1 Investigating Practice: The Cognitive-Cultural Systems of the Research Lab

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? I have been occupied with this problem in various formulations starting from when I was an aspiring theoretical physicist and then as a philosopher of science but also, in addition, as a cognitive scientist. Science is an activity with many dimensions, and for the last twenty years I have been following out my conviction that the answer lies in fathoming how these are integrated in the problem-solving practices scientists create to get a grip on the world.

By "many dimensions" I mean that there is, obviously, a *cognitive* side to science (there's no way to do science without using your mind/brain), but, equally obviously, also a *social* side (lab organizations, academic institutions, how people work together, and so forth), a *material* side (for example, computers, pipettes, instruments, cells, and chemicals), and a *cultural* side (for instance, locally maintained traditions). In each of these dimensions, scientists create resources through which to think. I mark these jointly as "cognitive-cultural" to indicate that in scientific practice these dimensions are integrated.

Scientists' problem-solving practices comprise a range of activities that are of great interest to philosophers as well as to cognitive scientists. Such activities include, among others, the manner in which scientists reason, how they make representations of phenomena, how they construe "understanding," how they use their imagination, and how they work together on a day-to-day basis. The problem-solving practices of scientists on the frontier are exploratory; they are incremental in the sense of moving step-by-step;

and they are nonlinear in that they do not follow any obvious, let alone preexistent, pathway from original problem to final resolution. It is in the processes of developing and using these practices that scientists create complex dynamical investigative systems of research in which cognition and culture are mutually implicated in epistemic practices. Thus, scientific discovery and creativity need to be understood as system phenomena.

The route I have taken toward understanding what is going on in this vast multidimensional realm is to develop fine-grained examinations of those problem-solving practices, their origins and advances over time, and the epistemic principles that guide them. As a professor in philosophy of science *and* cognitive science at Georgia Tech, I was presented with an opportunity to learn what was going on in research undertaken there in laboratories working at the exciting intersection of biology and engineering. It seemed to me from even a cursory initial look that the bioengineering sciences, as frontier research areas, provide an excellent locus for examining cognitive-cultural integration. They are paradigmatic of the interdisciplinary aspirations of frontier twenty-first-century science. Researchers in these fields aim to bring engineering, technological, and computational resources to bear on understanding and controlling complex biological systems. To achieve their ambitions, they would need to "integrate" cognitive-cultural resources in the form of concepts, methods, and materials from the domains of biology and engineering, which requires understanding their affordances and limitations, and to train new kinds of researchers.

I had no prior understanding of the domain, and the selection of labs to investigate was largely a matter of chance. Each specific field was undertaking research on complex biological systems about which there was little scientific understanding. In each lab, the members self-consciously referred to themselves as pioneering researchers on the frontiers of engineering and science. They expressed excitement at being on the forefront of research. As one told us, "*What we are doing right now is the most exciting thing I can think of, and I think most people know that too when they hear about it.*" We began with a lab where engineering techniques are applied to cells and tissues that constitute the blood vessel wall, simply because the director approached me in my role as director of the Program in Cognitive Science. He, along with other senior colleagues, wanted to explore what cognitive scientists might contribute to the development of a new educational vision they had: to train the emerging field's researchers as truly integrative, hybrid "bio-medical-engineers." I immediately saw in this offer the opportunities I had been seeking to both extend my research from historical studies in physics to ethnographic investigations, now, on frontier research in the bioengineering sciences, and to extend my advocacy that scientists should "preach what they practice" (Nersessian 1995) to an innovative educational program as it was being developed. However eager I was to join their endeavor, I responded that I needed first of all to develop some understanding of their research problems and their practices in attempting to solve them. I put together a research group that was (in gradually changing compositions) to investigate their research and learning practices for more than fifteen years.[1] We began to investigate the tissue engineering lab, and three subsequent labs, by conducting a pilot study. We interviewed the lab director about his background and the lab's history and current research, toured the lab while the director and the lab manager described artifacts and activities, and met with the group of researchers who would participate in the investigation. This preliminary research enabled us to focus our research project better.

That is how it began. This book is where it ends. I have two main goals with it. One goal is to establish how the modeling practices of bioengineering research labs exemplify, and work within, the cognitive-cultural framework.[2] I cast research labs and the problem-solving activities within them as what I call *distributed cognitive-cultural systems with epistemic aims*—that is, roughly, a complex system comprising researchers, artifacts, social structures, and practices through which specific epistemic aims are advanced. Within that special frame, which I discuss at some length in section 1.2. below, I examine the building of models (that is, of material and computational artifacts) as what drives the creation and evolution of these systems. On my analysis, models are loci of cognitive-cultural integration. The processes of model-building make a researcher a part of "the lab" socially, culturally, and cognitively.

The other goal is to build on my previous work to further articulate, and to develop the rationale for, the integrative cognitive-cultural framework itself. This framework is applicable to scientific practice, in a broader, more general sense, as well. That is, although my exploration is of the bioengineering sciences that provide my case material, the analytic framework, along with my methodological approach, applies to scientific research as such.

In this introduction, I first lay out the reasoning underlying the integrative framework, and also the methodological approach my research group and I have followed—what I call “cognitive ethnography of research labs.” Next, I describe the four labs we have investigated over many years, and how we adapted and used ethnographic methods for our specific purposes. Finally, I sum up the themes of the successive chapters of the book.

1.1 A Cognitive-Cultural Framework for Investigating Practice

Thomas Kuhn’s *The Structure of Scientific Revolutions* (Kuhn 1962) was a major impetus in the flourishing of case studies of the investigative practices of scientists since the 1970s. From the start of the “practice turn,” accounts of scientific practices have tended to focus on *either* “cognitive/rational” *or* “cultural/social” factors, where these are taken to be separate—often mutually exclusive—interpretive categories. There have been, on the one hand, philosophical and cognitive studies of science, both of which focus on individuals and on the cognitive representations and processes they use in problem-solving and, for philosophy, on the epistemic standards they employ in carrying out experimental or theoretical practices.[3] But there also have been, on the other hand, social and anthropological studies of practice (STS), which focus instead on interests, motivations, and a range of social, cultural, and material factors in play within communities of scientists. Both kinds of analyses have certainly produced valuable insights about scientific practice. Even so, their oppositional stances have prevented fruitful interaction. STS accounts have programmatically constrained relevant explanatory factors to social and cultural, thus downplaying (or denying entirely) the relevance of rational and cognitive factors. On the cognitive side there have been two main objections. Philosophers reject STS accounts for their epistemic relativism. In addition, both philosophers and cognitive scientists are opposed to the social reductionism, likewise characteristic of STS studies, that completely omits cognitive and rational factors from explanations of how science produces knowledge. Such reductionism was most famously announced in the “ten-year moratorium” on cognitive explanations issued first in 1986 by Bruno Latour and Steven Woolgar (Latour and Woolgar 1979, 280; Latour 1987, 247).

The moratorium is long over, but one still needs to ask: What, besides a penchant for rhetorical flourish, could explain such a pronouncement as the ten-year moratorium? One can agree that scientists have interests, motivations, and sociocultural loci in conducting research, but they also have human cognitive capabilities and epistemic aims. Scientists develop and use artifacts ("material culture"), but they also articulate and use evaluative procedures and standards. These "both/ands" look like such truisms that only the investigation of a mix of complex issues can lay bare the roots of the moratorium in twentieth-century intellectual history. I have attempted such an examination elsewhere (Nersessian 2005), but there is no need to detail it here. For our purposes, it suffices to agree that any such divide, though at times analytically useful, is artificial. It imposes boundaries on dimensions that are inherently integrated *in practice*.[4] What twenty-first-century studies of science require is *an analytical framework that facilitates theorizing the cognitive and cultural aspects of practice in relation to one another*.

Here is how I articulate such a cognitive-cultural framework for investigating practice. As with my previous methodological approach of cognitive-historical analysis (Nersessian 1987, 1992a, 2008), I draw, primarily, on resources in *both* the philosophy of science *and* the cognitive sciences. From a philosophical perspective, the project of integration fits into a philosophical tradition of epistemological naturalism (Quine 1969; Nersessian 2008). My construal of that naturalism has the following requirements. First of all, the project needs to be informed by the best available scientific understanding of humans that the biological, cognitive, and social sciences have to offer. Next, it needs to be informed by a grasp of the actual investigative practices as they are created and used by scientists, including their epistemic warrant. Finally, we need to avail ourselves of appropriate empirical methods for determining these practices.

Already in my earlier work on the cognitive basis of model-based reasoning I have established the fruitfulness of thinking about scientific practices through the lens of specific research in the cognitive sciences (Nersessian 1984, 1987, 1992a, 2002a,b, 2008). And, indeed, for the case at hand, there is a line of research that can provide important insights of an empirical, but also theoretical and even methodological, nature—insights that are most helpful if we want to attain an integrative study of scientific problem-solving practices as located with complex cognitive-cultural systems.

Within contemporary cognitive science, there is a movement across the disciplines it comprises that has advanced and persuasively argued for an integrated understanding of cognition and culture. This understanding did not begin with science as a specific object of investigation—that is where much of my own work over past decades, and also the present book, comes in. Even so, this specific direction in cognitive science has much to offer by way of a starting point from which to tackle the problem of integration in science. Indeed, participants in the movement have, themselves, advocated that cognitive science study real-world practices because, "it is in real practice that culture is produced and reproduced. In practice we see the connection between history and the future and between cultural structure and social structure" (Hutchins 1995a, xiv).

The argument in cognitive science begins from a critique of traditional research that locates and studies cognition in terms of what goes on "in the head," without taking into account the role of environmental resources (social, cultural, and material) in shaping and participating in cognitive processes. What I have called "environmental perspectives" (Nersessian 2005) are grounded in empirical evidence from a range of research in the cognitive, biological, and social sciences. Environmental perspectives comprise accounts of cognition as an embodied, artifact-using, and situated process that the social, cultural, and material environment does not just scaffold, but also provides resources that are integral to cognitive process.[5] Advocates for environmental perspectives contend that analyses of cognitive processes need to determine and incorporate salient dimensions of the contexts and activities in which cognition occurs.

Traditionally, the environment is construed as providing mental content on which "internal" cognitive processes operate. In contrast, pioneers of the new approach have argued that "cognition" and "culture" both need to be seen as processes that are integrated in human activity (Hutchins 1995a; Shore 1997). This construal points away from the traditional perspective that regards cognitive and cultural *factors* as independent variables in an explanation of intelligent behavior. Instead, human cognition is understood to be inherently cultural, as being shaped in ongoing processes of the evolutionary history of the human species, of the development of the human child, and of the various sociocultural environments in which learning and work take place—*processes* that are integral to intelligent behavior.

