they were able to create and control learning in the in vitro dish model-system. Again, this was a complex and extended problem-solving process, which included the development of several novel concepts for understanding neural network activity. They worked by first abstracting generic structure and processes from the in silico dish that could be represented mathematically, and, then, using that model as a source analogy for the target in vitro model, they were able, with appropriate adaptations, to develop and evaluate mathematical representations for a range of embodied dish model-systems. The researchers were able to combine these analyses and earlier stimulation techniques developed in the burst quieting research to formulate a set of novel stimulation patterns (a control structure), which led to the first instance of supervised learning in a living neuronal network in the field. As we were wrapping up, they began work to transfer this model-based research to real-world systems; specifically, to investigate in mice, with medical researchers, the hypotheses that Parkinson's disease is a bursting phenomenon.

As with conceptual and physical models (in vitro, in this case), computational models, as parts of D-cog systems, can provide environments for analogical, visual, and simulative reasoning in problem-solving processes, which, as in the case at hand, can lead to the formulation of novel scientific conceptual resources. However, computational models have cognitive, manipulative, and experimental affordances not available to other kinds of models. For instance, a computational model synthesizes a vast amount of data into a complex representation that can enact the dynamics of the target system. In effect, it provides the modelers with a running literature review from which to derive inferences that might otherwise be hidden in the disparate data. Researchers can run an unlimited number of experiments, which can include counterfactual scenarios that support the modelers in "what if…?" reasoning (a kind of thought experiment). They can take measures of significant variables at the level of detail often not accessible with the in vitro or in vivo target systems. They can visualize the dynamic behavior of the model in representations that support their thinking about the behaviors of three-dimensional phenomena across time, and visualizations can be recorded and viewed and compared repeatedly. Taken together, these affordances help researchers to form a global perspective on the phenomena. This global perspective is what informs the claim modelers often make to have developed "a feeling for the model," which facilitates predictive inferences about its behavior and, by analogy, about the potential behavior of the target system.

Specific to this case, the computational dish model provided a way for the researchers to envision the structure and behavior of the real-time propagation of the dynamic network activity. This visualization is significantly different from the per-channel MEAscope visualization, which does not capture the network behavior. As D4 later put it, "he [D11] was thinking like a wave, where we were thinking of a pattern." Of course, because learning in a network of neurons was the target of their investigation, everyone knew the dish activity was network activity. But no one had seen the network activity or a representation of it, so the computational visualization was built on a counterfactual scenario: "If we were able to see into the dish…" Indeed, the researchers spoke of this visualization as enabling them "to see into the dish." The visualization of the network activity and the capacity of the in silico model to run an unlimited number of simulations that could be recorded and played back repeatedly provided affordances that enabled first D11, and then the research group to notice the similar-looking patterns and work together to formulate their behavior mathematically. The manifest nature of the visualization served to align the mental models of the researchers, enabled the group to make and critique joint inferences, and facilitated joint exploitation of bursts as signals. Analyses of the computational model behavior, in interaction with the in vitro model, enabled them to modify and develop new conceptual resources. Finally, the computational model served as a driving force that brought the researchers together to form a highly effective distributed cognitive-cultural system. In particular, the in silico visualization generated many types of lab activity, which when put together, led to a solution of the problem of controlling the neuronal behavior, first in the computational dish model, then by analogical transfer and suitable modification in the target dish model-systems, and, then, to an extension of their research findings as hypotheses to neuron networks in real-world brains (target mice models).

Finally, there are many ways in which these two kinds of models provided a locus of cognitive-cultural integration in the neural engineering lab. For instance, concepts and methods from neuroscience, engineering (mechanical, electrical), optical imaging, and computational sciences, as well as epistemic values and norms, were combined and adapted in ways that enabled the researchers to build the various model-systems. Further, materials from biology and engineering were integrated to build the ontologically and epistemically hybrid in vitro models. And, as the research progressed, individual researchers who worked on different pieces of the research coalesced into a tightly connected distributed problem-solving system that achieved a major lab research goal. Finally, these graduate students, who arrived as an electrical engineer, a mechanical engineer, and a bioscientist, developed identities as hybrid neuroengineers (Osbeck & Nersessian, 2017).

