

Engineering Concepts

In Mind, Culture, and Activity 19, 2012 - on line May 2012

Engineering Concepts:

The interplay between concept formation and modeling practices in bio-engineering sciences

Nancy J. Nersessian

Georgia Institute of Technology

Abstract

This paper addresses “concept formation in the wild” through examining the relations between concept formation and physical and computational simulation modeling practices in two research laboratories in the bioengineering sciences. It argues that processes of concept formation and of building distributed cognitive systems are deeply entwined.

Introduction

As much research has demonstrated, novel scientific concepts do not arise fully formed in the head of a scientist but are created in problem-solving processes, which can extend for considerable periods and even span generations of scientists. To understand concept formation and change it is important to investigate these processes in the contexts of scientific practice, that is, “in the wild.” Historical records provide some insights into practice and I have argued that these records of concept formation provide significant evidence that novel scientific concepts arise from the interplay of attempts to solve specific problems, use of conceptual, material, and analytical resources provided by the problem situation, and often through the use of *model-based reasoning* (Nersessian, 1984, 2002, 2008).¹ In particular, the historical records provide evidence that conceptual models, developed in conjunction with external representations (especially graphic and mathematical), have been instrumental in problem solving leading to conceptual innovation and change. Ethnographic studies of problem-solving practices in research laboratories in the bio-engineering sciences conducted by my research group over the last ten years establish that physical and computational modeling practices also participate in concept formation. Together, these investigations into the modeling practices of scientists provide significant insight not just into how they *use* concepts and models, but into how scientists “create their cognitive powers by creating the environments in which they exercise those powers”

¹ Recently I have discovered that in characterizing concept formation in learning processes, Vygotsky claimed that “a concept emerges and takes shape in the course of a complex operation aimed at the solution of some problem” (Vygotsky, 1962). Further, his general notion of concept formation is that it is an on-going dynamic and socio-cultural process in each use or acquisition of a concept. Concepts are neither completely fixed units of representation nor solely mental representations, but arise, develop, and live in the interactions among people as they create and use them

((Hutchins, 1995), xvi). Designing and constructing conceptual, physical, and computational models are one way that scientists build thinking environments, that is, how they *build cognition* (Chandrasekharan & Nersessian, 2011). As our investigation of research labs underscores, “the lab” creates itself as a distributed cognitive system with specific affordances and constraints for problem solving as it creates knowledge. In this paper I focus on the interplay between the formation and articulation of concepts and the investigative practices of bio-engineering researchers in areas that require the indirect investigation of in vivo biological phenomena by means of physical and computational models. Although these models are intended to serve as sources from which hypotheses can be transferred to the in vivo phenomena, in much of this research the models themselves are objects of interest. That is research is directed towards building models that exemplify selected features of the in vivo phenomena, experimenting with these controlled environments, and trying to make sense of the phenomena manifested in them. Thus concept formation is situated in the context of modeling, where the need to articulate concepts drives the creation of novel models and the need to understand and interpret model phenomena drives concept formation.²

The modeling practices examined in this paper are drawn from ethnographic studies of two research laboratories in the biomedical engineering sciences (BME): tissue engineering and neural engineering. These labs share the practice of designing and constructing in vitro physical

² There is a large literature in science studies on social and cultural aspects of scientific modeling, and there is an equally large philosophical literature that focuses on the representational nature of models. However, little research has been directed towards the intellectual work performed by modeling (exceptions include (Boon & Knuuttila, 2009; Morgan & Morrison, 1999)) in general and, specifically, towards the interplay of modeling and concept formation. Although not focused on concept formation, Giere (Giere, 2002) also construes models are part of distributed cognitive systems.

simulation models that serve as sites of experimentation on selected aspects of in vivo biological phenomena. In effect, these researchers build “*parallel*” worlds in which models enact specific aspects of phenomena under controlled conditions that they cannot investigate directly due either to ethical or control issues. The labs are populated by *engineering scientists* who are a breed of researcher whose aim is to make fundamental contributions to basic science as well as to create novel artifacts and technologies for applications (medical in BME). For engineers, generally, “to engineer” means, in general, to means to conceive, design, and build artifacts in iterative processes. In the cases considered here, this practice extends to concepts. In this paper I examine the interaction between engineering simulation models and engineering concepts.

Both of the BME labs we studied were carrying out groundbreaking research that required developing new investigative practices. The primary means of investigations in both labs is through designing, building, and carrying out experiments with bio-engineered physical simulation models they call “*devices*.” These devices participate in experimental research in various configurations of hybrid “*model-systems*.” As one researcher commented, they “*use that [notion] as the integrated nature, the biological aspect coming together with the engineering aspect, so it’s a multifaceted model-system.*” Simulation models are designed to function as analogical sources for inference and prediction about target in vivo systems. They are constructed so as to enable the researcher “*to predict what is going to happen in a system [in vivo]. Like people use mathematical models... to predict what is going to happen in a mechanical system? Well, this [model-system she was designing] is an experimental model that predicts – or at least*

you hope it predicts – what will happen in real life.” That is, research is conducted with these in vitro devices and outcomes are transferred as candidate understandings and hypotheses to the in vivo phenomena. In this paper I will examine how in processes of problem solving by means of model-systems, researchers in both labs in fact engineered novel concepts.

These frontier areas of bio-engineering are good candidates for investigating practices implicated in concept formation because 1) the research requires interdisciplinary synthesis and 2) the phenomena they investigate are so novel that they have only been conceptualized partially or not at all. Thus such areas are likely sources for cases of concept transfer and modification and, potentially, the formation of scientifically novel concepts. During the period of our investigation, these labs actually did demonstrate and confirm the centrality of physical (and, unexpectedly, computational) modeling to concept formation. Further, they provided insight into how processes of articulating concepts can generate novel investigative practices, and configure and reconfigure “the lab.” I briefly discuss two cases, one from a tissue engineering lab (Lab A) which provides insight into how the attempt to articulate a rudimentary conceptualization can in turn articulate the lab itself an evolving distributed cognitive systems comprising researchers and model-systems; the other from the neural engineering lab (Lab D) where the cross-breeding of two models – one physical and one computational – built a cognitive system that created a cluster of scientifically novel (and potentially highly significant) conceptual innovations.