The route to attaining analytical integration for scientific practice follows a path similar to environmental perspectives in that it, too, moves the boundaries of scientific problem-solving beyond the individual to comprise complex cognitive-cultural systems. The grain size of the system—individual, group, groups of groups—depends on the focus of analysis. Here the analytical framework of "distributed cognition" proves particularly useful, in that it incorporates all of the dimensions of an integrated analysis advanced by the range of environmental perspectives. From that perspective, scientific researchers are understood to be embodied agents situated in problem-solving contexts. That is why I have adopted this analytical framework as the starting point of a cognitive-cultural analysis of scientific practice. In this project, we analyze scientific problem-solving as situated within contexts and distributed across researchers and specific artifacts. Extremely fruitful as it is for starters, the framework of distributed cognition is itself, however, in need of further development. What, for present purposes, it needs in particular is to think about what adjustments and enrichments are required to accommodate scientific practice.

1.2 Broadening the Framework of Distributed Cognition

The framework of distributed cognition (from here on, D-cog) has its roots in concerns that arose, starting in the mid-1980s, from the sense that traditional cognitive science had put too much of what goes on in cognition "in the head." Clearly humans are able to do some simple reasoning and problem-solving in the absence of external resources, but we are able to accomplish much more by drawing on those resources—either using or creating them as needed. As Daniel Dennett has remarked, "Just as you cannot do very much carpentry with your bare hands, there's not much thinking you can do with your bare mind" (Dennett 2000, 17). Thinking requires doing with resources in the mind and in the world. One question is how the mind incorporates external resources in thinking. Several streams of research, much of it using ethnographic methods, converge on the notion of cognitive processes as distributed across people and artifacts, including distributed cognition, distributed intelligence, activity theory, situated action, and extended mind theory.[6] Of these various approaches, Edwin Hutchins's research has been most influential on my project of integration because he has focused on the roles of representational artifacts in problem-solving processes ("tasks") situated in technologically rich environments ("socio-technical systems"). Although not themselves science, these systems have features that do align well with science. Further, I agree with Hutchins in that I understand D-cog to be a framework for analyzing cognitive-cultural processes, not an ontological thesis.[7]

The primary unit of analysis is what Hutchins calls a distributed socio-technical system. It consists of people working either together or individually to accomplish a task, along with the specific artifacts they use in the process. *Cognitive artifacts* are the artifacts that have been designed to perform cognitive functions and are used in problem-solving tasks by groups of people (e.g., piloting a ship, Hutchins 1995a) or individuals (e.g., piloting a plane, Hutchins 1995b). The classic exemplar is that of the speed bug on the airplane cockpit instrumentation panel. It eliminates the need for the pilots to remember the range of minimum maneuvering speeds while taking off and landing a plane and, when there are copilots, enables them to have a coordinated perspective (Hutchins 1995b). What makes the speed bug artifact "cognitive" is that it provides a representation that "remembers" the speed constraints for the pilots, thus reducing their cognitive load at a critical juncture in flying a plane. On Hutchins's account, cognitive artifacts, generally, are material media that possess the cognitive properties of generating or manipulating or propagating representations. D-cog analyses, then, focus on the functions of these representational artifacts in human activities.

Science, however, differs in significant ways from the socio-technical systems Hutchins has investigated. Therefore, to accommodate it to our purposes, the D-cog framework needs to be broadened.[8] The framework of distributed cognition has been developed largely from investigations of contexts quite different from research laboratories and of practices quite different from those of scientific problem-solving. In the sections that follow I elaborate on these differences.

First, much of the research that has contributed to the initial framework of D-cog has focused on highly defined task environments with well-defined problems and goals and ready-to-hand cognitive artifacts. Scientific environments, in contrast, are continually evolving along all dimensions.

Second, initial D-cog research has focused on the way artifacts/technologies, as "external" representations, contribute to accomplishing cognitive tasks, especially how they reduce internal cognitive load by "off-loading" cognitive processes, such as memory, to the environment. Scant attention has been paid to the nature of the resources the mind/brain, especially "internal" (human memory) representations and processes, contribute to problem-solving tasks.[9] If we wish to understand scientific problem-solving, we need to pay attention to how these two kinds of representations—artifact and memory—interact.

Third, unlike other practices that have been investigated in a D-cog framework, science is an epistemic practice that needs to justify its outcomes. A scientist cannot just say, "I have discovered this or that." She is bound to offer arguments grounded in the observations made, in the theories employed, and in the methods followed, for why she regards her discovery claims as justified. In D-cog analyses so far, attention to this dimension has, by and large, been lacking.[10]

In these three distinct (though naturally overlapping) regards, then, the D-cog framework needs to be extended: science as a process of ongoing growth and change; science as marked by ongoing interaction between memory and artifactual representations; science as standing in ongoing need of an epistemic warrant. I now successively discuss in more detail these three dimensions to be incorporated into the D-cog framework, always with a view to making it the best possible analytical tool for our investigation of present-day science, in general, and, in particular, bioengineering epistemic practices.

1.2.1 Scientific Task Environments

The distributed cognition framework was developed to analyze cognitive processes in complex problem-solving environments in which there are well-defined problems and tasks for which the requisite salient artifacts and technologies for problem-solving are already at hand. In these studies of environments, for instance the cockpit (Hutchins 1995b) or a naval ship (Hutchins 1995a), the problem-solving situations change in time, but, for instance, how to land a plane is a well-defined problem. Situational features of the problem faced by the pilot or crew can change in the process of landing the plane or bringing a ship into the harbor, and creative solutions can arise in novel situations. However, the technology (cognitive artifacts) required to accomplish the problem-solving, such as the alidade or the speed bug, the practices surrounding them, and the knowledge the pilot and crew bring to bear in those processes, are relatively stable. Even though the artifacts have a history in the field, as Hutchins documents for the instruments of ship pilotage, they do not change in the day-to-day problem-solving processes on board. Thus, this kind of cognitive system is dynamic but largely *synchronic*. By contrast, the cognitive systems of scientific problem-solving are not only dynamic but also *diachronic*.

Scientific problem-solving environments have open, ill-formed problems and goals, and tasks for which methods and artifacts for problem-solving often need to be created, depending on the science and its state of development. Although there are loci of stability, the salient artifacts, methods, knowledge, and the problem formulations themselves can all undergo development and change during problem-solving processes. This is especially so in frontier science environments. A bioengineering research lab provides a good example. Not only does this environment have open, ill-defined problems and goals, but novel methods are under development, and the most salient cognitive artifacts—in this case, physical and computational simulation models—need to be created. Indeed, one could say the lab itself needs to be built in the process of articulating its problems and goals (Nersessian 2012).

Further, university labs are sites of situated learning within a community of practice (Lave and Wenger 1991). In the labs we investigated, the researchers are primarily graduate students, who are "researcher-learners" with developmental trajectories that intersect with the trajectories of development of the technological resources and other dimensions of the problem space. Of course, problem-solving on the frontier requires everyone—including lab directors and managers—to be learning continuously, but student researchers have the additional challenge of acquiring the knowledge and skills to become scientists. A further challenge for everyone in interdisciplinary communities is that interactions among concepts, knowledge, practices, norms, and so forth from more than one field need to be formed into new cognitive-cultural practices. The dynamic and diachronic nature of research labs has led us to characterize them as *evolving distributed cognitive-cultural systems* (from here on to be abbreviated as cognitive-cultural systems or, simply, D-cog systems). Research labs, as problem-solving environments, provide the framework of D-cog with the opportunity to attain what Hutchins has called its "most ambitious goal": to investigate "a system in which adaptive processes that are continually operating are responsible for the production of both stability and change" (1996, 67). Lab researchers build their cognitive-cultural systems as the research moves along, especially through building the artifacts necessary for the activities he singles out: generating, manipulating, and propagating representations within the system.

Much of the work within D-cog has focused on constructing detailed descriptions of the way artifactual representations are used and how they change the nature of the cognitive tasks, especially by reducing cognitive load through "off-loading" memory to the environment. Less understood are the processes of generating/building artifacts to alter task environments in the course of problem-solving.[11] Indeed, as Daniel Schwartz and Thomas Martin have observed, “most cognitive research has been silent about the signature capacity of humans for altering the structure of their social and physical environment” (Schwartz and Martin 2006, 314). Rogers Hall et al. have also noted the lack of attention paid to the practices through which people actively *distribute* cognitive processes to the environment (Hall et al. 2002; Hall et al. 2010). Yet, a central premise of D-cog is precisely that. As Hutchins has succinctly stated, “Humans create their cognitive powers by creating the environments in which they exercise those powers” (1995a, 169). Scientific practices provide an especially good locus for examining the human capability to create or extend cognitive powers. After all, distributing cognition through creating problem-solving environments is a major component of scientific research. The problem-solving environments scientists create include material and conceptual artifacts, methodological practices, and communities of researchers (whether working together or alone), such as labs. By examining the processes through which these environments are built, we can begin to understand how the artifactual resources, in particular, are *incorporated into* problem-solving systems, and are not just external representations “used by” the mind/brain. In our analyses of the interactions between the mental and the artifactual components of the D-cog system in model-based reasoning processes, we cast the relationship as one of “coupling,” rather than “off-loading” and then “using.” This distinction will become clearer in the chapters of this book, but for present purposes, in general, we use the notion of coupling to indicate that in problem-solving processes, each component of the D-cog system—mental model and material model—interacts in a manner that can create change in the other. But to extend the D-cog framework in this way requires we attend to the nature of the resources the mind/brain contributes to problem-solving. That is our second missing dimension, to be taken up now.

1.2.2 Mental Modeling

To date in D-cog research there has been scant discussion on the nature of the mental resources the human component of the system contributes beyond attending to how the people involved perform coordinating functions. There are, however, many cognitive capacities that come into play in problem-solving. Important examples are memory, representation, reasoning, imagination, and executive functions. Hutchins is explicit that his is a representational view of cognition, in which intelligent behavior results from interactions among "internal" (human) and "external" (artifact) representations in cognitive processes. His analyses track what he calls the generation, manipulation, and propagation of representations in the D-cog system. However, Hutchins focuses primarily on the nature and function of the *artifact* representations, while remaining silent about the *human* representations.

Jaijie Zhang and Donald Norman (Zhang and Norman 1995) provide a rare analysis from a D-cog perspective of the interactions among external and internal representations, in this case, solving insight problems (e.g., tower of Hanoi) and reasoning in games (e.g., tic-tac-toe). They assume, on the basis of a substantial experimental literature, that the internal representations are mental models; however, they do not elaborate on their nature. In earlier work I have developed an account of the capacity for mental modeling, as it provides a cognitive basis for model-based reasoning practices in science (Nersessian 1992a,b, 2002, 2008). My pathway to constructing that account was to synthesize a wide range of cognitive science research from what are usually viewed as separate areas, but which I came to see as interrelated from the perspective of scientific thinking, and to add data from investigations of scientific model-based reasoning practices. I have argued that the cognitive basis for these practices lies in the mental ability to imagine real-world, counterfactual, and impossible situations, and to make inferences about future states of these situations through manipulation of a model of the situation. The remainder of this section provides the basic features of that account, by way of a brief summary of the full analysis to be found in chapters 4 and 5 of Nersessian (2008).