In sum, this case exemplifies (when developed in full) what we found across all the labs we studied: one way in which scientists create their cognitive powers is by creating the modeling environments in which they exercise those powers.

5. Conclusion

I conclude with the short answer to my title question: scientists think not only with their minds/brains, but by creating environments rich in artifacts, developing problem-solving methods, and building communities. The focus of my research program continues to be on how scientists create various kinds of environments that enable them to think analogically about complex real-world phenomena: modeling environments. The findings and theoretical analyses that derive from research on real-world scientific problem-solving provide unique insight into the human capacity to design and utilize resource-rich environments at the highly creative end of the cognitive spectrum.

To conduct this research required my evolution from a physicist to a philosopher and historian of science to a hybrid cognitive scientist. My evolution has been a combination of following a continually developing problem wherever it led and a willingness to learn whatever was needed to tackle it. As a first step into cognitive science, as a fellow at the Pittsburgh Center for Philosophy of Science, I approached cognitive scientists at the Learning Research and Development Center there and at Carnegie Mellon University about a postdoctoral position. Surprisingly, Herbert Simon had an abiding interest in the work of Maxwell, and that got my foot in the door. He and Lauren Resnick worked out a postdoctoral position that opened the unanticipated opportunity to bring my research on model-based reasoning in science to bear on science education—work that has continued throughout my career. Since that time, I have resided in various interdisciplinary units in history of science, psychology, media and cultural studies, public policy, computing, architecture, and ultimately, the cognitive science program I developed and directed at Georgia Tech. These different environments provided a rich range of affordances I could make use of in my quest. At Georgia Tech, for instance, I had the opportunity to work with graduate students from a wide range of backgrounds together with colleagues to develop AI models of scientific thinking with Ashok Goel (e.g., Griffith, Nersessian, & Goel, 1996), to conduct psychological experimentation with Richard Catrambone (e.g., Craig, Nersessian, & Catrambone, 2002), to conduct problem-solving protocol studies and ethnographic investigations of architectural design cognition with Craig Zimring and Chuck Eastman (e.g., Dogan & Nersessian, 2012; Kasali & Nersessian, 2015; Yagmur-Kilimci, 2010), and to conduct ethnographic investigations with Wendy Newstetter and Lisa Osbeck (e.g., Nersessian, Kurz-Milcke, Newstetter, & Davies, 2003; Osbeck & Nersessian, 2006; Osbeck & Nersessian, 2017; Osbeck et al., 2011). Because of its location in an engineering university, the program focused on cognition in real-world contexts of work and of learning. As a result, I learned much not only about traditional cognitive science, but also about environmental perspectives from my colleagues. All these experiences furthered my contributions to the cognitive science of science. Finally, my research has enabled me to use what I have been learning to contribute to the efforts of science and engineering faculty and learning researchers to develop educational programs in science and engineering for students across the educational lifespan, based on how scientists think in conducting their research.

Notes

1 As an interdisciplinary researcher, I have inhabited many academic positions in search of a "home." It took 18 years to find a tenured position. As requested by a reviewer, I note the major institutions in which I was situated during various phases of my career. The research in Section 3 was conducted mostly when I was a faculty member in the Program of History of Science and Department of History at Princeton University. There I also worked with Gilbert Harman (Philosophy) and George Miller (Psychology) to develop a Program in Cognitive Science.

2 The research in Section 4 was conducted mostly when I was a faculty member with a "cognitive science" position split among multiple Schools and Colleges at Georgia Institute of Technology. I was hired to develop a Program in Cognitive Science.

3 I received my AB with a double major in Physics and Philosophy from Boston University. I supported my undergraduate education by working on the Apollo 11 project at the MIT Instrumentation Lab, Displays and Human Factors Group. I note this because working with engineers who were focused on the interactions among the astronauts, the onboard computer guidance system, and the various ground control stations was a formative experience that influenced my later thinking.

4 Case Western Reserve University had just started a new Ph.D. program funded by the Rockefeller Foundation to train students with undergraduate degrees in the sciences or mathematics in the "philosophy of" that subject. My Ph.D. research was focused on the foundations of physics.