The Investigations

Methods

Our research group conducted an ethnographic study that sought to uncover the activities,

artifacts, and meaning-making that constitute research as it is situated in the on-going practices of each community. For each lab we conducted 2 years of intensive data collection, followed by 2 years of targeted follow-up and, thereafter, limited tracking of students through to their graduation. We took field notes on our observations, audio taped largely unguided interviews, and video and audio taped research meeting . As a group (6 ethnographers) we estimate over 800 hours of field observations. This “team ethnography” approach supplanted videotaping research activities in the labs. Although we were allowed to video research meetings, it proved not feasible to video research in the laboratories. Researchers expressed discomfort at being videotaped in the labs as they worked. Additionally, the kinds of observations we were after were not amenable to requesting to be notified so that we could film when "X" was happening. We used interpretive coding in analyzing interviews and field notes. Broadly consistent with the aims of grounded theory(Glaser & Strauss, 1967; Strauss & Corbin, 1998), we approached coding so as to enable core categories and interpretations to emerge from the data and remain grounded in it, while at the same time being guided by our initial research questions.

Additionally, these labs are evolving systems that reconfigure as the research program moves along and takes new directions in response to what occurs both in the lab and the broader community of which the lab is a part. The signature technologies of the labs are designed and redesigned in the context of research problems and projects, new methods are developed or adopted, and at any slice in time the lab is populated largely by students at various points in their development into full-fledged researchers. To capture the historical dimension of these lab communities we use also interpretive methods of cognitive-historical analysis (Nersessian, 1995,

2008). In this investigation cognitive-historical analysis aims at understanding the integrated cognitive - social - cultural nature of the research practices. Data collection included a range of customary historical records.

Case 1: Articulating ‘Arterial Shear’

Lab A dates from 1987, when the director moved to a new university to take advantage of the opportunity to begin research on tissue engineering. During our investigation the main members included the director, one laboratory manager, one postdoctoral researcher, seven PhD graduate students (three graduated while we were there, the other four, after we concluded formal data collection), two MS graduate students, four long-term undergraduates. The laboratory director's background is in aeronautical engineering, but he was by then a senior, highly renowned pioneer in the field of biomedical engineering and the emerging subfield of tissue engineering. All of the researchers came from engineering backgrounds, mainly mechanical or chemical engineering.

During a period of research (starting in the mid-1960's) for the space program on the effects of vibration along the axis of the Saturn launch vehicle (“pogo stick vibration”) on the cardiovascular system of astronauts, the now director of Lab A developed the hypothesis that the physical forces to which blood vessels are naturally exposed, such as pressure and forces associated with blood flow through the arteries, could adversely affect the endothelial cells that line the vessels and, thus, be implicated in disease processes such as atherosclerosis. He embarked on a program of research into how and under what conditions the physical forces might create disease. Along with a handful of medical researchers, he formulated a rudimentary concept,

‘arterial shear’: the frictional force of blood flow parallel to the plane through the lumen of an artery.³ At the outset this preliminary understanding functioned as a “germ cell” (Engestrom, 1999) or “placeholder” concept (Carey, 2009) – one that captures a rudimentary idea but is in need of articulation, which in this case is on-going forty years later. What the research community now understands to be among the features of arterial shear are: it regulates endothelial cell migration, morphology, and proliferation; laminar flow is needed for these functions; and turbulent flow creates changes that promote vascular constriction and platelet aggregation. During the time Lab A has been in existence, its research has expanded from investigating macroscopic features, such as morphology, to the functioning of the cells themselves, such as in the expression of anti-coagulant proteins.

The future Lab A director decided early in this research to focus his efforts on the endothelial cells that line the arteries: “*It made sense to me that if there was this influence of flow on the underlying biology of the vessel wall, that somehow the cell type had to be involved, the endothelium.*” In the 1970's vascular biologists were focused on biochemical processes and those he contacted were skeptical of the hypothesis and that investigating arterial shear would lead to anything worthwhile. So, he ended up in the laboratory of a veterinary physiologist where they surgically created *animal models* to induce pathologies in native arteries and investigated the

³ It is unclear when and where this term came first into use, but the Lab A director did find some medical researchers to collaborate with at that time, and so the germ cell concept arose within a small community. For our purposes what is more important is how concept formation processes drove the setup of Lab A and the line of in vitro research we witnessed during our investigation. In the cases discussed in this paper, a close connection will be evident between concepts and what might be considered the developing theories of the phenomena. This is in line with philosophical accounts that advocate understanding a concept to be represented by the set of interrelated features or characteristics ascribed to it by a theory, or by the scientists who are using it when no developed theory exists (see, e.g. (Arabatzi, 2006; Nersessian, 1984)). It is also in line with cognitive science accounts, most notably the “dynamic frames” analysis of (Barsalou, 1992) It has been applied to scientific concepts by (Andersen, 2006)

nature and effects of arterial shear. Problems of control were significant in these modeling practices and led him to set on a program of in vitro research which would require developing physical models through which to simulate biological phenomena with the desired experimental control. Thus, nearly all problem-solving activities in Lab A are model-based, that is they require building (designing and constructing) physical models, assembling them in various model-system configurations, and performing simulations under various controlled experimental conditions.

The initial configuration of Lab A centered around one physical model, the *flow channel device* (“*flow loop*”), that is designed to enact selected in vivo blood flow conditions, normal and pathological. It consists of a flow channel (designed in a physiologically meaningful range) with accompanying flow inducing components (pump, pulse dampener, liquid the viscosity of blood) designed to represent to a 1st order approximation shear stresses that can occur during blood flow in an artery. When cells mounted on slides are “*flowed*” under different conditions, changes in cell morphology can be related directly to the controlled wall shear stresses. This device started its life as a large, cumbersome artifact on a stand, for which contamination was a constant problem since it could not be assembled under the sterile workbench hood. Within a few years it was re-engineered into a compact design that fits under the hood and experiments can be run in an incubator.

