Nersessian, N. J., Kurz-Milcke, E., Newstetter, W. C., & Davies, J. (2003). Research laboratories as evolving distributed cognitive systems. In R. Alterman & D. Kirsh (Eds.) *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society.*

Research Laboratories as Evolving Distributed Cognitive Systems

Nancy J. Nersessian (nancyn@cc.gatech.edu)
College of Computing, Program in Cognitive Science, Georgia Institute of Technology
Atlanta, GA 30332-0280 USA

Elke Kurz-Milcke (kurzmi@cc.gatech.edu)
College of Computing, Georgia Institute of Technology
Atlanta, GA 30332-0280 USA

Wendy C. Newstetter (wendy@bme.gatech.edu)
Department of Biomedical Engineering, Georgia Institute of Technology
Atlanta, GA 30332-0535 USA

Jim Davies (jim@jimdavies.org)
College of Computing, Georgia Institute of Technology
Atlanta, GA 30332-0280 USA

Abstract

We are carrying out a research project aimed at understanding reasoning and representational practices employed in problem solving in biomedical engineering (BME) laboratories. These laboratories are best construed as *evolving* distributed cognitive systems: the laboratory is not simply a physical space, but a *problem space*, the components of which change over time; cognition is distributed among people and artifacts; and the *cognitive partnerships* between the technological artifacts and the researchers in the system evolve. To investigate this evolving cognitive system we use both ethnography and cognitive-historical analysis. Understanding practices in innovative research laboratories requires in-depth observation of the lab as it presently exists, as well as research into the histories of the experimental devices used in it. We are aiming here for relational accounts (‘*biographies’*) of the distributed cognitive systems within the lab as they change in time. In this we find that one cannot divorce research from learning in the context of the laboratory, where learning involves building relationships with artifacts.

1. Introduction

Science and engineering research laboratories are prime locations for studying the role of environment in cognition. Clearly laboratory practices are located in mico and macro social, cultural, and material environments. So too, these practices employ cognitive representations and processes in developing and using knowledge to solve research problems. We are carrying out a research project aimed at understanding reasoning and representational practices employed in problem solving in biomedical engineering (BME) laboratories. We have begun working in multiple sites, but here report on our research on a specific tissue engineering laboratory, ‘Lab A’, that has as its ultimate objective the eventual development of artificial blood vessels. The daily research is directed towards solving problems that are smaller pieces of that grand objective. The problem solving in these laboratories is best characterized as situated and distributed. These activities are situated in that they lie in localized interactions among humans, and among humans and technological artifacts. They are distributed in that they take place across systems of humans and artifacts.

BME is an *interdiscipline* in which the melding of knowledge and practices from more than one discipline is so extensive that significantly new ways of thinking and working are emerging. Significantly, innovation in technology and lab practices occurs continually, as does learning, development, and change in lab researchers. Thus, we characterize the labs as “evolving distributed cognitive systems”. Investigating and interpreting these kinds of systems requires innovation on the part of cognitive science researchers studying them as well. As an interdisciplinary team of researchers, ourselves, we struggle with myriad methodological and conceptual issues. Two issues we will discuss in this paper are (1) to capture the “evolving” dimension of the labs we have developed a “mixed-method” approach, integrating ethnography and cognitive-historical analysis and (2) to develop the “distributed cognitive system” analysis of the labs we are recasting some useful traditional interpretive notions employed in cognitive science, especially as they relate to the customary internal/external distinction; in particular, ‘problem space’ and ‘mental model’.