5 At this early post-Ph.D. point in my career, I was in the Netherlands, where I went with as a Fulbright Scholar to Leiden University and the Museum Boerhaave to investigate the contributions of physicist H. A. Lorentz to the formation of the field concept. After that, I accepted a part-time position on the faculty at the Technical University of Twente in the History, Philosophy, and Sociology of Science Unit.

References

Alac, M., & Hutchins, E. (2004). I see what you are saying: Action as cognition in fMRI brain mapping practice. *Journal of Cognition and Culture*, *4*(3–4), 629–662.

Becvar, A., Hollan, J., & Hutchins, E. (2008). Representing gestures as cognitive artifacts. In M. S. Ackerman, C. Halverson, T. Erickson, & W. A. Kellog (Eds.), *Resources, co-evolution, and artifacts: Theory in CSCW* (pp. 117–143). New York: Springer.

Carey, S. (2009). *The origin of concepts*. Oxford University Press.

Chandrasekharan, S. & Nersessian, N. J. (2015). Building cognition: The construction of computational representations for scientific discovery. *Cognitive Science*, *39*, 1727–1763.

Charbonneau, M. (2013). The cognitive life of mechanical molecular models. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *44*(4, Part A), 585–594.

Clement, J. J. (2008). *Creative model construction in scientists and students*. Springer.

Clement, J. J. (2022). Multiple levels of heuristic reasoning processes in scientific model construction. *Frontiers in Psychology*, *13*, 750713. https://doi.org/10.3389/fpsyg.2022.750713

Collins, H., & Evans, R. (2015). Expertise revisited, Part I—Interactional expertise. *Studies in History and Philosophy of Science Part A*, *54*, 113–123. https://doi.org/10.1016/j.shpsa.2015.07.004

Collins, H., Evans, R., & Weinel, M. (2016). Expertise revisited, Part II: Contributory expertise. *Studies in History and Philosophy of Science Part A*, *56*, 103–110. https://doi.org/10.1016/j.shpsa.2015.07.003

Craig, D. L., Nersessian, N. J., & Catrambone, R. (2002). Perceptual simulation in analogical problem solving. In L. Magnani & N. J. Nersessian (Eds.), *Model-based reasoning: Science, technology, values* (pp. 167–190). New York: Kluwer Academic/Plenum Publishers.

Darden, L. (1991). *Theory change in science: Strategies from Mendelian genetics*. New York: Oxford University.

De Benedetto, M., & Poth, N. (2024). Bootstrapping concepts via hybridization: A step-by-step guide. *Review of Philosophy and Psychology*, https://doi.org/10.1007/s13164-024-00752-x

DeKleer, J., & Brown, J. S. (1983). Assumptions and ambiguities in mechanistic mental models. In D. G. a. A. Stevens (Ed.), *Mental models* (pp. 155–190). Hillsdale, NJ: Lawrence Erlbaum.

Dennett, D. C. (2000). Making tools for thinking. In D. Sperber (Ed.), *Metarepresentations: A multidisciplinary perspective* (pp. 17–29). New York: Oxford University Press.

Dogan, F., & Nersessian, N. J. (2012). Conceptual diagrams in creative architectural practice: The case of Daniel Libeskind's Jewish Museum. *Architectural Research Quarterly*, *16*(01), 15–27.

Eisner, E. W. (2003). On the art and science of qualitative research in psychology. In P. Carnic, J. Rhodes, & l. Yardly (Eds.), *Qualitative research in psychology: Expanding perspectives on methodology and design* (pp. 17–29). Washington, DC: American Psychological Association.

Faraday, M. (1932). *Diary*. London: G. Bell & Sons Ltd.

Feist, G. J., & Gorman, M. (1998). The psychology of science. *Review of General Psychology*, *2*, 3–47.

Feyerabend, P. (1962). Explanation, reduction and empiricism. In H. Feigl & G. Maxwell (Eds.), *Minnesota studies in the philosophy of science* (Vol. III, pp. 28–97). Pittsburgh, PA: University of Pittsburgh Press.