Several important lines of theoretical and experimental research concerning mental models started with the 1967 reissue of the book on explanation by Kenneth Craik (Craik 1943). One specific line focuses on working-memory processes of dynamic and mechanistic mental modeling. My analysis draws on this strand of research in the literature that, like Craik's, examines the processes of constructing and manipulating a mental model during reasoning and problem-solving. Inferences, on this account, are made by means of manipulating both static and dynamic features of the model. That literature, including my own contribution, addresses working memory representations and does not make any claims about the nature of long-term memory representations, which some have claimed to be mental models as well. I have argued that to accommodate the complex nature of scientific reasoning requires an account of model-based inference to move beyond the mental modeling literature per se, so as to construct a synthesis of an extensive range of experimental literature that consists of research on a whole range of issues that implicates mental modeling. These issues comprise discourse and situation modeling, mental animation, mental spatial simulation, and, finally, perceptual simulation in embodied mental representation.[12] When taken together, this research supports a "minimalist hypothesis" that "in certain problem-solving situations humans reason by constructing a mental model . . . in working memory that in dynamic cases can be manipulated by simulation" (Nersessian 2002, 143). That is, people have the capacity to perform what I call simulative model-based reasoning. Such a mental model "is an organized unit of knowledge that embodies representations of spatiotemporal relations of situations, entities, and processes, as well as representations of other pertinent information, such as causal structure" (Nersessian 2008, 128).

Mental modeling is subject to the representational and processing constraints/capacities of the brain, chief among which is memory. Thus, mental model representations are limited in how much detail they can contain, and for how long. The experimental research associated with mental model simulations of physical and mechanistic representations (for instance, pulley systems) identifies specific features of these as being qualitative (DeKleer and Brown 1983; Roschelle and Greeno 1987) and with animation/simulation being piecemeal (Roschelle and Greeno 1987; Hegarty 1992, 2004; Schwartz and Black 1996). Long-term memory provides background knowledge to support mental model-building and simulation processes (Roschelle and Greeno 1987). Information in various formats can be used to construct, manipulate, and revise a model. Such formats may cover language but also formulae, pictures, sounds, and kinesthetic phenomena.

In line with Dennett's comment, simple mental simulations might be possible "in the head" alone, yet mental modeling of any complexity needs to be carried out in the presence of real-world resources. For example, consider a mundane case. It is much easier to mentally simulate how to get an awkward piece of furniture through the door when it is in front of the reasoner in the doorway, rather than recalled from a visit to the furniture store. This is even more true in the case of science. A wide range of data, which comes from historical records, from protocol studies, and from ethnographic investigations, establish that many kinds of external representations are integral to reasoning in scientific problem-solving processes. Such representations can be linguistic (descriptions, narratives, written and oral communications), or mathematical (equations), or visual (diagrams, graphs, sketches, computer). They also include gestures, physical models, and computational models.

How might the capacity for mental modeling interface with the relevant resources in the external world? Much of the experimental research on this interface has been directed toward the use of diagrams and other visual representations. The research noted above by Zhang and Norman, for instance, analyzes diagrams as external representations that are coupled as information sources with mental models in problem-solving. Mary Hegarty (Hegarty 2004) has argued that the corpus of research on mental animation in the context of external visual representations leads to the conclusion that these and internal representations form a coupled system in inferential processing (see also, Gorman 1997; Greeno 1989b). Although limited, the experimental research on scientific reasoning, in particular with computational representations, promotes the coupled system view as well (Christensen and Schunn, 2008; Trafton et al. 2005; Trickett and Trafton 2007). Our research extends the notion of mental model and external representational coupling to include physical and computational models in processes of simulative model-based reasoning (Osbeck and Nersessian 2006; Nersessian 2008, 2009; Chandrasekharan and Nersessian 2015; Nersessian et al. 2003). Diagrams and other visual representations (with the exception of computational) considered in the experimental literature are static, but our analyses concern also the interface between dynamic representations: physical model simulations, computational model simulations, and mental model simulations. Unlike most cases examined in the experimental literature, our research looks at cases in which each part of the system is subject to change in inferential processes—mental models can improve from interaction with material models and vice versa. We understand "coupling" on analogy with how it is understood, generally, in mechanics as heterogeneous components interacting dynamically in a feedback loop to improve the function of a system, in our case, model-based reasoning in a D-cog system.

One way to accommodate this notion of coupling would be to expand what is understood as memory to encompass external representations and cues. If memory is so distributed, problem-solving affordances and constraints in the environment are ab initio part of cognitive processes, which now incorporate both kinds of representations (see, e.g., Donald 1991). In this case, during model-building processes, mental and real-world models iteratively develop correspondences among - their features, and in reasoning, information is co-processed in human memory and in the environment. Although cognitive science stands in need of an account of the nature of the mechanisms of internal/external representational coupling (see, e.g., Chandrasekharan and Stewart 2007; Rahaman et al. 2018), we contend that coupling provides a better metaphor than off-loading (itself in need of such an account of mechanisms), since coupling intimates that cognitive artifacts become incorporated into a D-cog system. We would not want to say that reasoning processes are off-loaded to an artifact model as memory processes are to the speed bug; rather the model and modeler form a dynamic, coupled system that performs reasoning, and, over time, and in interaction with other elements in the system, changes in one lead to changes in the other. Indeed, we have been arguing and I will show in this book, building coupling between mental and artifact models in a D-cog system is a major means through which scientists "create cognitive powers," which extend beyond enhanced memory. To paraphrase Hutchins, *one way in which scientists create their cognitive powers is by creating modeling environments*.

1.2.3 Scientific Practice Is Epistemic

Finally, unlike most other practices that have been investigated within the D-cog framework, science is an epistemic practice. Research labs, and other configurations of scientific practice, can be cast as what Karin Knorr Cetina has called epistemic cultures (1999). These "are cultures that create and warrant knowledge" (Cetina 1999, 1). She chose that designation to contrast with the customary terms of "discipline" and "specialty," which typically refer, in the social sciences, to institutional organizations of knowledge. "Epistemic culture" is used to shift the focus of attention to "knowledge-in-action" (Cetina 1999, 3), or practice. As an approach to the study of science, according to Cetina, to analyze an epistemic culture requires one attend to the differences of "knowledge-making machineries" in different scientific cultures and subcultures (Cetina 1999, 3). Her case studies of practices in particle physics and molecular biology make clear that "machineries" comprise sociocultural practices as well as the technologies of research. Her analysis of these epistemic practices, however, is too limited. It largely omits the cognitive and rational dimensions of the practices.[13] In particular, she does not attend to how a culture provides, or develops, warrant for its practices, including its "machineries." What makes science epistemic, however, is not only that it makes claims to create knowledge, but also that it provides warrant for those claims and the investigative methods leading to them. Nor does she attend to differences in epistemological assumptions, norms, and values among cultures, which Evelyn Fox Keller has pointed out are equally significant for individuating scientific cultures and understanding their practices (Keller 2002). My understanding of an "epistemic culture," and use of that designation, include all of these dimensions.

In the context of extending the framework of D-cog to accommodate science, it is important underscore the normative dimension of a practice. Joseph Rouse's general analysis of practice is useful to understand what this means. Rouse characterizes practices as "situated patterns of activity" (Rouse 1996, 150). Such a pattern "constitutes a practice rather than some other kind of regularity to the extent that it is a pattern of correct or appropriate performance" (Rouse 1996, 137). So, the idiosyncratic activities of a pilot or a scientist do not constitute a practice. Rather, a practice has correct and incorrect performance standards, and is evaluated in light of these. Unlike other practices investigated by D-cog, however, science makes epistemic claims on the basis of its practices. Just as a claim to scientific discovery is not a "discovery" until it has been accepted as justified by something approaching consensus in the relevant field (see, e.g., Arabatzis 2006), a scientific practice is not a "practice" until it is acknowledged in the field as a warranted means of investigation. As Helen Longino has pointed out in her analyses aimed at bridging what she calls the "rational-social" divide, the normative dimension of practice is inherently social, in that it relies on the critical scrutiny of the relevant scientific community (Longino 1990, 2001).

The decision to warrant a scientific practice or system of practices is based on an assessment of the reasons advanced in favor of the practice or the system, and on a history of reliable success at achieving its epistemic aims. In developing novel practices, researchers are required to advance reasons for why their specific application is warranted, and also for their more general applicability in the field. For instance, the computational systems biologists we studied provided reasons in support of the abstractions they used to represent biological interactions while they were building specific models. They came up, likewise, with arguments to account for why their practice of building certain midlevel ("mesoscopic") models with less fidelity to the biological details than the field customarily aims at is warranted. They argued, and demonstrated, that such models can produce significant understanding and possibilities for intervention (e.g., to create biofuels or kill cancer cells), while also providing insight into how one might build in fidelity as capabilities for modeling develop in the field.

In our investigations of problem-solving in research labs, we consider "normativity," in its most general sense, to mean that there are specific constraints on practice. These constraints are of three kinds. Some are material, in that they are tied to the composition and behavior of entities and objects. Others, of a cognitive nature, are imposed by the processes and structures by which humans, for instance, categorize and make inferences. Again, others are imposed socially, in accordance with standards of research and of professional conduct. All these constraints can be shared by other kinds of practices, but science alone is subject, in addition, to constraints that are imposed by the epistemic aims and justification of the practice. That is why our examinations of innovative methodological practices for modeling attend to how warrant is built for novel practices in the course of their development, as well as for specific models, which are intended to function epistemically.

We, of course, realize that the wider community is always implicated in the development a practice. Even so, for practical reasons we limited the scope of our investigations mainly to the participants in our labs, including external collaborators, where present and feasible. We did nonetheless attend to how the researchers expressed the norms and values of their communities, and also to how their work was received and critiqued in community responses, notably in the context of conference presentations and of reviewer responses to publication submissions and grant proposals (see, e.g., Osbeck et al. 2011, chapters 5 and 6).