After several years of research, the Lab sought a better model-system. Using cell cultures on slides can provide only limited understanding of arterial shear stress. Specifically, as the director noted, “*putting cells in plastic and exposing them to flow is not a very good simulation*”

of what is actually happening in the body.... If you look within the vessel wall you have smooth muscle cells and then inside the lining is [sic]the endothelial cells, but these cells types communicate with one another. So we had an idea: let's try to tissue engineer a better model-system for using cell cultures." The idea was to create "a more physiological model" where the effects of shear could be studied on more components of the blood vessel wall than the endothelial cells in isolation, and help to better understand the functional properties related to arterial shear. To expand the possibilities for studying these properties, the director took "the big gamble" – to "tissue engineer" a model of the blood vessel wall constructed from living tissue, now called the "construct" or the "tissue engineered blood vessel wall model" or, underscoring its application potential, the "tissue engineered vascular graft" within the lab.

The in vivo blood vessel comprises several layers and the in vitro construct can be built into a range of models with different levels of approximation for simulating in vivo processes. The construct, in turn, has given rise to several other physical models through which to manipulate and examine construct properties under various conditions and also to a novel applied goal: to tissue engineer a viable replacement blood vessel for human implantation. To be either a functional model or an implant requires among other things that the cells that are embedded in the scaffolding material must replicate the capabilities and behaviors of in vivo cells so that higher level tissue functions, such as expressing the right proteins and genetic markers, can be achieved. These problems, in turn, opened new lines of lab research and led to building new physical models.

When we entered the lab, the construct, was the focal point of the interconnected research

problems directed towards both what the director called the lab's "*basic biology*" research aimed at continuing to articulate 'arterial shear' and the new application goal of creating a viable vascular implant. Early on, we asked the director to draw a picture of the current lab research (Figure 1). His intention was to depict how his research problems ("barriers" listed at top), researchers, and technologies (listed at bottom) are interconnected. A diagram is a static representation, but for the purposes of my analysis, the director's representation provides a schematic of what I have been calling "the lab as an evolving distributed cognitive system" – a dynamic constellation of problems, researchers, simulation models, and other technologies, reconfiguring as the research moves along. Although the director did not include himself on the diagram, he is, of course, an integral part of the cognitive system. I will outline, briefly, how the system formed and evolved in the course of addressing the research problems by designing and building model-based simulations.

[insert Figure 1 here]

At the top of the diagram the director categorized the major problems ("barriers") the research was addressing. "Mechanical properties and the influence of mechanical forces" are problems directed towards understanding the nature of arterial shear and its role in normal and disease processes and towards creating an implant with the requisite mechanical biological properties to function within the body. "Cell source strategies" research addresses the problem

that endothelial cells are among the most immune sensitive in the body and a viable implant would need to be seeded with cells that would not be rejected by the individual's body. This problem opened lines of research for all the lab members, for example, into stem cell differentiation (A8) and progenitor cells (A7). The lab-built simulation models are designated by "collagen gel technology" (construct), "flow chamber studies" (flow loop), and "mechanical testing" (equibiaxial strain device and a mechanical strength tester). The kinds of investigations along the bottom of the diagram implicate the lab-built simulation models and the technologies through which simulation outcomes are examined. For instance, after a *flow chamber* study in which the *construct* would be subjected to controlled shear stresses, the effects on the endothelial cells can be examined for various biological markers or gene profiling, which implicate a range of technologies such as the confocal microscope to study morphology and migration and DNA microarray technology for gene expression.

The director intended that thick lines to denote the interconnections among the individual research projects with respect to the major problems of the lab. A8, a postdoctoral researcher, is represented as not connected to the students because she started a new line of lab research into the possibility of stem cell differentiation by means of mechanical forces that only later has become integral to the research. She did interact with other lab members about her and their research during daily the lab activities and at the research meeting sessions. Each student had an individual research project, and the problem-solving processes associated with it could be explicated as performed by a distributed cognitive system (see Nersessian 2009 for an analysis of

A7's research). But the lab's dual problems of articulating 'arterial shear' and creating a viable implant built the lab itself into an evolving distributed cognitive system that afforded and constrained the student research.

All the lab members noted on the diagram entered the lab within the same year, with the exception of the MS student A22, who entered with the following academic year. When they entered the lab, the flow loop model was a well-established technology of research, but several of their research problems required some re-design of it. The construct model was a recent development and they played significant roles in furthering its design, in directions related to their specific projects. A5's research was directed towards correlating the development of arteriosclerosis with the genetic behavior of the endothelial cells and progenitor endothelial cells that circulate in the blood stream by simulating various flow conditions with both model-systems. A10's research was investigating the effects of shear stress on aortic valve function, using valvular endothelial cells and also a novel aortic construct model that he designed. Mechanical integrity and strength were primary concerns for him, and, although he ended up not using it himself, he designed and built a new model for the lab, an *equibiaxial strain device* that would simulate the strain (deformation from stress) experienced by vessels as blood flows through them. A4's research was examining specific biological markers in relation to controlled mechanical stimulation of constructs, as compared with their behavior in native tissue. A22's MS research was focusing on improving the mechanical strength of constructs.