Feyerabend, P. (1970). Against method: An outline of an anarchistic theory of knowledge. In H. Feigl & G. Maxwell (Eds.), *Minnesota studies in the philosophy of science* (Vol. IV, pp. 17–130). Minneapolis, MN: University of Minnesota Press.

Geertz, C. (1983). Thick description: Toward an interpretive theory of culture. In R. Emerson (Ed.), *Contemporary field research* (pp. 37–59). Prospect Heights, IL: Waveland Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.

Gooding, D. (1990). *Experiment and the making of meaning: Human agency in scientific observation and experiment*. Dordrecht: Kluwer.

Goodwin, C. (1995). Seeing in depth. *Social Studies of Science*, *25*, 237–274.

Gorman, M. E., & Carlson, W. B. (1990). Interpreting invention as a cognitive process: The case of Alexander Graham Bell, Thomas Edison, and the telephone. *Science, Technology, and Human Values*, *15*, 131–164.

Griffith, T. W., Nersessian, N. J., & Goel, A. K. (1996). The role of generic models in conceptual change. In G. W. Cottrell (Ed.), *Proceedings of the 18th Annual Conference of the Cognitive Science Society* (pp. 312–317). Lawrence Erlbaum.

Hall, R., Stevens, R., & Torralba, T. (2002). Disrupting representational infrastructure in conversation across disciplines. *Mind, Culture, and Activity*, *9*, 179–210.

Hall, R., Wieckert, K., & Wright, K. (2010). How does cognition get distributed? Case studies of making concepts general in technical and scientific work. In M. Banich & D. Caccmise (Eds.), *Generalization of knowledge: Multidisciplinary perspectives* (pp. 225–246). New York: Psychology Press.

Hegarty, M., & Just, M. A. (1994). Constructing mental models of machines from text and diagrams. *Journal of Memory and Language*, *32*, 717–742.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Kasali, A., & Nersessian, N. J. (2015). Architects in interdisciplinary contexts: Representational practices in healthcare design. *Design Studies*, *41*, 205–223.

Klahr, D., & Simon, H. A. (1999). Studies of scientific discovery: Complimentary approaches and divergent findings. *Psychological Bulletin*, *125*, 524–543.

Köhler, W. (1929). *Gestalt psychology*. Oxford, England: Liveright.

Kuhn, T. (1962). *The structure of scientific revolutions: International encyclopedia of unified science* (Vol. 2). Chicago, IL: University of Chicago Press.

Langley, P., Simon, H. A., Bradshaw, G. L., & Zytkow, J. M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge, MA: MIT Press.

Lave, J. (1988). *Cognition in practice: Mind, mathematics, and culture in everyday life*. New York: Cambridge University Press.

Lincoln, Y., & Guba, E. (1985). *Naturalistic inquiry*. Newberry Park, CA: Sage.

Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Adams-Bannister-Cox.

Nersessian, N. J. (1979). The roots of epistemological ‘anarchy’. *Inquiry*, *22*, 423–440.

Nersessian, N. J. (1984a). Aether/or: The creation of scientific concepts. *Studies in the History & Philosophy of Science*, *15*, 175–212.

Nersessian, N. J. (1984b). *Faraday to Einstein: Constructing meaning in scientific theories*. Dordrecht: Martinus Nijhoff/Kluwer Academic Publishers.

Nersessian, N. J. (1987). A cognitive-historical approach to meaning in scientific theories. In N. J. Nersessian (Ed.), *The process of science* (pp. 161–179). Dordrecht: Kluwer Academic Publishers.

Nersessian, N. J. (1988). Reasoning from imagery and analogy in scientific concept formation. In A. Fine & J. Leplin (Eds.), *PSA 1988* (Vol. 1, pp. 41–47). East Lansing, MI: Philosophy of Science Association.

Nersessian, N. J. (1991). Discussion: The method to ‘meaning’: A reply to Leplin. *Philosophy of Science*, *58*, 678–687.

Nersessian, N. J. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In R. Giere (Ed.), *Minnesota studies in the philosophy of science* (pp. 3–45). Minneapolis, MN: University of Minnesota Press.