2. Integrating Methods to Investigate Evolving Cognitive Systems

In carrying out this case study we have been conducting both cognitive-historical analyses of the problems, objects, and models employed in the research and ethnographic analyses of the day-to-day practices in the lab. Cognitive-historical analysis uses the customary range of historical records to recover how the representational, methodological, and reasoning practices have been developed and used by practitioners in a domain. The practices are examined over time spans of varying length, ranging from shorter spans defined by the activity itself to spans of decades or more. Cognitive-historical analysis aims to interpret and explain the generativity of these practices in light of salient cognitive science investigations and results (Nersessian 1992, 1995). Saliency is determined by the nature of the practices under scrutiny. Cognitive-historical analyses are not historical narratives. Rather, the objective is to enrich understanding of cognition through examining the development of cognitive practices in science and engineering domains.[1] Existing studies have tended to focus on historical individuals, including Faraday, Maxwell, and Bell, and on developing explanatory accounts of concept formation, concept use, and conceptual change (Gooding 1990, Gorman & Carlson, 1994, Nersessian 1992, 1999, Tweney 1985). In our study of BME practices thus far, the cognitive-historical analyses are focused on the technological artifacts that push BME research activity and are shaped and re-shaped by that activity. These artifacts become and remain part of the lab's history. How the members of the lab appropriate the history and employ the artifacts in their daily research is subject to ethnographic analysis.

Ethnographic analysis seeks to uncover the situated activities, tools, and interpretive frameworks utilized in an environment that support the work and the on-going meaning-making of a community. Ethnography of science and engineering practices aims to describe and interpret the relations between observed practices and the social, cultural, and material contexts in which they occur. Ethnographic studies of science and engineering practices abound (See, e.g, Bucciarelli 1994, Latour & Woolgar, 1986, Lynch, 1985). However, studies that focus on situated *cognitive* practices in these areas are few in number. And, existing observational (Dunbar 1995) and ethnographic studies (See, e.g., Goodwin 1995, Hall *et al.*, in press, Ochs & Jacoby 1997) of cognition lack attention to the historical dimension that we find important to our case study.

We find the 'mixed methodology' approach essential to developing an integrated understanding of the reasoning and representational practices in the cognitive systems of the BME lab. Ethnographic analysis allows us to develop traces of transient arrangements of the components of the cognitive system, such as evidenced in laboratory routines, the organization of the workspace, the cultural artifacts in use, and the social organization of the lab at a time, as these unfold in the daily research activities and ground those activities. Cognitive-historical analysis enables us to follow trajectories of the human and technological components of the cognitive system on multiple levels, including their physical shaping and re-shaping in response to problems, their changing contributions to the models that are developed in the lab, and the nature of the concepts that are at play in the research activity at any particular time.

[1] For a comparison of cognitive-historical analysis to other methodologies – laboratory studies, observation, computational modeling – employed in research on scientific discovery, see Klahr & Simon, 1999.

None of the conceptions of distributed cognition in the current literature account for systems the components of which are evolving in time. In studies of distributed cognition in work environments, for instance the cockpit or on board a ship, the problem solving situations change in time. The problems faced, for example, by the pilot change as she is in the process of landing the plane or bringing a ship into the harbor. However, the nature of the technology and the knowledge the pilot and crew bring to bear in those processes is relatively stable. Even though the artifacts have a history within the field, such as Hutchins documents for the instruments aboard a ship, they do not evolve in the day-to-day problem solving processes on board. Thus, the cognitive system is dynamic but largely *synchronic*. The cognitive systems of the BME research laboratory are dynamic and *diachronic* in that, although there loci of stability, during problem solving processes they undergo development and change over time. The technology and the researchers have evolving, relational trajectories that must be factored into the understanding of the cognition at any point in time. It is the cognitive-historical analyses performed on these trajectories that we refer to as *biographies*. BME researchers, assistants, and students have biographies that, in part, become a piece of the lab history; as do the multiple and diverse objects that are manipulated and transformed in the lab. The researchers, for instance, include Post-docs, Ph.D. students, and undergraduates, all of whom have learning trajectories. These trajectories, in turn, intersect with the developmental trajectories of the technological resources within the lab. In this case, the user of the artifacts also re-designs some of them. In order to begin research, a new participant must first master the relevant aspects of the biography of an artifact necessary to the research and then figure ways to alter it to carry out her research project as the new research problems demand.