As indicated by the thick lines on the lab diagram, all of the systems components are

connected to A7, who, in an early interview with us, noted that she had been designated as “*the person who would take the construct in vivo,*” by which she meant she would need to create a model-system in which a construct would be connected to a living animal. To be successful, the project would need to “*obviously integrate the results of colleagues here in the lab.*” Her research project evolved, specifically, into investigating whether shear stress conditioning of endothelial progenitor cells with the flow loop would make them function as mature endothelial cells in the production of thrombomodulin (a protein that prevents platelet formation) when attached to an animal circulatory system. To address this problem she designed a model-system that could connect the construct to the blood stream of a baboon by means of an exterior shunt between the femoral artery and vein of the animal. Designing and running this model-system with the requisite experimental controls was the most complex problem undertaken by the lab to date. As A7 noted, “*in the lab we can control... exactly what the flow is like... But when we move to an animal model, it’s more physiologic—the challenge then is that is a much more complex system.*” Despite the complexity, she was able to determine that flow loop shear at the normal human in vivo rate of arterial shear (15 dynes/cm²) enhances the ability of progenitor cells to express anti-coagulant proteins. This finding made significant contribution to both the research community’s understanding of arterial shear and the problem of cell sources for implantation.

Discussion. This case illustrates how problem-solving efforts directed towards concept formation can build an evolving distributed cognitive system, which creates the possibilities for problem solving. The studies that first led to the development of the notion of distributed

cognition were of highly structured task environments (plane cockpit, naval ship) in which people created their cognitive powers by making use of existing representational artifacts (Hutchins 1992, 1995). These are dynamic problem-solving environments, but the artifact components of the system are relatively stable. The research labs we have been studying are ill-structured problem-solving environments in which people design and build representational artifacts that in turn serve to articulate the nature of the distributed cognitive system itself. The need to develop concepts for understanding novel phenomena drives practices that create the problem-solving systems from which the concepts emerge and develop. The comparison is akin to that between flying a plane and building the plane while it's flying – and with only a vague idea of what a flying vehicle might look like.

At the outset, the lab director did not envision, for instance, his lab tissue-engineering vascular construct models or conducting stem cell research and gene profiling. The “germ cell” concept of *arterial shear* stemmed from early mathematical modeling of vibratory forces on the human vascular system and experimental modeling with animals. The problems associated with further articulating the concept led to taking the research in vitro, which afforded more control and opened the possibility to examine selected features of arterial shear in relation to the endothelial cells that had been under investigation with the animals. Building the flow loop model enabled them to focus largely on structural properties and proliferation behavior of cells under shear. Devising the construct family of models provided not only a range of more physiologically accurate models, but also the possibility of investigating the functional properties

of blood vessels in relation to shear, such as gene regulation, and led to the lab's ability to create a completely different kind of animal model-system from those of the director's initial research. In sum, over time building physical simulation models articulated a fabric of interlocking researchers, models, and technologies for visualization and analysis that constitutes Lab A, within which the formation of the concept of arterial shear is on-going.

The cognitive powers created by constructing physical simulation models are much the same as in my earlier account of model-based reasoning by means of conceptual modeling (Nersessian 2008), including abstractive ability, integration of knowledge and constraints from diverse domains, conceptualization, and changes in representational format affording analogical, visual, and simulative reasoning. What is novel here is that investigating concept formation in the wild enables us to discern facets of how these powers are created as the material and socio-cultural environment is being built.

Case 2: The Emergence of 'CAT'

The second case I consider provides an exemplar of when novel concepts emerge as a distributed cognitive system takes shape and evolves its modeling practices. As an institution, the neural engineering laboratory was in existence for a few months and still very much in the process of forming when we entered. During our study the main members included a director, one laboratory manager, four PhD graduate students, one MS student, and six long-term undergraduates. When we began, the laboratory director was a new tenure-track assistant professor, fresh from a postdoc in a biophysics laboratory that develops techniques and

technologies for studying in vitro network cultures of neurons. He already had achieved some recognition as a pioneer. His background is in chemistry and biochemistry, with his engineering experience largely self-taught, though highly sophisticated. The PhD students engaged in the research discussed here had backgrounds in mechanical engineering (D2), electrical engineering (D4), and a joint degree in life sciences and chemistry (D11)..

During the period of our investigation, Lab D's overarching research problems were to understand the mechanisms through which living networks of neurons learn and, potentially, to use this knowledge to develop aids for neurological deficits and diseases. At the time the Lab D director began research, the major approach to neuron learning was through single neuron studies on living animals, where one might gauge responses to stimuli, but not control supervised learning. Parallel to the Lab A case, a problem of conceptual articulation also played a role in the formation of Lab D as an evolving distributed cognitive system. In this case, the problem was how to conceptualize 'learning' at the neural level as a network phenomenon. The director believed that to study learning there needed to be a way to investigate the network properties of neurons, since it is networks that learn in the brain. While in graduate school, working on a completely unrelated project, he had the idea that "*perhaps you can make cell culture systems that learn.*" Such a culture would more closely model learning in the brain and also enable emergent properties to arise and be studied. His postdoctoral research was directed towards developing an in vitro physical model, *the MEA dish*, and associated technologies for stimulating, recording, and optically imaging activity in "*the dish*" of cultured neurons. The dish

model-system is a living entity composed of cortical neurons from embryonic rats, dissociated and plated on an specially designed set of 64 electrodes called a “*multi-electrode array*” (MEA) where the neurons generate new connections to become a living network. At the time of our investigation the dish was a completely new kind of model-system – this lab was one of the first to investigate its properties and behavior. The researchers’ immediate focus was on trying to understand and conceptualize the novel phenomena produced by experimenting with the dish. The concept formation processes discussed here involved both transfer and modification of concepts from single neuron studies and engineering and the formation of emergence of novel concepts specific to network behavior. In the process of solving problems that arose in trying to understand dish behavior, one researcher built a computational model of the dish and the interaction between two kinds of models, *in vitro* and *in silico*, gave rise to a cluster of fundamentally novel concepts for understanding neural activity, ‘CAT’ (*center of activity trajectory*) being chief among them.

When the research started there was no established understanding of neural network communication and no knowledge of dish behavior or how to control and use it to study learning. The lab conceptualized *learning* in terms of the Hebbian notion of *learning as plasticity* (basically, changes in the brain from adding or removing neural connections or adding cells in response to experience) and the mathematical formulation known as the Hebbian rule (“*neurons that fire together wire together*”), which “*talks about learning between two neurons.*” They transferred this concept, expecting that “*our data will show something has to be added to the*

known equation in order for it to manifest is a population of neurons.”