1.3 Cognitive Ethnography of Research Labs

To determine how warrant is built for practices, and indeed to move beyond the perceived cognitive-cultural divide more broadly, requires more than theoretical argumentation of the kind broached above. It requires likewise fine-grained empirical investigations of the epistemic practices of scientists to determine how integration takes place in these practices. Until quite recently (and definitely when we began our investigations), philosophers out to determine practices have been relying primarily on historical data of two kinds—archival records and publications. Historical records, however, leave us at the mercy of "what's left behind"—a problem even more significant with late twentieth and early twenty-first-century science—and rarely are they sufficient to provide details on the day-to-day research problems and paths to solution implicated in what become discoveries. Established canons of modern science also require that such details be omitted from published accounts of research. Impasses and obstacles encountered along the way, in particular, rarely make an appearance in published accounts, which frame the problem-solving processes and reasoning as linear and relatively straightforward (Bazerman 1988). Historical data do afford a means to contextualize the practices of scientists with respect to their historical situatedness from the viewpoint of the problem situations of the traditions in which scientists carried out their work, and so can advance the project of integration (Nersessian 2008). But full analysis of integration requires more. It necessitates that we move beyond exclusive use of historical records, which by their nature place limitations on thinking about the interplay of cognition and culture, to carrying out empirical investigations of practices *as they are enacted in situ* in order to determine the range and kind of resources that contribute to the epistemic accomplishments of science. In other words, we need to move beyond historically informed philosophy to ethnographically informed philosophy.

What understanding philosophers of science have of ethnography derives, primarily, from how it has been portrayed and used in STS. STS researchers have established that ethnography, which comprises field observations and interviews, provides a fruitful means of collecting data on day-to-day scientific practices in research labs.[14] The main objection of philosophers to STS accounts has been, and still is, their tendency toward epistemic relativism, which stems from their alignment with the so-called strong program of the "sociology of scientific knowledge (SSK)" (see, e.g., Bloor 1991). Since their inception and in line with the "ten-year moratorium" mentioned above, these accounts have programmatically constrained relevant explanatory factors in their analyses to social and cultural factors, including personal motivations and interests, while downplaying or denying entirely the relevance of rational and cognitive factors.[15] There is nothing in principle, however, about ethnographic methods as such that would prohibit addressing those aspects of practice. Indeed, a major contribution of STS ethnographies of research labs is that they have demonstrated repeatedly the value of ethnographic methods for investigating scientific practices across a range of sciences. I began the project presented here with the conviction that philosophers need not cede this important methodological tool to sociocultural science studies fields, but that it could be adapted and used to address philosophical issues (Nersessian and Macleod 2022). This conviction was supported by its use by cognitive scientists to investigate other kinds of problem-solving practices in D-cog. Ethnography conducted for philosophical objectives can be placed within the perspective of what is called “cognitive ethnography” within the D-cog framework.

1.3.1 Cognitive Ethnography of Scientific Practices as a Method

D-cog has been using ethnographic methods to move the study of human cognitive processes out of the experimental psychology lab and into real-world settings, ranging from ordinary activities to sophisticated work practices. As Edwin Hutchins argued, “We need to look in the wild, not because that is where cognition is, but because it is a place where it is easier to see the cultural nature of cognition” (Hutchins 1996, 67).[16] Originally, ethnography was developed by anthropologists as a method by which to study and interpret cultural and social practices of indigenous communities. In the late 1970s, it began to be used to study the practices of other kinds of communities as situated in their natural settings. As methods of “qualitative analysis” began to develop, Egon Guba, a pioneer in promoting ethnographic educational research, characterized ethnography, generally, as a form of “naturalistic inquiry” (1978; Lincoln and Guba 1985). This characterization was introduced to contrast it, and qualitative methods, with empirical inquiry by means of controlled experimental design. A naturalistic inquiry aims to collect in situ data and extract information on practices and their relations to context through an intensive and detailed description and systematic analysis of those practices and their contextual relations. A naturalistic inquiry is ecologically valid in that there is at most minimal manipulation of existing settings, and no strict constraints, such as predetermined categories of interpretation, are placed on outcomes. Such studies are principally inductive rather than hypothetico-deductive.

An ethnographic investigation is geared toward the open exploration of practices, rather than the testing of hypotheses. The scope and focus of an inquiry are, of course, framed by its research questions, which serve to focus the ethnography.

Ethnography is interpretive research. The anthropologist Clifford Geertz characterized a main objective of ethnographic analysis as “thick description”—a term he claims to have borrowed from the philosopher Gilbert Ryle (Geertz 1973). Thick descriptions interweave description and explanation of an observed phenomenon or practice by unpacking it layer by layer with respect to its context. In general, ethnographic investigations are built around a family of tools for gathering data, mainly field observations and interviews, and around interpretive data analysis methods, such a grounded coding (Corbin and Strauss 2008) and thematic analysis (Braun and Clark 2006). Ethnography provides systematic methods of data collection and analysis, some of which are discussed in section 1.3.2, in order to establish that the interpretations are robust and consistent across a range of evidential sources (“triangulation”), thereby establishing warrant for the claims advanced from the investigation (see, e.g., Guba 1981).

“Cognitive ethnography” (dubbed thus by Hutchins 1995a) is used to gather data in real-world settings on how conceptual, social, and material resources are integrated in cognitive processes. What makes this “cognitive” is, among other things, that the focus of the ethnography is on how individuals and communities solve problems by reasoning about them, by seeking to understand them, by altering the concepts they use, by working together, by using their imaginations, and by learning. These practices are investigated as situated in contexts, with their attendant resources, which include, importantly, material artifacts. For example, one pioneer in the approach, Jean Lave (1988), rooted her critique of traditional experimental cognitive science in her own ethnographic investigations focused on mathematical problem-solving by “just plain folks” in their natural environments, such as home and the grocery store. Her studies showed in what ways people integrate environmental resources into their mathematical reasoning and problem-solving, which has the effect of making them generally much more competent at these tasks in the real world than they demonstrate as subjects in experimental psychology labs or in traditional school settings, where they are usually deprived of such resources. Meanwhile the most widely influential cognitive ethnographic research, and the one best known among philosophers, is Hutchins's research on technologically rich and well-defined problem-solving environments, where he extends "natural" to comprise specific work contexts. In keeping with the traditional cognitive science framing, Hutchins conceives of problem-solving as a form of information processing that uses representations and reasoning in pursuit of goals. However, his conception diverges from the traditional framing in that, from the D-cog perspective, the relevant representations and reasoning processes are located not only "in the head" of an individual, but also situated in the problem-solving environment and distributed across one or more individuals and select artifacts.

"Naturally occurring culturally constituted human activity" (Hutchins 1995a, xiii) of any kind can, in principle, be investigated with ethnographic methods. Research labs like the ones we have investigated certainly constitute such natural environments of scientific practice. In a cognitive ethnographic investigation, philosophers of science, too, are likely to focus on problem-solving contexts. These contexts can provide detailed information on many issues of interest to philosophers, including the nature and structure of scientific problems; how these are modified in the course of research; how scientists develop and use methods and concepts; how they create and evaluate claims and explanations; and how they communicate results.

Another way in which D-cog ethnography is "cognitive" is that a central aim of analysis besides providing richly nuanced thick descriptions of the particularities of a given case is to advance a more general, theoretical account of cognitive processes. As Hutchins has framed this objective, "There are powerful regularities to be described at the level of analysis that transcends the details of the specific domain. It is not possible to discover these regularities without understanding the details of the domain, but the regularities are not about the domain specific details, they are about the nature of cognition in human activity" (quoted in Woods 1997, 177).[17] Cognitive ethnography, although rooted in the concrete, can make use of several kinds of qualitative data analysis methods to abstract to the extent warranted beyond the details, such as "grounded theory" and "thematic analysis" (see, e.g., Corbin and Strauss 2008; Patton 2002; Braun and Clarke 2006). Such analyses aim to move from the specificity of the case to build a broader interpretive account by using systematic procedures to abstract and coalesce interpretive categories and, where appropriate, formulate candidate hypotheses to transfer and assess across cases, using multiple cases to work back and forth between data and theory to attain a warranted degree of generality. This is a different kind of process from inductive generalization. Insights from philosophy of science on how to use case-study material to build theory help to illuminate the difference.

As with cognitive science, philosophy of science, too, is interested in using empirical insights from data on scientific practices to develop or examine theoretical notions, while avoiding unwarranted generality. Early critiques, including my own, of a simple inductivist perspective on case data advocated that the way to understand the relation between specific cases and theory is as a bootstrapping method customarily used in the sciences (Nersessian 1991a), which we will see ample examples of in the case studies developed in this book. Roughly, in such bootstrapping processes, “hypotheses are made within a background of beliefs and problems. . . . They are refined, made more specific, modified, or rejected in light of more constraining data (a detailed case study). Surviving hypotheses are then tested against other data and other hypotheses to determine the extent of their validity” (Nersessian 1991a, 683). Bootstrapping entails working back and forth between data and theory, until a satisfactory accommodation is achieved. It is an iterative and incremental, open-ended process.

Recently, in thinking about the use of qualitative data on scientific practices, such as from ethnographies or interview studies, Erika Mansnerus and Susann Wagenknecht (Mansnerus and Wagenknecht 2015) follow up on a recommendation of Hasok Chang (2012, 111) to construe the relation between historical case studies and philosophical theorizing in terms of the “concrete” and the “abstract” instead of the customary inductive categories of the “particular” and the “general.” They use this suggestion to further articulate the bootstrapping account and argue that the way philosophers can arrive at “limited generalizations,” while they “avoid unwarranted generality,” is to “create a *dialogue between the abstract and the concrete*” (emphasis original). That is, to work back and forth between data and theory, which they, too, note is a bootstrapping procedure. Further, they contend, such “productive interplay” (Mansnerus and Wagenknecht 2015, 40) makes it possible to examine and further develop philosophical concepts and theories with qualitative case study data, while avoiding the pitfalls philosophy has often succumbed to of fitting the data to the theory. With ethnographic studies of scientific practice, the context for “productive interplay” is established in the way philosophers of science frame the investigation, how and what data are collected, and how analysis proceeds.

Such dialogue is in line with traditional practices in ethnographic analysis. As Geertz has emphasized, “one does not start intellectually empty-handed. Theoretical ideas are not created wholly anew in each study. . . . They are adopted and refined from other, related studies, and, refined in the process, applied to new interpretive problems. If they cease being useful . . . they stop being used. . . . If they continue to be useful, throwing up new understandings, they are further elaborated and go on being used” (1983, 57).

To consider here, in general, how ethnography can be adapted from its social science roots to serve as a method of philosophical investigation, in line with philosophers’ interests, goals, and values, would take us too far afield from the subject of this book. I have undertaken that analysis elsewhere in order to promote an ethnographic approach in philosophy of science (Nersessian and Macleod 2022). Instead, I next detail, specifically, how we framed and developed the multiyear ethnographic investigation of epistemic practices in bioengineering sciences research labs that is discussed in this book. The chapters that follow provide the fruits of our investigation.