For example, one highly significant artifact in Lab A is the *flow loop*, a device that emulates the shear stresses experienced by cells within blood vessels. A Ph.D. student we interviewed discussed how the researcher prior to her

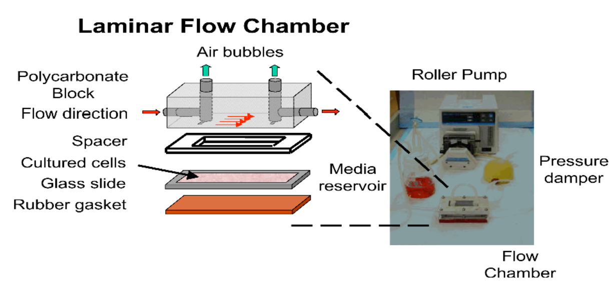

Figure 1: Diagram and photograph of a flow loop

had modified the block to solve some technical problems associated with bacterial contamination - a constant problem in this line of research. The flow loop, as inherited by the new student had previously been used on smooth muscle cells. The new student was planning to use the flow loop to experiment with vascular constructs of endothelial cells that are thicker than the muscle cells, and not flat. To begin that research, she, together with another new student, had to re-engineer the flow loop by changing the width of the flow slit that holds the spacers. Because the vascular constructs are not flat, spacers need to be used between the block and the glass slides in order to improve the flow pattern around the boundary to bring the *in vitro* model more in accord with the *in vivo* model.

Making sense of the day to day cognitive practices in a BME laboratory and bringing biographies of scientific objects to life are prime facie separate tasks. We experience, however, that the research process in this distributed cognitive system evolves at a fast pace, which necessitates going back and forth between the two endeavors. The ethnographic study of the development, understanding, and use of particular devices by various lab members, as well as ethnographic interviews have allowed us to conjoin the cognitive-historical study of biographies of lab members, lab objects, and the lab itself with an eye on the perception of these entities by the lab members themselves. We are using the notion 'biography', in distinction to 'history', to emphasize that at this point we are most interested in the developing relationship between the researchers and their artifacts. We are aiming here for an account of the *lived relation* of the researchers with specific artifacts, rather than for an account of the developing knowledge about these artifacts *per se.*

Setting 'biography' and 'history' apart in this fashion, however, does not suffice for the relational account of distributed cognitive systems that we are attempting to articulate. In fact, the mixed-method approach that we have been utilizing entails the distinction between the two categories as well as recognition of the need for their orchestration in the meaning-making processes occurring around the artifacts in the laboratory. On this account, relating to an artifact entails an appropriation of its history, which chronicles the development of the problem situation including what is known about the artifacts in question. In a way then biography comes to include history.

3. The BME Lab as an Evolving Distributed Cognitive System

3.1. The Lab as Problem Space

The laboratory, as we construe it, is not simply a physical space existing in the present, but rather a *problem space*, constrained by the research program of the lab director, that is reconfiguring itself almost continually as the research program moves along and takes new directions in response to what occurs both in the lab and in the wider community of which the research is a part. Researchers and artifacts move back and forth between the wider community and the physical space of the lab. Thus the problem space has permeable boundaries.

For instance, among the most notable and recent artifacts (initiated in 1996) in Lab A are the tubular-shaped, bio-engineered cell-seeded vascular grafts, locally called ‘constructs’. These are models of native blood vessels the lab hopes to engineer, eventually, as viable implants for the human vascular system. The endothelial cells the lab uses in seeding constructs are obtained by researchers going to a distant medical school and bringing them back into the problem space of the lab. Occasionally, the constructs or substrates of constructs travel with lab members to places outside of the lab. For instance, one of the graduate students takes substrates of constructs to a laboratory at a nearby medical school that has the elaborate instrumentation to perform certain kinds of genetic analysis (microarrays). This line of research is dependent on resources that are currently only available outside Lab A, here in the literal, spatial sense. The information produced in this locale is brought back into the problem space of the lab by the researcher.