Figure 2 is our schematic of the distributed cognitive system of Lab D as it evolved during the period of the episode discussed here. The system comprised the three primary researchers involved in the day to day activities, simulation models, and other technologies. Of course, here too, on-going interactions with the director were an integral part of the system. During year 1, after an initial period of getting the lab up and running (learning to build dishes, formulating lab protocols, and developing associated technologies and a suite of software tools they called MEAbench), the researchers all spent a considerable period together “*playing with the dish,*” which consisted of exploring the problem space through stimulating the neuronal network using different electrical signals and tracking the output (“*open-loop physiology*”). Additionally, D2 and D11 were engaged in developing ways of expanding the dish model-system into various kinds of “*embodied*” model-systems (“*closed-loop*”), computational animals (“*animats*”) and robotic embodiments (“*hybrots*”) that could be connected to the dish and would enable real-time feedback experiments. These embodied dish model-systems were also entirely novel to the research field.

[insert Figure 2 here]

My analysis begins in year 2 when the researchers had encountered a major problem: spontaneous dish-wide electrical activity seemed to be getting in the way of dish learning. To

interpret this behavior, the group borrowed two concepts: the notion of *burst*, transferred from single neuron studies, and the engineering notion of *noise*. They called the dish-wide spontaneous phenomena “*bursting*,” extending the meaning from single neuron studies, where it means the spontaneous activity of *one* neuron, to the *population* of dish neurons. Figure 3 shows burst activity in each channel of the MEA dish as recorded using their software MEAScope that produces a visual display that mimics a standard oscilloscope type representation of electrical activity.

[Insert Figure 3 here]

Bursting created a problem investigating plasticity, because it prevented the detection of any systematic change that arose due to controlled stimulation of the network. The researchers interpreted bursts as “*noise in the data...noise, interference in the way...so it is clouding the effects of the learning that we want to induce.*” They hypothesized that bursting arose because the neurons lacked the sensory inputs they would ordinarily get if they were in a live animal’s brain. Given this view, and the problem of the noise generated by bursts, the group decided that they needed to get rid of bursts. D4 began working on “*quieting*” bursts in the dish by providing the network with artificial sensory input in the form of electrical stimulation. She tried a range of stimulation patterns to lower the bursting activity in the networks and, after about a year, achieved a breakthrough, managing to stop the bursting entirely. However, despite the quieted

network, for the next six months she was unsuccessful in inducing plasticity in the network. The activity pattern evoked by a stimulus in the quieted dish did not stay constant across trials, but “drifted” away to another pattern. This drift prevented tracking the effect of a stimulus, because the network never responded the same to a constant stimulus.

During the period D4 was trying to quiet the network and induce plasticity, the researchers were engaged in largely separate though interrelated activities. D2 was working on an embodiment software module – the translation between the neuron network and the motor commands for controlling the animats and robots. This would be part of the closed-loop model-systems to be used in training the dish once they had the sought-after control structure. The control structure, however, had an unexpected origin: it was derived ultimately through interaction with a computational modeling simulation of the in vitro dish model-system. This in silico dish model was a 2nd order model built in order to try to understand the behavior of the living dish model.

Early in the quieting period, when that research seemed at an impasse, D11 decided to branch away from working with the in vitro model-system entirely, and develop a computational model that would simulate dish phenomena. As he put it, “*the advantage of modeling [computational] is that you can measure everything, every detail of the network.....I felt that modeling could give us some information about the problem [bursting and control] we could not solve at the time [using the in vitro dish model-system].*” D11 felt that to understand the phenomena of bursting he needed to be able to “see” the dish activity at the level of individual

neurons, to make precise measurements of variables such as synaptic strength, and to run more controlled experiments than could be conducted with the physical dish. Computational modeling was not part of the practices of the lab, largely because the director had rejected computational neural networks as a means of understanding neural activity in the brain, and he was skeptical it would be of any benefit. Fortunately he allows his students considerable freedom in their research. D11 moved to a different physical space where he could work on a computer, and only moved back into the physical space of the lab (represented by the dashed line on Figure 2) after he had successfully replicated experimental results from their dish, and the researchers could work together to exploit the computational model's results.

It is not possible to go into the details of building the computational model here (see, (Chandrasekharan & Nersessian, 2009)), however it is important to understand that D11 built the initial *in silico* dish not on the experimental data from their dish, but by drawing from other sources in neuroscience; in particular, from studies involving single neurons, brain slices, and other computationally simulated networks. The initial constraints the model adopted from their *in vitro* dish were not based on their experimental outcomes (i.e. the behavior of their dish), but had to do only with the physical construction of the dish. The computational model was tested and optimized with data from other MEA dishes first, and then their own. The model, as developed, is an computational model of the activity of a *generic* *in vitro* dish. Figure 4 provides a schematic of the computational modeling processes.

[insert figure 4 here]

After several months of building tentative versions of the computational simulation and probing these, D11 started to get what he called some “*feeling*” about the behavior of the simulated network under different conditions. When he was able to replicate results reported in the literature, he then used their in vitro dish data and was also successful in replicating their results. During this period he developed a visualization that captured the activity of the network as it ran (Figure 5). This visualization figured centrally in solving the supervised learning problem in the process of articulating a cluster of conceptual innovations, including CAT. Sometime during this period he moved back into the physical space of Lab D and all three researchers began working together in a reconfigured distributed cognitive system (depicted on the right side of Figure 2).