Nersessian, N. J. (1995). How do scientists think? Capturing the dynamics of conceptual change in science. In B. Chandrasekaran, J. Glashow, & H. Narayanan (Eds.), *Diagrammatic reasoning: Computational and cognitive perspectives* (pp. 137–182). Cambridge, MA: AAAI/MIT Press.

Nersessian, N. J. (1999). Model-based Reasoning in Conceptual Change. In L. Magnani, N. J. Nersessian, & P. Thagard (Eds.), *Model-Based Reasoning in Scientific Discovery* (pp. 5–22). New York: Kluwer Academic/Plenum Publishers.

Nersessian, N. J. (2002a). The cognitive basis of model-based reasoning in science. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The cognitive basis of science* (pp. 133–153). Cambridge: Cambridge University Press.

Nersessian, N. J. (2002b). Maxwell and the “method of physical analogy”: Model-based reasoning, generic abstraction, and conceptual change. In D. Malament (Ed.), *Reading natural philosophy: Essays in the history and philosophy of science and mathematics* (pp. 129–165). Lasalle, IL: Open Court.

Nersessian, N. J. (2005). Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In M. Gorman, R. D. Tweney, D. Gooding, & A. Kincannon (Eds.), *Scientific and technological thinking* (pp. 17–56). Hillsdale, NJ: Lawrence Erlbaum.

Nersessian, N. J. (2008a). *Creating scientific concepts*. Cambridge, MA: MIT Press.

Nersessian, N. J. (2008b). Mental modeling in conceptual change. In S. Vosniadou (Ed.), *International handbook of conceptual change (forthcoming)* (pp. 391–416). London: Routledge.

Nersessian, N. J. (2022). *Interdisciplinarity in the making: Models and methods in frontier science*. Cambridge, MA: MIT Press.

Nersessian, N. J., & Chandrasekharan, S. (2009). Hybrid analogies in conceptual innovation in science. *Cognitive Systems Research*, *10*, 178–188.

Nersessian, N. J., Kurz-Milcke, E., Newstetter, W., & Davies, J. (2003). Research laboratories as evolving distributed cognitive systems. In D. Alterman & D. Kirsch (Eds.), *Proceedings of the Cognitive Science Society* (Vol. 25, pp. 857–862). Hillsdale, NJ: Lawrence Erlbaum Associates.

Nersessian, N. J., & MacLeod, M. (2022). Rethinking ethnography for philosophy of science. *Philosophy of Science*, *89*(4), 721–741.

Newton, I. (1687). *Mathematical principles of natural philosophy and his system of the world (A. Motte, Trans.)*. Berkeley, CA: University of California Press.

O'Doherty, K., Osbeck, L., Schraube, E., & Yen, J. (2019). *Psychological studies of science and technology*. New York: Palgrave MacMillon.

Osbeck, L., & Nersessian, N. J. (2006). The distribution of representation. *Journal for the Theory of Social Behaviour*, *36*, 141–160.

Osbeck, L. M., & Nersessian, N. J. (2017). Epistemic identities in interdisciplinary science. *Perspectives on Science*, *25*, 226–260.

Osbeck, L. M., Nersessian, N. J., Malone, K. R., & Newstetter, W. C. (2011). *Science as psychology: Sense-making and identity in science practice*. Cambridge: Cambridge University Press.

Popper, K. R. (1962). *Conjectures and refutations*. New York: Basic Books.

Quine, W. V. O. (1969). Naturalized epistemology. In *Ontological relativity and other essays* (pp. 69–91). Cambridge, MA: Harvard University Press.

Schwartz, D. L., & Black, J. B. (1996). Analog imagery in mental model reasoning: Depictive models. *Cognitive Psychology*, *30*, 154–219.

Thagard, P. (1992). *Conceptual revolutions*. Princeton, NJ: Princeton University Press.

Tweney, R. D. (1985). Faraday's discovery of induction: A cognitive approach. In D. Gooding, & F. A. J. L. James (Eds.), *Faraday rediscovered* (pp. 189–210). New York: Stockton Press.

Wertheimer, M. (1959). *Productive thinking*. Oxford, England: Harper.

Yagmur-Kilimci, E. (2010). *3D mental visualization in architectural design*. Georgia Institute of Technology.