At any point in time the lab-as-problem-space contains resources for problem solving which comprise people, technology, techniques, knowledge resources (e.g. articles, books, the internet), problems, and relationships. Construed in this way, the notion of ‘problem space’ takes on an expanded meaning from that customarily employed by the traditional cognitive science characterization of problem solving as search through an internally represented problem space. Most importantly, in our characterization, the problem space comprises models and artifacts together with a repertoire of activities in which simulative model-based reasoning assumes a central place (Nersessian 1999).

Following Hutchins (1995), we analyze the cognitive processes implicated in a problem-solving episode as residing in the *cognitive systems* comprising both one or more researcher and the *cognitive artifacts* involved in the episode. In line with his analysis, ‘cognitive systems’ are understood to be socio-technical in nature and ‘cognitive artifacts’ are material media possessing the cognitive properties of generating, manipulating, or propagating representations. In contrast to the systems he as studied, though, ours are communities focused on innovation.

Determining the cognitive artifacts within any cognitive system involves issues of agency and intention that are pressing questions for cognitive science research, both in the development of the theoretical foundations of distributed cognition and in relation to a specific case study. To better understand such issues we have been focusing on the technology employed in experimentation. During a research meeting with the lab members, including the PI, we asked them to sort the material artifacts in the lab according to categories of their own devising and rank the importance of the various pieces to their research. Their classification is represented below (Table 1).

Additional ethnographic observations have led us to formulate working definitions of the categories employed by Lab A’s researchers. ‘Devices’ are engineered facsimiles

Nersessian, N. J., Kurz-Milcke, E., Newstetter, W. C., & Davies, J. (2003). Research laboratories as evolving distributed cognitive systems. In R. Alterman & D. Kirsh (Eds.) *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society.*

that serve as *in vitro* models and sites of simulation.[2] 'Instruments' generate measured output in visual, quantitative, or graphical form. 'Equipment' assists with manual or mental labor. Much to the surprise of the PI, the

Table 1: Sorting of lab artifacts by the lab members

ONTOLOGY OF ARTIFACTS

DEVICES	INSTRUMENTS	EQUIPMENT
flow loop	confocal	pipette
bioreactor	flow cytometer	flask
bi-axial strain	coulter counter	refrigerator
construct	"beauty and beast"	sterile hood
	mechanical tester	water bath
	LM5 (program)	camera
	computer	

newer Ph.D. students initially wanted to rank some of the equipment, such as the pipette, as the most important to their research; whereas for him the devices the lab engineers for simulation purposes are the most important.

The cognitive artifacts in the distributed systems in the lab cut across these distinctions, though most are devices or instruments. Analysis of the ethnographic data has focused our attention on the devices, all of which we classify as cognitive artifacts. Devices serve as sites of experimentation with cells and constructs under conditions simulating those found in the vascular systems of organisms. It is in relation to the researcher(s)'s intent of performing a simulation with the device in order to create new situations that parallel potential real-world situations, and the activity of the device in so doing, that qualify a device as a cognitive artifact within the system that employs it. For example, as a device, the flow loop *represent*s blood flow in the artery. In the process of simulating, it *manipulates* constructs which are representations of blood vessel walls. After being manipulated, the constructs are then removed and examined with the aid of instruments, such as the confocal microscope, which generates images for many color channels, at multiple locations, magnifications, and gains. Thus, the manipulated representation of the flow loop is *propagated* within the cognitive system.

[2] We are using the term 'device' because this is how the researchers in the lab categorized the *in vitro* simulation technology. This notion differs from the notion of "inscription devices" that Latour & Woolgar (1987, p. 51) introduced and that has been discussed widely in the science studies literature. The latter are devices for literally creating figures or diagrams of phenomena. The former are sites of *in vitro* simulation, and further processing with instruments is necessary to transform the information provided by these devices into visual representations or quantitative measures.