1.3.2 Our Own Cognitive Ethnography of Bioengineering Labs

The “wild” of our ethnographic investigations is the university bioengineering sciences research lab and the researchers within it. We have been making use of cognitive ethnography and the broadened framework of D-cog to investigate scientific labs and the problem-solving practices within them as distributed cognitive-cultural systems. As such, “the lab” is not simply a physical space existing in the present, but a dynamic problem space that reconfigures itself as the research program moves along in time and takes new directions in response to what occurs both within the lab and in the wider community of which the research is part.[18] My choice to study specific research labs was largely serendipitous, but once the opportunity presented itself it was apparent they were ideal loci in which to investigate how cognitive-cultural integration proceeds in the reality of everyday research. The bioengineering sciences are a hotbed of creativity and innovation, including of ways to go about doing research on complex biological systems. The pioneering, interdisciplinary nature of the research offers a perfect opportunity to study investigative practices as they are created, as well as how they are used.

The central epistemic practice across all labs is building some kind of model of a complex biological system to use as a basis for inference and understanding about the system. Our approach to cognitive-cultural integration was to frame these labs and the problem-solving configurations within them as D-cog systems and to examine the various components these comprise and their interrelationships. The project of cognitive-cultural integration is rich, multidimensional, and difficult. I do not claim that the analyses presented here fully cover this richness or dimensionality in the problem-solving practices of the labs we investigated. Among the specific components of a distributed cognitive-cultural system we attended to are conceptual and methodological resources, artifacts central to the research, and epistemic norms and values. We carried out several kinds of analysis with respect to these, as will be discussed throughout the book.[19] In the analyses I present in the chapters of this book, I first attend to the specific nature and configuration of these aspects, as well as the interactions among them, in each lab. I then examine specific *epistemic affordances* of each D-cog system. Epistemic affordances, as I characterize them here, are those features of the D-cog system that enable or facilitate epistemic access to the phenomena under investigation.

The bioengineering sciences are inherently complex in at least three ways: how researchers think, what technologies they work with, and how they work together. Researchers in these fields make use of a range of conceptual, methodological, theoretical, and material resources drawn not only from various engineering fields but also from the biosciences and from computational sciences—resources they use to conduct groundbreaking, basic biological research in the context of potential application. It is a field initiated largely by engineers, who are recasting specific biological problems as bioengineering problems. For example, they recast the problem of changes in functional properties of endothelial cells in the cardiovascular system in terms of the effects of mechanical forces on them. They aim, ultimately, to *get a grip* on complex biological systems in both senses of the word “grip”: to understand and to control. The understanding they aim at is to develop a model of a complex biological phenomenon, for instance, a model of its underlying mechanism or a mathematical model of the interactions among the components of a system. Researchers in bioengineering sciences hope that such a model will offer the possibility to control specific processes, such as disease processes (usually by others, for instance, medical researchers). However, given the frontier nature of the research, the initial aim of the researchers is to develop, understand, and control a physical or computational model that has the potential to be informative about the behavior of the biological system, as was the case in the labs we investigated.

The movement of engineering into biology has given rise to a multifaceted interplay of quite disparate conceptual frameworks, methodological approaches, and epistemic values (see also Boon 2011). Figuring out what resources to draw from and how to adapt them is in no way straightforward. For instance, concepts developed in the context of engineered systems, such as robustness, modularity, and noise, are being transferred to the study of biological systems, but their attempted transfer has required not only modification, but often conceptual innovation (Knuuttila and Loettgers 2011; Loettgers 2007; Nersessian 2012a,b). Further, mathematical engineering theories and frameworks, as well as engineering methodologies, are being imported and adapted to perform as tools of biological representation and analysis (Wimsatt 2007). Along with these tools, bioengineers transfer certain values embedded in the practices of the engineer, such as precision, control, isolation, and abstraction—values that often conflict with assumptions and epistemic values of their collaborators in the biological sciences.

At the outset, bioengineering scientists face a major challenge in that they usually cannot experiment on biological systems directly. The complexity of the real-world (in vivo) phenomena makes experimentation too difficult, or even impossible, to control. Also, to intervene on in vivo systems often presents significant ethical issues. Thus, researchers in these fields need to devise means to model the phenomena in sufficient measure to enable experimentation on the model to yield plausible understanding of the in vivo system. Physical simulation models (in vitro) and computational simulation models (in silico) are artifacts designed to function as epistemic tools (Boon and Knuuttila 2009; Knuuttila 2011). They are part of the epistemic infrastructure through which bioengineering scientists manage and probe the nature of complex biological systems. But to devise an appropriate modeling practice as well as specific models faces all of the challenges of biological-engineering integration noted above. Thus, by investigating specific practices we are also laying out the basic epistemic structure of biological engineering; namely, to bring the conceptual, methodological, technological, and material resources of engineering to bear on the problem of managing the complexity of biological systems so as to be able to study them.

Generally speaking, bioengineering scientists investigate biological systems in either of two ways. They work through iterative and incremental processes of designing, constructing, redesigning, evaluating, and experimenting either with surrogate in vitro physical simulation models, which comprise biological and engineering materials, or with in silico computational simulation models. We refer to these basic processes as *building to discover*. The processes themselves, and the question of how they advance their epistemic goals, have been the focus of our research on cognitive-cultural integration. It would not be far-fetched to conclude that understanding in bioengineering sciences, and engineering sciences in general, derives largely from *building*. It is no accident that the famous saying attributed to Richard Feynman, "what I cannot create, I do not understand," is misquoted by bioengineering scientists as "what I cannot build, I do not understand" (famously encoded in the first synthetic cell by Craig Venter; see also Voit et al. 2012). However, this process of building to discover has not received much attention in either philosophy of science or cognitive science. Our research shows that—and how—building simulation models is a major means through which bioengineering scientists build understanding of complex biological phenomena. Model-building is the means by which they actively distribute cognition in their environment, thereby creating complex distributed cognitive-cultural systems of people, practices, and artifacts.

The bioengineered models thus function as cognitive artifacts, which participate in the reasoning processes and the representational processes of a distributed cognitive system. They function equally as what researchers in STS studies of science refer to as the "material culture" of communities, which participate in social and cultural processes. Our research demonstrates on a day-to-day basis that it is simply not possible to fathom the epistemic work such models enable by focusing exclusively on one or the other aspect. Bioengineered models are representations, and as such they play a role in reasoning processes. They are central to social practices related to community membership, such as mentoring and identity formation. They are sites of learning. They embed epistemological norms and values. They are repositories of lab history. They perform as cognitive-cultural "ratchets" (Tomasello 1999) in an epistemic community, which enable one generation to build on the results of the previous, and so serve as loci of stability in the context of innovation, while moving the problem-solving forward. In sum, *models are central components of the cognitive-cultural fabric of creative problem-solving in bioengineering sciences*. Examining how that fabric is built in specific cases provides insight into how cognition and culture are intertwined in scientific practice. It is for all these reasons that we have focused our investigations on the modeling practices we have encountered in the research labs.

1.4 Four Cognitive Ethnographies: An Overview of Our Bioengineering Sciences Project

Our investigational settings comprise four pioneering university research laboratories. Two labs work in biomedical engineering (BME)—tissue engineering and neural engineering, respectively. The other two specialize in integrative systems biology (ISB)—one solely computational, the other a combined computational and wet experimentation lab. We chose university labs because they are largely populated by graduate students, who are pioneers in research and at the same time learning to become scientists. Many of the graduate students were in an educational program aimed at moving beyond collaboration between engineers and biologists through producing *hybrid* biomedical engineering researchers. The educational program was itself under development at the time, and we were assisting the faculty in this undertaking. As a central part of our NSF-funded project on the labs, we had proposed not only to investigate the nature of emerging practices in these fields, but also to determine important requirements for learning them, and to work with bioengineering faculty to "translate" findings from our research into educational experiences. Our investigations taught us, and the faculty as well, that the forms of interdisciplinarity practiced in BME and in ISB are quite different, leading to different environments, requirements, and challenges for problem-solving and learning.

We chose to investigate interdisciplinary fields at the intersection of biology and engineering, but the selection of labs was largely a matter of chance, and we had no prior understanding of the fields. Each field was conducting research on complex biological systems about which there was little scientific understanding. In our first pilot research, the tissue engineering lab, we were surprised that its physical space had the look of a biology wet lab, with pipettes, flasks, sterile hoods, and cell cultures, but also included strange-looking engineering artifacts, which members were referring to as "devices." We quickly realized that these artifacts were the focal point of the research life of the lab, as they provided the means of what the director called "*taking the research in vitro*." These lab-built models are intended to replicate selective processes of complex biological systems. We also quickly realized that conducting physical simulation experiments with these hybrid models—part living tissue and cells and part engineered materials—constitutes a novel modeling practice not previously investigated in the philosophical literature. We made these models and the practices surrounding them our focus in BME. Similar pilots with each lab directed our attention to the salient modeling practices that became the focus of our data collection and analysis.

What labs? Now that we had discovered the novel practice of building in vitro physical simulation models, we wanted a second BME lab with that practice in a different domain. We also thought that a newly established lab would be good for contrast, which is why we chose the neuroengineering lab. We began our ISB study because we wanted to move into a pioneering area in computational modeling and simulation of biological systems, and the director of the purely computational lab was interested in how we might assist research and learning in that area. By then we had established an excellent reputation in the department, and lab directors often asked if we could "do" their lab. In addition to our contributions to the development of their educational programs, the faculty began to notice that the researchers in the labs we studied had become more reflective about their research. Indeed, the researchers, themselves, told us that our interviews always provided an opportunity to reflect on their research problems and, on occasion, served to rekindle motivation when they were tired and lagging. As one researcher told us, "*talking about it is good 'cause it also reinforces what you're doing. So, I can go back and feel motivated about it now*." For a second ISB lab we sought one that conducted experimental research as well as computational modeling, and also was newly established. When an assistant professor offered her lab, we concluded from her description that it would contain experimentalists as well as modelers. In our pilot we were surprised to discover that the modelers themselves were being trained to conduct the experimental work. Since the challenges of the dual nature of the practice were novel and interesting, we decided to stay with that lab. We coded the practice as that of the "bimodal strategy" (MacLeod and Nersessian 2013).

In both BME labs, the researchers tackled problems of bringing cells and tissues together with engineered materials in processes of building living hybrid models to simulate in vivo processes of complex biological systems.

Their aim was to understand basic processes within the biological systems, with the hope of providing a basis for future application, chiefly in medicine. Their research problems required little collaboration with researchers outside the lab.

Both ISB labs aim to understand biological systems that comprise integrated and interacting complex networks of genes, proteins, and biochemical reactions. When this is achieved in sufficient measure, they expect it will allow collaborators to attempt interventions such as to produce a better biofuel or to make cancer cells receptive to treatment. Solutions to the problems they tackle require that they construct computational simulation models in need of rich experimental data, which creates an essential epistemic interdependence with their collaborators in the biological sciences.