[insert figure 5 here]

A major contribution of the visualization is that it enabled D11 to notice - literally to see - interesting patterns in the way his model responded to different stimuli. These behavioral patterns were completely novel and distinct from what they understood about in vitro dishes. The visualization of the computational network’s activity shows the movement of an activity pattern across the entire network, in real time. In the visual display of the in vitro model-system the

individual neuronal activity is hidden, as is the propagation of activity through the dish. One can see activity across each channel as in Figure 3, but this display can only track activity at each electrode. That is, the MEAscope representation of activity does not have a representation of the entire network itself – it does not capture *burst movement across the network*. Thus, it was not possible to see from the MEAscope visualization whether there were patterns moving across the network, so the idea did not occur to them. Such patterns, however, did show up in the visualization of the in silico dish, leading them to wonder if they were also present in the in vitro dish. As expressed by D11, “*I can visualize these fifty thousand synapses...so you can see...after you deliver a certain stimulation, you can see those distributions of synaptic weights change.*” It is important to realize that computational visualizations are largely arbitrary; he could have visualized the in silico dish in a number of ways, including the one the group was familiar with: a per channel spike representation (Figure 3). However D11 imagined the dish as a network and visualized the model activity that way (Figure 5).

The computational model offers additional advantages in exploring burst phenomena. The simulated network could be stopped at any point and started again from there. Further, it is possible to provide detailed measures of significant variables, such as synaptic strength, which are not accessible for the in vitro model. Finally, a large number of experiments could be run at no cost, since the computational model does not require the careful and laborious processes involved in setting up and maintaining a living dish. When coupled with the visualization that enables visually tracking the activity of the network as it is happening, these features proved to

be a powerful combination. They gave D11 immediate access to a range of configurations and data that the in vitro dish could not provide. He made movies of the dish visualization as it ran, which showed the movement of activity patterns over time. He showed these to the other researchers (and us) so they too could “*come away with the same thing.*” In the process of running numerous simulations, he began to notice something interesting: there were repeated *spatial patterns* in the activity, both when spontaneous bursts arose in the in silico network and when it responded to stimuli. Basically, he found that there were “*similar looking bursts*” that propagated across the network, and a limited number of what he called “*burst types.*” The understanding he reached from these findings was that “*the spontaneous activity or spontaneous bursts are very stable.*” The others agreed and together they worked to develop a means of tracking and mathematically representing the activity of possible “*stable*” bursts across the network.

From this point, the group worked together on statistical analyses, experiments to see whether the “*drift immune*” measures developed for the in silico network transferred to the in vitro, and whether “*burst feedback*” in the in vitro dish could be used for supervised learning with the embodiments. Each of these is a novel conceptualization of the newly recognized dish behaviors. This phase of the research began from the idea that “*bursts don’t seem as evil as they once did;*” indeed, the group now began to develop the concept of *bursts as potential signals* (rather than only noise). Developing this idea led to the formation of several interrelated novel concepts, chief among them is ‘CAT’ (‘*center of activity trajectory*’) which provides an entirely

novel way of understanding neural activity and, if transferrable to in vivo phenomena, could prove of major importance to neuroscience.

Briefly, CAT is an averaging notion, similar to the notion of a population vector. A population vector captures how the firing rates of a group of neurons that are only broadly tuned to a stimulus, when taken together, provide an accurate representation of the stimulus. However, CAT is more complex than the population vector because it tracks the *spatial* properties of activity as it *moves* through the network. For instance, if the network is firing homogeneously, the CAT will be at the center of the dish, but if the network fires mainly at the left corner, then the CAT will *move* in that direction. CAT thus tracks the *flow of activity* (not just activity) at the population scale, and on a much quicker time scale than population vectors. The CAT is like a “signature” for a burst type in that each burst type has a corresponding range of similar-looking CATs specific to that type. The CAT notion was developed first for the in silico model and then mapped, adapted, and transferred to the in vitro dish model. Although our account ends here, in finishing their dissertation projects, the researchers were able to combine the CAT analysis and the earlier techniques developed for burst quieting to develop a set of stimulation patterns (a control structure) that has led to supervised learning by the living neuronal network.

Discussion. This case, too, demonstrates how the need to conceptualize novel phenomena built the lab into an evolving distributed cognitive system. The hybrid, bio-engineered in vitro dish mode-system is a novel research technology, designed to move beyond single neuron studies to more closely parallel neuronal network activity in the brain. At the start, concepts were

transferred provisionally from single neuron studies and engineering to understand dish phenomena with the understanding that they would likely need revision, and these both facilitated and impeded the research. Conceptualizing the spontaneous dish-wide activity ('bursts') as 'noise' created problems. A computational model of the dish was constructed as a way of envisioning the activity of the living networks. The in silico dish has different affordances, which provided new insights into the in vitro behavior and gave rise to novel concepts that enabled solving the main problem of the lab: to create supervised learning in the in vitro dish. Once the in silico dish could replicate the lab's in vitro dish results, it created a different kind of problem-solving system where all researchers now directed their efforts towards trying to understand 'bursts' as signals that could be controlled and made use of in supervised learning.

One significant affordance of the computational model is that it allowed D11 to choose to visualize the activity the way he envisioned it, as a network. Of course, since they were investigating learning as a network phenomenon, everyone knew that activity in the dish was network activity. But, no one had seen the network activity or a representation of it, so the computational visualization was built on a counterfactual scenario – a thought experiment: “If we were able to see into the dish...” Indeed, the researchers spoke of the visualization as enabling them to “*see into the dish.*” The network visualization provided a significant representational change from the per channel representation with which they had been working. The MEAScope visualization does not capture the *network* features of the in vitro dish and in vivo phenomena,

The computational network visualization captures the structure and behavior of the network. CAT is a novel notion that derives from the visualization of the movement of patterns of network activity and the capacity of the in silico model for running an unlimited number of simulations, which together enabled D11 to notice the repeated similar-looking patterns and for the group to seek to mathematically represent their behavior. The manifest nature of the visualization served to align the mental models of all the researchers - how they now envisioned the dish activity and to communally exploit the possibilities of *bursts as signals*. The in silico visualization worked as a generator of many types of lab activity, which when put together, created new concepts.