3.2 Distributed Mental Modeling[3]

An *in vivo/in vitro* division is a significant component of the cognitive framework guiding practice in Lab A. Because the test bed environment for developing artificial blood vessels cannot be the human body in which they will ultimately be implanted, the BME researchers have to design facsimiles of the *in vivo* environment where the experiments can occur. These devices provide locally constructed sites of experimentation where *in vitro* models are used to screen and control specific aspects of the *in vivo* phenomena they want to examine. The researchers in the lab call the process of constructing and manipulating these *in vitro* sites "putting a thought into the bench top and seeing whether it works or not." These instantiated "thoughts" allow researchers to perform controlled simulations of an *in vivo* context, for example, of the local forces at work in the artery. The "bench top", as one researcher explained, is not the flat table surface but comprises all the locales where experimentation takes place.

Devices perform as models instantiating current understanding of properties and behaviors of biological systems. For example, the flow loop is constructed so that the behavior of the fluid is such as to create the kinds of mechanical stresses in the vascular system. But devices are systems themselves, possessing engineering constraints that often require simplification and idealization in instantiating the biological system they are modeling. For example, the bioreactor, as used in Lab A, was designed for "mimicking the wall motions of the natural artery." It is used to expose the constructs to mechanical loads in order to improve their overall mechanical properties. The researchers call this process "mechanical conditioning" or as one researcher put it "exercising the cells." Preferably, this is done at an early stage of the formation of the construct, shortly after seeding the cells onto a prepared tubular silicon sleeve. *In vivo*, arterial wall motion is conditioned upon pulsatile blood flow. With the bioreactor, though, which consists of a rectangular reservoir containing a fluid medium (blood-mimicking fluid) in which the tubular constructs are immersed and connected to inlet and outlet ports off the walls of the reservoir, "fluid doesn't actually move," as one

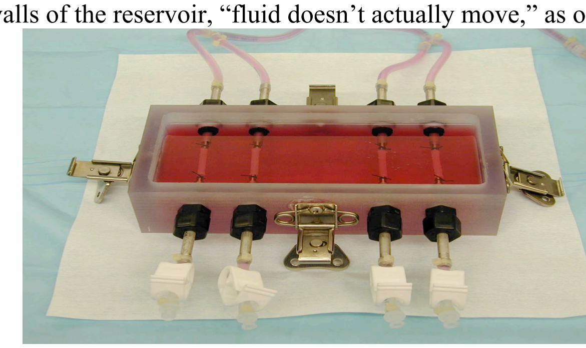

Figure 2: Photograph of a bioreactor

[3] Of course, we use the term 'mental' metaphorically here, as a rhetorical move to connect our discussion with aspects of the traditional notion of mental modeling and extend the notion for use in the distributed cognition framework.

lab member put it, "which is somewhat different from the actual, uh, you know, real life situation that flows." The bioreactor is a functional model of pulsatile blood flow. The sleeves are inflated with pressurized culture medium, under pneumatic control (produced by an air pump). The medium functions as an incompressible fluid, similar to blood. By pressurizing the medium within the sleeves, the diameter of the silicon sleeve is changed, producing strain on the cells, similar to that experienced *in vivo*.

Many instances of model-based reasoning in science and engineering employ 'external' representations that are constructed during the reasoning process, such as diagrams, sketches, and physical models. These can be viewed as providing constraints and affordances essential to problem solving that augment those available in whatever 'internal' representations are used by the reasoner during the process. A device is a kind of physical model employed in the problem solving in Lab A. Within the cognitive systems in the lab, devices instantiate part of the current model of the phenomena and allow simulation and manipulation. The intent of the simulation is to create new situations *in vitro* that parallel potential *in vivo* situations.

In previous research Nersessian (2002) characterized the reasoning involved in simulative model-based reasoning as a form of dynamic mental modeling employing iconic representations. That analysis used the traditional notion of mental modeling as an internal thought process. Here we expand the notion of simulation of a mental model to comprise both what are customarily held to be the internal thought processes of the human agent and the processing of the external device. Simulative model-based reasoning involves a process of co-constructing the 'internal' researcher models of the phenomena and of the device and the 'external' model that is the device, each incomplete. Understood in this way, simulating the mental model would consist of processing information both in memory and in the environment (See Greeno 1989 for a similar view). That is, the mental modeling process is distributed in the cognitive system.