1.4.1 Research Questions and Data Collection

Our objective in data collection can be summarized as follows. Starting from an open and broad stance about what might prove relevant to our research questions, we aimed to conduct a systematic longitudinal investigation involving numerous bioengineering scientists across a broad range of perspectives, problems, and lab organizations. In each case we aimed to collect a range of data from different sources with which to triangulate the analyses.

Our investigations in each lab began with a basic set of questions motivated by our combined philosophical and cognitive science interests:[20]

- What are the representational and reasoning practices used in problem-solving in this community?
- How is epistemic warrant developed for novel practices?
- What are the epistemic assumptions, values, and norms at play in each of these interdisciplinary communities?
- What concepts, methods, and theories are being used from engineering, and how? Ditto from biology? What is the nature, or are the results, of their interaction?
- In what ways might cognitive, social, cultural, and material “factors” be mutually implicated in these epistemic practices?

These questions enabled us to focus our research while remaining open in sufficient measure to guide, but not totally constrain, our data collection and our analysis. Within the course of the investigation, we intended that specific issues to be addressed with respect to these questions would emerge from our findings, as indeed they did.

We investigated each lab for approximately five years. We collected data between 2000 and 2014, and data analysis continues in the present. Data collection in all labs comprised the following main items:

- audio-taped open and semi-structured interviews
- participant field observation with note-taking
- lab tours (given for us and for visitors)
- arranged demonstrations of experimental procedures and technologies involved in the lab researchers' data collection and analysis
- video and audio recorded lab meetings
- "journal club" meetings in which pertinent articles were discussed
- photographs of white boards
- diagrams of the spatial layout of each lab and photographs of how lab space changed over time
- artifact collection: grant proposals, paper drafts, presentations, dissertation proposals, emails, diagrams/sketches, and so forth

For each lab we compiled an extensive "technology document" that surveyed all the technologies in the lab, which researchers used them, and what their functions were within the research.

The extent of our interview and observational data is summarized in table 1. All interviews and some (or parts of) meetings we thought especially significant have been transcribed. In our interviews and discussions with the researchers, we probed the nature of their research, how they went about doing their research, what problems they were encountering, how they responded to those, how and what kind of learning was needed as they went along with their research, and how they positioned their research with respect to the broader field. We began with unstructured interviews, with the initial ones focused on their background, motivation for choosing that bioengineering science, and an overview of their research. As we gathered more information on their projects and as we developed a better understanding of the scientific/engineering content and methods, we conducted more targeted interviews both to probe their reasoning as they were working on specific problems and to probe specific issues that arose as we began to analyze the transcripts and other data. We learned a great deal about their practices and about the content and context of the research. Often, when making field observations and when it would not interfere with the research, we were able to ask questions about what they were doing at that time and why. They also were willing to set up meetings with us to demonstrate procedures they used in model-building and data analysis. In return, as I noted, our probing provided the opportunity for them to articulate and reflect on their research. One student likened our interviews to "research therapy" appointments, and several even expressed the desire to have them more frequently. Finally, we had many informal interactions with them. Our student researchers went along on their group hikes and bike rides when invited, and sometimes ate lunch with them. We were all invited to holiday parties and to dissertation defense celebrations. Some students asked for career guidance or letters of recommendation.

1.4.2 Research Sites

I now provide a brief overview of the makeup and the kinds of problems addressed in each lab in order to aid the reader in understanding the different challenges of interdisciplinary problem-solving faced by the researchers as they attempt to integrate engineering and biology in these environments. All the labs we investigated were conducting research for which there was little or no precedent when they began. The members repeatedly told us about the pioneering nature of their research, using expressions such as "it had never been done before" or "no one has approached it this way before" or "no one was thinking this way" or similar such expressions.[21]

In the two BME labs we conducted intensive data collection over the first two years. For approximately five years, we further continued data collection on selected dissertation projects up until they were completed, including

Table 1.1
Data summary

Laboratory	Interviews	Meetings	Field observations (hours)
BME A	72	15	~350
BME D	75	40	~450
ISB G	44	7	~40
ISB C	62	22 (plus 2 joint C and G)	~250

additional interviews.[22] We began data collection in lab D approximately a year after we began collecting in lab A. In all, we worked on the BME project for ten years, before moving on to ISB. As noted before, both labs designed, built, and conducted experiments on hybrid living physical models, locally called "devices," which is the shorthand they use for "bioengineered modeling devices."[23] Given the distant nature of their respective research, there was little interaction between the directors of these labs, beyond the attention and the informal mentoring one would expect a quite senior member of a department (lab A) would provide to a quite junior member (lab D) of his department.

Lab A was a tissue engineering lab. Its overarching research problems were to understand mechanical dimensions of cell biology, such as the effects of the forces of blood flow on morphology, proliferation, and gene expression in cardiovascular endothelial cells. The researchers saw their research also as a contribution to the eventual medical application goal of creating a living substitute blood vessel to implant in the human cardiovascular system. Examples of intermediate problems that contributed to the daily work included constructing specific living tissue models ("constructs") that mimic properties of natural blood vessels; using biomechanical forces to create endothelial cells from adult stem cells and progenitor cells; designing environments for mechanically conditioning constructs; and designing means for testing their mechanical strength and functionality.

When we entered lab A, it had been in existence for thirteen years. It closed ten years after our study, when the director retired. During our study, the main members included a male director, a male laboratory manager, a female postdoctoral researcher, seven PhD students (two male, five female, two of whom, one male and one female, graduated early in our study, the other five after we concluded our formal data collection), two MS graduate students, and four long-term undergraduates. Additional undergraduates from around the country participated in summer internships, and international graduate students and postdocs visited for short periods. The laboratory director was a senior, highly renowned pioneer in the field of biomedical engineering, who had started his career as a mechanical engineer in aeronautical engineering. Near the end of his career, he liked to characterize its trajectory as "from astronauts to stem cells." All of the researchers had engineering backgrounds, mainly in mechanical or chemical engineering, and some were currently students in the BME program that was just starting. Some had spent time in industry before joining the lab. The lab manager had an MS in biochemistry. The researchers frequently consulted with a histologist located in the building, and some traveled to other institutions for various purposes, including to collect animal tissues and to run gene microarray analyses. Lab meetings were held irregularly, when the director, who traveled a significant amount of time, would be in town (approximately every three to four weeks).

Lab D was a neuroengineering lab. Its primary research problem was to understand the mechanisms through which neurons learn as networks in the brain. Here, again, the researchers had dual scientific and engineering goals. They aspired to use this knowledge to develop aids for neurological deficits or, more generally (as the director liked to say), “to make people smarter.” Here are some examples of intermediate problems that contributed to the daily work. They developed ways to culture, stimulate, control, record, and image neuron arrays. They designed and constructed feedback environments (robotic and simulated) through which the main device (the model-system comprising a “dish” of cultured neurons) could learn. They used electrophysiology and optical imaging to study “plasticity.” One researcher developed a computational model of the dish model-system that played an unanticipated pivotal role in the research. All the projects centered around the “dish,” and, as the research unfolded, there developed significantly more interaction among research projects than we witnessed in lab A.

Lab D was just taking shape as we began our research. It closed when the director moved to another position, which was nine years after our study ended. During our study the main members included a male director, a male postdoctoral researcher, four PhD students in residence (one female, three male; one male left after two years to pursue neuroscience, and the remaining three graduated after we concluded formal data collection), one PhD student at another institution who occasionally visited the lab and was available via video link, one MS student, six undergraduates, and one volunteer for nearly two years, who was not pursuing a degree (already possessed a BS in engineering) but who helped out with breeding mice. Because the lab was new and had limited funding at the start, the director made more use of undergraduates, who usually had short-term research projects, which were supervised by the director for course credit.

When we began, the laboratory director was a new tenure-track assistant professor, fresh from a lengthy postdoc in a biophysics laboratory that develops techniques and technologies for studying cultures of neurons. He already had attained some recognition as a pioneer. His background was in chemistry and biochemistry, with his engineering knowledge largely self-taught, though highly sophisticated. The backgrounds of the researchers in lab D were more diverse than those in lab A and included mechanical engineering, electrical engineering, physics, life sciences, chemistry, and microbiology; some were currently students in the BME program, but also in electrical engineering and mechanical engineering. The wet lab was in a separate room. The main lab had the look of a computer lab, with copious wires connecting the incubator for its main object of study—the dish—to computers, and with small robotic devices, connected with a dish, scattered (or rolling) around the lab. They held lab meetings and a journal club (to discuss recently published research) weekly. Unlike the traditional configuration of a stand-alone lab, lab D was embedded in an open space designed to promote interdisciplinary collaboration among neuroengineering labs. It was shared by seven faculty members, their postdoctoral researchers, and graduate and undergraduate students.

The members of the two labs we studied in systems biology, lab G and lab C, preferred the name “integrative systems biology” (ISB) for the area in which they worked. They explained that “integrative” stressed both the integrating function of building a model of a biological system, as well as their research aim to integrate a range of resources from biosciences, engineering, and computational sciences in their investigations. In the ISB study we had less funding and fewer researchers for our project, so we conducted intensive data collection in both labs over the first year and followed selected dissertation projects through to completion for a total of five years. We started data collection in the two labs at the same time. In both labs the primary focus was building computational simulation models. There was significant interaction between the directors of these labs, though not much among their students. There were few ISB researchers in the department, and lab directors were hopeful they could build the area together, and possibly start an educational program aimed, specifically, at aspiring ISB researchers. They worked together with us to develop a graduate-level introduction to biosystems modeling course, which they co-taught. The lab G director provided some mentoring to the quite junior lab C director, even though, as we will see, they had quite different “philosophies” about how to conduct ISsB research and for what purposes.

Lab G is a purely computational systems biology lab, with the clever motto, “where life becomes numbers and numbers come to life.” Its research problems focus on computational simulation modeling of biological systems at the genetic, metabolic, and cellular levels. The focus of the modeling is on the interactions among different components of biological systems (such as metabolic and signaling pathways), rather than on structural properties of specific components (such as DNA and ribosomes). The problems addressed are wide-ranging, and usually brought to the lab by biological researchers from universities and industry because of the outstanding reputation of the lab director as a pioneer in ISB. For instance, one of the problems tackled by the lab was to develop a model of the production and transport of dopamine and of how this system is affected in Parkinson’s disease. In this research, the lab worked with experimental data provided by a medical research group specializing in neurodegenerative disorders. Another problem was to develop a model of ethanol production using algae, based on data provided by researchers at a biofuels company. In general, the domain-driven problems are provided by bioscience researchers of various kinds who approach the lab, asking the director to “model our data,” usually with little understanding of what that means or entails. The overarching focus of the lab’s own agenda is on methodological problems specific to computational modeling of biological systems, especially developing mathematical techniques and algorithms to improve the estimation of model parameters and the optimization of these parameters.