Conclusions

The frontier nature of the research conducted in labs in the biomedical engineering makes them prime locations for studying scientific practices surrounding concept formation in the wild. Endeavoring to understand novel phenomena makes concept formation an on-going dynamic within these research labs. Concepts are formed and articulated with respect to particular constructions of problems that become commitments of the specific community of practice. These commitments both afford and constrain future possibilities for research and include engineering specific kinds of technologies to address these problems. Designing (and re-designing), building, and experimenting with physical simulation models that parallel in vivo phenomena is a central practice of these communities. Simulation models constrain ways of proceeding without rigidly specifying in advance what moves can be made. These models embody constraints from engineering and biological sciences that interact in ways such that novel

structures and behaviors can emerge and the problem of representing these can lead to the formation of novel concepts.

In the labs we investigated, modeling practices associated with concept formation both drove the formation of, and arose from, complex distributed problem-solving systems. Simulation models constitute central components of these systems. These models provide concrete ways of imagining and envisioning potential scenarios for *in vivo* phenomena. Concept formation is not something that takes place just “in the head” of researchers, rather concepts are formed by a coupling of researchers and models in problem-solving processes. As we documented in our investigations in the tissue engineering Lab A, a germ cell conceptualization led to the development of a 40+ year research program, the last 20 of which articulated an evolving distributed cognitive system of *in vitro* research, of which we provide a snap shot mid-way. A major reconfiguration took place shortly before we entered because of the researchers’ felt need to have a better model of the blood vessel wall than could be provided by studying the reactions of “*cells in plastic*” to fluid flow. The lab’s development of the construct model opened a range of new possibilities for research, including constructing a novel animal model-system to further articulate ‘arterial shear’. Similarly, in neural engineering Lab D, the problem of conceptualizing learning in a network of neurons articulated a system that was at first configured around model-systems comprising a dish of living neurons and potential embodiments. Problem solving in this system included transferring concepts from engineering and single neuron studies to investigations of the lab’s *in vitro* models – living cultures of neuron networks – which both

facilitated and impeded problem solving. The sense that the lab's then current practices of investigating the dish were not yielding adequate information to understand its behavior led to adding a novel modeling practice, computational simulation, which, in turn, facilitated the lab's development of a cluster of fundamentally novel concepts. One way in which modeling practices promote conceptual innovation is that the processes of building and transforming models shift among representational formats. Different kinds of formats (e.g. physical model vs. computational model; visual representation vs. equation) afford different kinds of manipulations and different kinds of inferential processes.

Investigating scientific concept formation "in the wild" adds a significant dimension to science studies. Within the science studies fields there remains a divide between cognitive and socio-cultural accounts of knowledge-creating practices. There is scant research on concept formation on either side. The situation stems in part from a lingering mind – body dichotomy that still sees concepts as situated "in the head," that is, as exclusively mental entities and processes (Nersessian, 2005). However, in studying concept formation "in the wild" we see that it plays a significant role in the on-going dynamics of science practices – as scientists grapple with the dual tasks of trying to understand and make sense of complex, novel phenomena and of conceiving and building the artifacts by means of which they do their sense-making. An integrated understanding of science practice demands attention to concept formation. By situating cognitive processing in environments comprising people (with embodied minds) and artifacts, the analytical framework of distributed cognition affords a means of rapprochement. Scientists build

cognitive powers by actively *distributing cognition* (Hall, Wieckert, & Wright, 2010) to what Centina calls the “machineries” of knowledge-making or “epistemic machineries” ((Cetina, 1999), 2-5). Through studying “concept formation in the wild,” we see how concept formation processes are deeply entwined with “the manufacture” of epistemic machineries. The practices of science create complex social structures, rich material cultures, and societal implications, but they also create one of the most sophisticated expressions of human cognitive powers in large part through building epistemic machineries. The cases briefly developed here examine how the bioengineering sciences use complex modeling practices in the first instance to investigate in vitro phenomena, but also show how the models themselves become objects of study by creating novel phenomena to investigate, understand, and control. As these cases exemplify, a budding concept can drive the building of novel epistemic machineries that configure into distributed cognitive systems and scientifically novel concepts can emerge from systems in the course of experimenting via these machineries.

Acknowledgements

I thank the members of the Cognition and Learning in Interdisciplinary Cultures research group (www.clic.gatech.edu) who worked on this research project, Wendy Newstetter (co-PI), Lisa Osbeck, Elke Kurz-Milcke, Ellie Harmon, Chris Patton, and Sanjay Chandrasekharan. The Lab A case analysis, specifically was developed together with Ellie; the Lab D case analysis with Chris

and Sanjay. I thank also the researchers in these labs for being generous with their time in letting us observe and interview them about their research over an extended period. This research was made possible with the support of the US National Science Foundation grants DRL0106773, DRL0411825, and DRL097394084.

References

- Andersen, H., Barker, P., and Chen, X. (2006). *The Cognitive Structure of Scientific Revolutions*. Cambridge: Cambridge University Press.
- Arabatzis, T. (2006). *Representing Electrons: A Biographical Approach to Theoretical Entities*. Chicago: University of Chicago Press.
- Barsalou, L. W. (1992). Frames, concepts, and conceptual fields. In A. Lehrer & E. Kittay (Eds.), *Frames, Fields, and Contrasts: New essays in Semantical and Lexical Organization* (pp. 325-340). Hillsdale, NJ: Erlbaum.
- Boon, M., & Knuuttila, T. (2009). Models as epistemic tools in engineering sciences. In A. Meijers (Ed.), *Philosophy of Technology and the Engineering Sciences* (pp. 693-723). Amsterdam: Elsevier.
- Carey, S. (2009). *The Origin of Concepts*: Oxford University Press.
- Cetina, K. K. (1999). *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press.
- Chandrasekharan, S., & Nersessian, N. J. (2009). Hybrid analogies in conceptual innovation in