3.3 Evolving Cognitive Partnerships

Devices, such as the construct, the flow loop, and the bioreactor are constructed and modified in the course of research with respect to problems encountered and changes in understanding. These devices have a history within the research of the lab. For example, the flow loop was first created in this particular lab to simulate "known fluid mechanically imposed wall sheer stress," in other words to perform as a model of hemodynamics. We have traced aspects of its development since 1985. The constructs were first devised in 1996 as an important step in the overall objective of creating implantable vascular substitutes. They afford experimentation not only on cells, but on structures more closely related to the *in vivo* model. The bioreactor, though having a longer and more varied history outside the lab, first made its appearance in this lab in conjunction with the tubular constructs and was not used anywhere before for that purpose.

Newcomers to the lab, who are seeking to find their place in the evolving system, initially encounter these devices as materially circumscribed objects. As they begin to interact with these devices, the newcomers, almost as a rule, experience them as fraught with intricacies that withstand their easy handling: Tubes leak, sutures don't keep, reservoirs overflow, pumps malfunction, the available spacers don't fit, and, as if this were not enough, cells "go bad." On the other hand, newcomers find themselves in an environment in which everybody else, including the most experienced old-timer, is constantly struggling to get things to work—a serious fact about laboratory life (not always handled with dead seriousness in this environment).

Growing cognitive membership in the lab involves a gradual process of coming to understand these objects as devices - as objects with evolving trajectories, constructed and employed to respond to problems, to help answer questions, and to generate new ones. Thus, we find that one cannot divorce research from learning in the context of the laboratory, and learning involves building relationships with artifacts. We characterize the relationships between the technological artifacts in the cognitive system and the researchers as *cognitive partnerships*. Over time understandings are constructed, revised, enhanced, and transformed through partnerships with the artifacts in the community. As relationships change, so too do knowledge and participation. It is because these lived relationships play out in time that we have dubbed our analyses of them 'biographies'.

Learning is central in the partnering process. One new undergraduate research scholar from mechanical engineering, who was about to use the *mechanical tester* – a device used to test the strength of the constructs – responded:

> A2: ….I know that we are pulling little slices of the construct – they are round, we are just pulling them. It's the machine that is right there before the computer in the lab. The one that has the big "DO NOT TOUCH" on it.
> I: Is it the axial strain (mechanical tester)?
> A2: I know it has a hook in it and pulls

That the novice researcher describes the mechanical tester at this point in time as nothing more than parts suggests that she has yet to partner with the device. In contrast, a senior Ph.D. researcher, at that point in time considered the "resident expert" on the bioreactor, was able easily to discuss the artifact and some of its history:

> I: Do you sometimes go back and make modifications? Does that mean you have some generations of this?
> A 12: Uh yes I do. The first generation and the second generation or an offshoot I guess of the first generation. Well the first one I made was to do mechanical loading and perfusion. And then we realized that perfusion was a much more intricate problem than we had - or interesting thing to look at - than we had guessed. And so we decided okay we will make a bioreactor that just does perfusion on a smaller scale, doesn't take up much space, can be used more easily, can have a larger number of replicates, and so I came up with this idea.

He continued by pulling down previous versions of bioreactor (made by earlier researchers as well) and explaining the modifications and problems for which design changes were made. His account suggests a developed partnership.