During our study the main lab members included a male director, four postdoctoral researchers (two female, two male), and four PhD students (one female, three male). The members of the lab varied widely in terms of educational background, although most were from engineering (mechanical, electrical, telecommunications, biomedical, computer). Other backgrounds included pharmacy, applied physics, bioinformatics, and information sciences. The main criteria for being accepted into the lab were applied mathematical and/or systems computational skills. A postdoctoral student who was from a collaborating experimental lab in Europe visited periodically for a month or so at a time; he had a PhD in biochemistry and was transitioning to modeling. A striking feature of the lab is that all members were from outside the United States. Eight were from Asia (China, Taiwan, Japan), two from Europe (including the director), and one from the Middle East. As a consequence of the sophisticated computational modeling skills the research requires, there were no undergraduates. The lab director is a senior pioneer in ISB, with an undergraduate degree in natural sciences and mathematics, two master's degrees (one in biology and the other in mathematics), a certification in philosophy and education, and a PhD in theoretical biology. The lab had been in existence for five years when we entered (the director's previous lab ran for fifteen years at another institution). The lab space consisted of desks with computers, and was quite often empty, since lab members could just as easily work at home on their laptops. For this reason, we stopped aiming for field observations after a few months. Research meetings were largely conducted one-on-one with the director, and they did not hold lab research meetings, though they did organize a few so we could get an overview of the research, and the researchers noted that these were beneficial to them as well. The lab researchers had a range of biosciences collaborators external to the lab (some of whom we interviewed).

Lab C is an ISB lab that conducts both computational modeling and biological experimentation in the service of model-building. Its research is guided by an overarching biological problem: to understand the impact of redox (reduction-oxidation) environments on proteins through systems modeling approaches. Under normal physiological conditions cells maintain a reduced internal environment. However, oxidizing molecules and free radicals that are produced in the cell as a part of physiological processes, or that enter the cell, can react with cellular components such as DNA, cell membranes, and proteins. Such reactions have physiological consequences and have been implicated in several diseases. Lab C's research focus is on the impact of alterations made by oxidants on proteins, which are part of signaling pathways, and on the dynamics and outcomes of these pathways. Based on her own training, the director has been training the graduate students, who have engineering backgrounds, to do biological experimentation in the service of building and testing their computational models. One student also engaged in engineering design through a collaboration to develop a microfluidic device ("lab-on-a-chip") to produce high-throughput single-cell and population data, which are the time-series data more amenable to quantitative investigation. The lab's overarching problem translates into specific research projects as varied as modeling chemotherapeutic drug resistance in acute lymphoblastic leukemia cells and modeling senescence in T cells. However, everyone in the lab was aware of what the others were working on, and provided feedback on the research projects of others in the weekly lab meetings and in weekly journal club meetings. We witnessed many instances of joint troubleshooting, in particular.

Lab C was just taking shape as we started our research. During our study the main lab members included a female lab director who was a new assistant professor, five PhD students, two of whom joined the lab after we started our observations (three male, two female), six undergraduates, and a female research technologist/lab manager with an MS in biology (who transitioned to a PhD student, while remaining manager, late in our study). A striking feature of this lab is that the researchers spanned four continents (North America, Europe, Africa, Asia). The lab director has an undergraduate degree in nuclear engineering (with a minor in biomedical engineering) and a doctoral degree in bioengineering, during which she first trained as a modeler and then as an experimentalist, followed by a postdoctoral period in a bioengineering lab that comprises both computational modelers and bioscientists. The graduate student backgrounds were predominantly engineering-related (electrical, biomedical, biotechnology, material science). A joint MD/PhD student had a background in chemistry and mathematics. The undergraduates mainly ran western blots and other experimental procedures for the graduate students. The lab's experimental biology research is conducted in-house, but they had a few external engineering and bioscience collaborators and bioscientists with whom they consulted during the period of our investigation, some of whom we interviewed.

1.4.3 Data Analysis

Numerous qualitative methods can be used singly or jointly in cognitive ethnographic data analysis. We have been using a variety of mutually complementary qualitative methods, specifically, qualitative data coding, case study analysis, thematic analysis, and cognitive-historical analysis. These are among a wide range of qualitative methods that have been developed and critiqued extensively over the last half century, especially in psychology and sociology (for an overview, see Patton 2002).[24] There are no formulas or recipes for how best to apply those qualitative methods in any specific case, so we have needed to tailor and innovate our data analysis with respect to our research goals and questions—as is standard in qualitative analysis—while adhering to accepted canons of what constitutes "trustworthy" (Lincoln and Guba 1985) and "validated" data collection and analysis procedures. Although "valid" is reserved in philosophy for logical argumentation, it is often used to signify credible qualitative research in that field. I prefer to use Guba's more neutral term, "trustworthy," when considering issues of warrant.

To establish trustworthiness, we have, in particular, taken into account the American Psychological Association standards (see, e.g., Eisner 2003, who argues standards need to take into account that qualitative analysis is both science and an art). We have followed three standard principles in particular: structural corroboration, referential adequacy, and consensual validation (Eisner 2003). Structural corroboration requires that a sufficient number of data points converge on a conclusion to support an interpretation. Referential adequacy addresses the richness and clarity of the description and interpretation, and how it aligns with member understanding. Consensual validation refers to the level of agreement that can be reached among two or more researchers in developing and using the coding schemes ("interrater reliability"). Adherence to these principles required that we would systematically collect the range and kinds of data sufficient to triangulate data from multiple sources in order to corroborate and determine the referential adequacy of interpretations. "Triangulation" in qualitative analysis refers to the processes of building warrant for an account through establishing consistency of findings across methods and sources of data collection. Our research conducted long-term studies that provided a variety of longitudinal data, which (as noted) consisted of persistent observations, of multiple interviews of each participant, and of the kinds of archival data previously mentioned.

Data collection in an ethnographic study always risks the dual charge of being not representative and/or subject to bias stemming from the researcher's own interests, values, and motivations. Ethnographic investigation demands continual self-scrutiny so as to mitigate researcher bias, which is an issue in all empirical research. Such self-scrutiny, for example, would control for asking leading questions in an interview or for importing favored notions into data analysis. In general, it is important to keep in mind that all ethnographic research is interpretive. As such, the researcher is the instrument of data collection and analysis, and, so, the researcher's interests, values, and motivations are always present, and it is a necessary part of good research to be explicit about and confront these (Osbeck and Nersessian 2015, Nersessian and MacLeod 2022). We were aided in our attention to potential researcher bias in data collection and analysis by the unusual approach we took to conducting ethnography. Unlike traditional practices, where the ethnographer is a single researcher, we decided to practice what we dubbed “team ethnography.” In each given lab, more than one ethnographer was responsible for observations and interviews, and our more senior members worked across the labs. As project director, I oversaw that in all labs we collected comparable data to the extent possible. Our research group varied in size and composition over time (undergraduates through senior faculty) but remained highly interdisciplinary and thus provided multiple lenses through which we could examine the data.[25] Our weekly research group meetings provided the venue for scrutinizing and evaluating the ethnographic work together as it unfolded, and for reaching consensus on coding, theme development, and other forms of data interpretation. As data analysis progressed, we related our findings to the appropriate philosophical and cognitive science theoretical frameworks. We also formulated hypotheses with respect to issues within these frameworks, and together we evaluated, revised, or refined these in comparison to our empirical analyses.

Since coding is the first method through which one starts to make sense of the data, I next briefly describe a few of the procedures we used in our coding analysis of data. Although we used a variety of complementary methods of data analyses, as noted earlier, procedures that relate to systematic, fine-grained open coding and to grounded theory development (Corbin and Strauss 2008; Glaser and Strauss 1967; Strauss and Corbin 1998) provide the primary basis for our interpretations.

Coding is an interpretive procedure by which to partition the data by attaching descriptive categories to units of interview texts and field observation notes. Our approach to coding can be broadly characterized as “grounded” in the sense described by Corbin and Strauss (2008). We understood this to mean, in particular, that we remain open to seeing what categories/themes might emerge from the data. While, obviously, our coding was guided by our research questions and objectives, it was by no means restricted by them. We developed our coding procedures in several phases. We established coding procedures to mitigate, to the extent possible, issues of subjectivity—we even hired an external coding auditor midway in our investigations, by way of a check on our procedures.

We began with “open coding” directed toward identifying, categorizing, and describing what the text of the interview is about. During this process, coding pairs worked together on each transcript. We analyzed a subset of interviews progressively, line by line, with the aim to provide an initial description for as many textual passages or "meaning units" as seemed appropriate. In our research meetings, the entire group discussed the clarity, fit, and logic of the codes assigned. In early coding, we presented interpretations to the research lab members by way of checking whether their views aligned with our understanding. More than that, we used feedback from all pertinent sources to make adjustments.

We continued coding additional interviews, revisiting previous coding, and assessing descriptions for adequacy and for fit throughout the process, as is consistent with the goals of analytic induction (codes emerging from data and leading to hypotheses) and constant comparison (codes compared against possible alternative interpretations) (Lincoln and Guba 1985; Corbin and Strauss 2008). After about 20 percent (the standard) of the interviews were coded intensively in this manner, we coded the rest more selectively, focusing on categories of most relevance to our research questions and building out those categories. During research group meetings we reviewed all codes, and further grouped and arranged codes into superordinate categories and subcategories. We then related the codes to each other and developed the categories/concepts more directly with respect to our research questions as a start toward building "theory." In this context, we understood this process, broadly, as formulating "a set of well-developed categories (themes, concepts) that are systematically interrelated through statements of relationship to form a . . . framework that explains some phenomenon" (Corbin and Strauss 2008, 55) and allows forming hypotheses. Theory development in effect, then, took the form of developing increasingly refined conceptual models.

We coded separately for each lab, and then assessed the candidates for transfer across the labs in BME and ISB. Exemplars of lower-level codes from the BME study include analogy; model-based reasoning, understanding, or explanation; problem formulation; anthropomorphism; epistemic values; and pragmatic focus. Exemplars of superordinate categories include model-based cognition; seeking coherence; norms; and affect. In all, we developed seventy-three codes, and, with respect to codes that transferred across the BME practices, we organized these into thirteen superordinate categories. We did not use a coding software, preferring instead the traditional method of coding by hand. We developed our own coding database, using

MS Excel, which lists each category and code with an associated description and memo discussing it, and the codes have multiple exemplars from the interview texts attached to them from each lab. Codes can be easily organized and reorganized into categories with this software. We used codes and categories to create case analyses, which are finely detailed descriptions that follow practices of a specific researcher, or small group, as they worked toward solving a complex problem. We also developed longitudinal case studies specific to learning in the labs.