The cognitive partnerships transform both researcher and artifact. A researcher who some months earlier was a newcomer and who saw the artifacts as just many kinds of machines and objects piled on shelves and on the bench top, now can see a device as an *in vitro* site for "putting a thought [his thought] into the bench top and seeing whether it works or not." During the problem-solving processes involved in instantiating a thought and seeing if it works, devices are re-engineered as exemplified with the flow loop in Section 2. And, potential device transformations can be envisioned, as with one undergraduate research scholar we interviewed about the bioreactor:

> A16: ..I wish we could accomplish - would be to actually suture the actual construct in there somehow. To find a way not to use the silicon sleeve….That would really be neat. Um, simply because the silicon sleeves add the next level of doubt. They're – they are a variable thing that we use, they're not always 100% consistent. Um the construct itself is not actually seeing the pressure that the sleeve does. And because of that you know, it doesn't actually see a – a pressure. It feels the distention but it doesn't really feel the pressure. It doesn't have to withstand the pressure. That's the whole idea of the sleeve. And so, um, I think that it would provide a little bit more realism to it. And uh, because that also, a surgeon would actually want to suture the construct into a patient. And um, because of that you're also mimicking the patient as well - if you actually have the construct in the path. I think another thing is to actually have the flow because um, so this flow wouldn't be important with just the sleeve in there. But if you had the construct in contact with the – with the liquid that's on the inside, you could actually start to flow media through there.

In this case an undergraduate student has been transformed over the course of several semesters to a BME researcher, contributing to immediate research goals; who transforms artifacts in his immediate research; who understands the outstanding problems and objectives; and who can envision how a device might change from a functional model to a model more closely paralleling the *in vivo* situation to push the research along. At this point in evolution, thinking is taking place through the cognitive partnering of the researcher and device. The partnership provides the means for generating, manipulating, and propagating representations within the distributed cognitive systems of this research laboratory.

4. Conclusions

A relational account of distributed cognitive systems characterizes cognition in terms of the lived relationships between the "players" in these systems, people and artifacts. In the research laboratories that we have been studying, these relationships evolve in significant ways for the individual lab members and for the community as a whole. In their established form, relationships with artifacts entail cognitive partnerships that live in interlocking models performing internally, as well as externally.

Acknowledgments

We thank Kareen Malone and Eteinne Pelaprat for their contributions to the research on this project. We thank our research subjects for allowing us into their work environment and granting us numerous interviews. We gratefully acknowledge the support of the National Science Foundation ROLE Grant REC0106773.

References

Bucciarelli, L. L. (1994). *Designing Engineers*. Cambridge, MA: MIT Press.

Dunbar, K. (1995). How scientists really reason: Scientific reasoning in real-world laboratories. In R. J. Sternberg & J. E. Davidson (Eds.), *The Nature of Insight*. Cambridge, MA: MIT Press.

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

Gooding, D. (1990). *Experiment and the Making of Meaning: Human Agency in Scientific Observation and Experiment*. Dordrecht: Kluwer.

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

Greeno, J. G. (1989). Situations, mental models, and generative knowledge. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing*. Hillsdale, NJ: Lawrence Erlbaum.

Hall, R., Stevens, R., & Torralba, T. (in press). Disrupting representational infrastructure in conversations across disciplines. *Mind, Culture, and Activity.*

Klahr, D. & Simon, H. A. (1999). Studies of Scientific Discovery: Complementary Approaches and Convergent Findings. *Psychological Bulletin*, 125 (5), 524-543.

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

Latour, B., & Woolgar, S. (1986). *Laboratory Life: The Construction of Scientific Facts*. Princeton: Princeton University Press.

Lynch, M. (1985). *Art and Artifact in Laboratory Science: A Study of Shop Work and Shop Talk in a Research Laboratory*. London: Routledge and Kegan Paul.

Nersessian, N. J. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In R. Giere (Ed.), *Minnesota Studies in the Philosophy of Science*). Minneapolis: University of Minnesota Press.

Nersessian, N. J. (1995). Opening the black box: Cognitive science and the history of science. *Osiris, 10 (*Constructing Knowledge in the History of Science, A. Thackray, ed.), 194-211.

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

Nersessian, N. J. (2002). The cognitive basis of model-based reasoning in science. In P. Carruthers, S. Stich, & M. Siegal, eds., *The Cognitive Basis of Science*. Cambridge: Cambridge University Press.

Ochs, E. & Jacoby, S. (1997). Down to the wire: The cultural clock of physicists and the discourse of consensus. *Language in Society, 26*, 479-505.

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