Code development, however, is more than mere description. It is an abstractive process, in which a code is both derived from and scrutinized in light of multiple exemplars across different interview texts within the study. Codes provide, also, the basis for cross-study comparison and for developing hypotheses to consider and assess for transfer, when detached from case-specific details. As our research progressed, we continued to assess transfer of selected major categories and themes across the labs. We were especially interested in what commonalities there might be in the general features of the model-building practices in these subfields of biological engineering and how these advance the epistemic goals of the subfields, in practices developed to support learning in the context of research, and in the challenges presented by the kind of interdisciplinarity. A major example of a cross-cutting category—or "theme"—that emerged from the BME labs and is developed in the chapters that follow is the multidimensional system notion of "interlocking models." This notion serves to articulate how multiple dimensions of these interdisciplinary research labs are built and fitted together as cognitive-cultural systems. Models interlock biological and engineering concepts, methods, and materials. They interlock in their design and construction and in experimental processes. Mental and material models interlock in model-based inference. In the latter instance, "interlocking models" is a specific kind of coupling between researcher mental models and artifact models as components of a distributed model-based reasoning system. Further, epistemic and sociocultural practices interlock in building models. We found the challenges of building systems of interlocking models to be central to research and learning for hybrid researchers. Table 1.2 provides a schematic overview of the major interlocking models with respect to the tissue engineering lab A, which is elaborated in chapters 2 and 4. The models interlock both within and across the categories.

Table 1.2
Interlocking models in lab A

Interlocking models			
Biological, engineering, medical models in the wider community (as detailed in journals, textbooks, etc.)			
cell biology	electrical engineering		
biochemistry	mechanical engineering		
fluid dynamics	disease processes		
Bioengineered in vitro artifact models			
flow loop	pulsatile bioreactor		
construct	baboon model-system		
Researcher mental models			
in vivo and in vitro phenomena			
devices qua in vitro models			
devices qua engineered models			
Sociocultural models			
Mentoring	Identity	History	Epistemic values

Finally, we never ceased using our codes and categories to analyze the data further and to examine them through various theoretical lenses in order to develop thick descriptions and analytical insights—in particular with a goal to extend, enrich, and critique philosophical concepts and theories in "productive interplay." As any ethnographer would agree, an ethnographic analysis is never complete. The greater the depth of the analysis, the more one sees what needs to be analyzed, as well as what additional data it would have been useful to have collected. This book presents yet a further analysis, which, hopefully, provides an exemplar of how to develop a conceptual model of the dynamics of cognitive-cultural integration in scientific problem-solving. Although undoubtedly there are other strategies to develop a cognitive ethnography, ours shows, in particular, that philosophers are well able, and well positioned, to work with ethnographic norms while pursuing philosophical targets of investigation.

1.5 Overview of the Chapters

Chapter 2: Building Hybrid Simulation Devices: Distributed Model-Based Reasoning. This chapter focuses on the BME in vitro simulation devices as cognitive-cultural artifacts that enable distributed model-based reasoning. The chapter provides an analysis of the iterative and incremental processes of *building*: designing, constructing, redesigning, evaluating, and experimenting with in vitro devices. The devices are physical simulation models comprising part-living, part-engineered materials. One tissue engineer called the practice of building these devices "*putting a thought into the bench top and seeing if it works*," which the chapter interprets as building "distributed model-based reasoning" systems. The hybrid model-systems simultaneously provide simulations of biological processes, the researcher's current understanding of these, and the epistemic culture's norms and values. The chapter introduces the analytical theme of "interlocking models" and examines how devices provide hubs for interlocking many dimensions of practice. It further examines how these in vitro models are "built analogies" that are designed to simulate the behaviors or functions of selected in vivo biological processes. This form of model-building expands the epistemic practice of "building the source analogy" (Nersessian 2008), wherein analogies are designed for the purposes of scientific investigation by analogical displacement. It details several examples of how the researchers build epistemic warrant for in vitro devices, as well as for the methodological practice itself, in the processes through which they create the models, focusing, specifically, on the relationship between analogy and exemplification. (Lab A and lab D data)

Chapter 3: Engineering Concepts: Conceptual Innovation in a Neuroengineering Lab. This chapter focuses on the interplay between conceptual innovation and modeling practices. The pioneering nature of the labs leads to researchers investigating novel phenomena that have been conceptualized only partially or not at all. Thus, frontier problem-solving often requires conceptual innovation. The chapter follows the researchers in lab D in their quest to understand and control the behavior of a living network of neurons. At the outset of the research, they transferred concepts from engineering and single-neuron studies to get a grip on the model. These resources both facilitated and hindered their problem-solving. Ultimately, the research led them to develop fundamentally novel concepts. The analysis starts at the point where the researchers were failing to understand and control perplexing in vitro model behavior, which led one to introduce a novel practice for the lab: computational (in silico) modeling of the in vitro model. The epistemic affordances of the in silico model promoted concept formation and change as the researchers worked toward interpreting both the in silico and the in vitro behaviors, including the relations between the models. The chapter details how, over the course of two years, the cross-breeding of these two kinds of simulation models created a cluster of scientifically novel (and potentially highly significant) concepts, while also building a D-cog system comprising all of the researchers. Armed with these new representations of the behavior, this system was able to leverage the affordances of both models to productively control the behavior of the dish model-system, and, ultimately, attain the lab's goal to establish and demonstrate that the in vitro network of neurons could learn. (Lab D data)

Chapter 4: Interlude: Building "the Lab." This chapter focuses on the theme of how "the lab" builds itself as a cognitive-cultural system. It analyzes the dynamics of how, starting from broadly framed complex interdisciplinary problems, a research lab on the frontiers of science creates and develops the cognitive-cultural structures for productive research. It examines this building process in detail for the tissue engineering lab. The chapter examines how intersecting trajectories of problems, methods, and researcher-learners develop in relation to the practice of building in vitro devices, and details how the lab's *signature* devices (examined in chapter 2), in particular, provide structuring constraints for articulation of the lab as a distributed cognitive-cultural system in ongoing flux. It ends with a brief look at the educational infrastructure developed to foster the BME goal of interdisciplinary hybridization. (Lab A data)

Chapter 5: Managing Complexity: Modeling Biological Systems Computationally. This chapter focuses on the challenges of computationally modeling complex dynamical biological systems in the absence of domain theories that can provide significant resources for building models, such as for physics-based modeling. It examines how ISB researchers, engineers with limited biological knowledge, develop practices around computational modeling and simulation that enable them to manage the complexities of modeling biological systems. The analysis shows how a close examination of the processes of building models, rather than a focus on the finished products, is needed to fathom the epistemic achievements of this emerging approach to discovery in systems biology. The chapter details a case in which an engineer with little experience in biological systems modeling and little biological knowledge was able to make a fundamental discovery in biology by means of his "adaptive problem-solving" processes. It examines, in particular, the epistemic affordances of in silico simulation for building the model while also developing a close coupling between the modeler's mental modeling processes and the biological system model. The developing modeler-model D-cog system provides another instance of what Hutchins called creating "cognitive powers," which in ISB enables a modeler to make novel, verifiable inferences about the biological phenomena that outstrip his own understanding. (Lab G data)

Chapter 6: The Bimodal Model-Building Strategy. This chapter focuses on a novel method for managing the complexity of building models of biological systems. In this practice, modelers conduct their own biological experimentation in the service of building their models. The chapter highlights the methodological flexibility available to ISB as an emerging field, which affords researchers the opportunity to tailor methods to manage complexity. The chapter develops two case studies of modelers following this strategy. The first examines, briefly, how a modeler collaborated with engineers to design a microfluidic "lab-on-a-chip device" (LOC) to integrate complex activities, actions, processes, and operations in wet-lab experimentation that would usually be carried out in many steps, by many persons, and using a range of equipment. The modeler built the LOC to solve the difficult problem of collecting time-series data needed to develop her computational model of T-cell signaling. The second case examines, in detail, how one modeler built a tightly coupled methodological system that used computational model-building and simulation to direct and focus her wet-lab experimental investigation of a biological system, while also using the experimentation to further develop her model. The epistemic affordances of this D-cog system, in particular, helped her to triangulate uncertainties and missing elements in her models without having to deal with complex problem spaces of many open parameters. In both cases the novel methodological strategies enabled the modelers to manage a range of constraints—data, computational, cognitive, collaborative—prevalent in ISB research. The chapter illustrates and provides further insights into the possibilities for adaptive problem-solving in this emerging field and how they provide researchers with considerable flexibility to create different methodological strategies and lab organizations to manage the complexities of modeling biological systems. (Lab C data)

Chapter 7: Interdisciplinarities in Action. After providing a high-level summary of the major insights gleaned from the previous chapters, this chapter goes on to consider implications with respect to the epistemic situation of interdisciplinary science as such. The analysis offers insights gleaned from all our investigations of what we call the “adaptive problem spaces” of biological engineering: spaces where interdisciplinarity is enacted in research and learning. It examines challenges, differences, and similarities across fields, and assesses their implications for broader application to interdisciplinary practice. The chapter provides a nuanced account of two major kinds of interdisciplinary practices: hybridization and symbiosis (“epistemic interdependence”). It proposes specific characteristics, *interdisciplinary epistemic virtues*, that foster creativity and collaboration in twenty-first-century interdisciplinary science, at least of those varieties, and illustrates how these can be cultivated in different research communities with targeted interventions, such as those we developed for BME and ISB. Although the focus of the analysis is on the cases we have investigated, I hope the insights in this chapter, and the book as a whole, lay the ground for future research into interdisciplinary epistemic virtues in situations beyond those cases, and beyond science to interdisciplinary practice per se. (Data from all labs)

Interdisciplinarity in the Making

Models and Methods in Frontier Science

By: Nancy J. Nersessian

Citation:

Interdisciplinarity in the Making: Models and Methods in Frontier Science
By: Nancy J. Nersessian
DOI: 10.7551/mitpress/14667.001.0001
ISBN (electronic): 9780262372275
Publisher: The MIT Press
Published: 2022

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Nersessian, Nancy J., author.
Title: Interdisciplinarity in the making : models and methods in frontier science / Nancy J. Nersessian.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.
Identifiers: LCCN 2021061880 (print) | LCCN 2021061881 (ebook) | ISBN 9780262544665 | ISBN 9780262372268 (epub) | ISBN 9780262372275 (pdf)
Subjects: LCSH: Biotechnology—Methodology. | Bioengineering—Methodology. | Biotechnology—Research—Case studies. | Bioengineering—Research—Case studies. | Biotechnology laboratories. | Scientific surveys. | Interdisciplinary research.
Classification: LCC TP248.24 .N47 2022 (print) | LCC TP248.24 (ebook) | DDC 660.6—dc23/eng/20220720
LC record available at https://lccn.loc.gov/2021061880
LC ebook record available at https://lccn.loc.gov/2021061881