- science. *Cognitive Systems Research*, 10, 178-188.
- Chandrasekharan, S., & Nersessian, N. J. (2011). Building cognition: The construction of external representations for discovery. In L. Carlson, C. Hölscher & T. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 267-272). Austin, TX: Cognitive Science Society.
- Engestrom, Y. (1999). Activity theory and individual and social transformation. In Y. Engestrom, R. Miettintn & R. Punamaki (Eds.), *Perspectives on Activity Theory*. Cambridge: Cambridge University Press.
- Giere, R. N. (2002). Models as parts of distributed cognitive systems. In L. Magnani & N. J. Nersessian (Eds.), *Model-Based Reasoning: Science, Technology, Values* (pp. 227-242). N. Y.: Kluwer Academic/Plenum Publishers.
- Glaser, B., & Strauss, A. (1967). *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Piscataway, N. J.: Aldine Transaction.
- Hall, R., Wieckert, K., & Wright, K. (2010). How does cognition get distributed? Case studies of making concepts general in technical and scientific work. In M. Banich & D. Caccmise (Eds.), *Generalization of Knowledge: Multidisciplinary perspectives*. NY: Psychology Press.
- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Morgan, M. S., & Morrison, M. (Eds.). (1999). *Models as Mediators*. Cambridge: Cambridge University Press.

Nersessian, N. J. (1984). *Faraday to Einstein: Constructing Meaning in Scientific Theories*.

Dordrecht: Martinus Nijhoff/Kluwer Academic Publishers.

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. (2002). The cognitive basis of model-based reasoning in science. In P.

Carruthers, S. Stich & M. Siegal (Eds.), *The Cognitive Basis of Science* (pp. 133-153).

Cambridge: Cambridge University Press.

Nersessian, N. J. (2005). Interpreting scientific and engineering practices: Integrating the

cognitive, social, and cultural dimensions. In M. Gorman, R. D. Tweney, D. Gooding &

A. Kincannon (Eds.), *Scientific and Technological Thinking* (pp. 17-56). Hillsdale, N. J.:

Lawrence Erlbaum.

Nersessian, N. J. (2008). *Creating Scientific Concepts*. Cambridge, MA: MIT Press.

Strauss, A., & Corbin, I. (1998). *Basics of Qualitative Research Techniques and Procedures for*

Developing Grounded Theory (2nd ed.). London: Sage.

Vygotsky, L. S. (1962). *Thought and language*. Cambridge, MA: MIT Press.

Figure Captions

Figure 1

Lab A director's diagram of how he conceived of the interconnections among the lab's problems, researchers, and investigative strategies.

Figure 2

Our representation of the approximate time line of the research leading to the conceptual innovations and development of the control structure for supervised learning with the robotic and computational "embodiments" of the in vitro dish (years 2-4 of the existence of Lab D). The dashed line represents the period after D11 moved back into the main part of physical space of the lab and all three researchers began to actively collaborate on exploiting the findings stemming from the in silico dish.

Figure 3

The MEAscope per channel visualization of in vivo dish activity showing spontaneous bursting across the channels of the dish. Spontaneous bursting activity is represented by the spikes appearing in the channels. A relatively "quiet" dish would have no spikes in the channels, with all channels looking closer to channel 15.

Figure 4

Our representation of the bootstrapping processes involved in constructing, evaluating, and adapting the computational dish through numerous iterations. Once the in silico dish was able to replicate the in vitro dish behavior and the novel concepts were developed for it, the analysis was

mapped (adapted to the specifics of its design) and transferred to the in vitro dish, and evaluated for it.

Figure 5

A screen shot of the network visualization of bursting in the in silico dish.

Figure 1

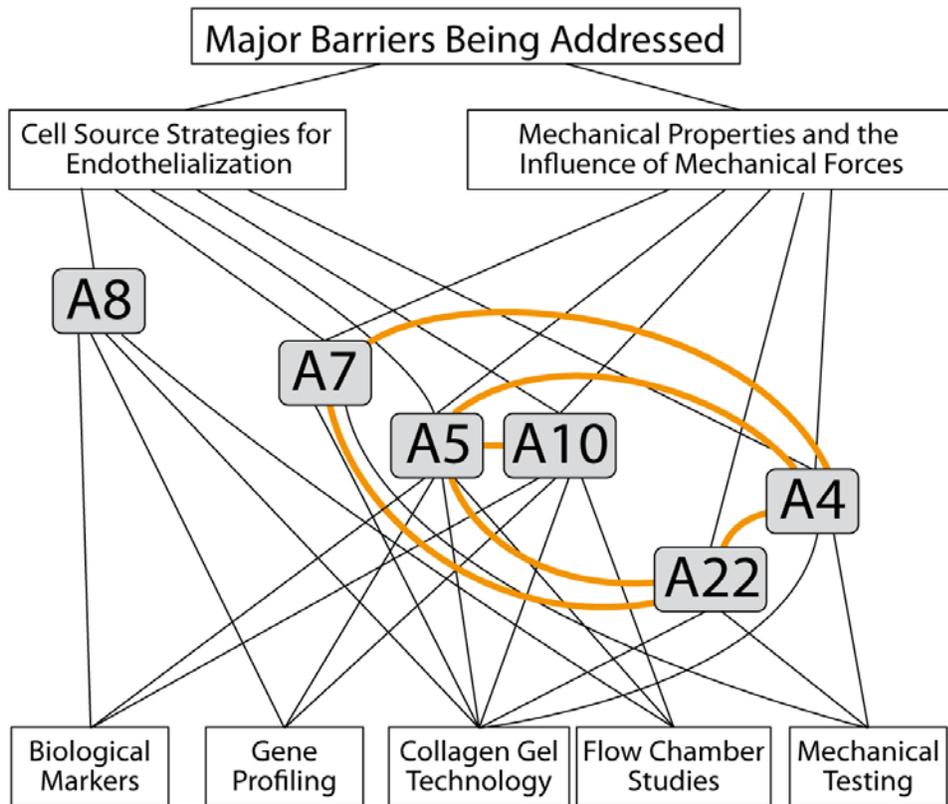


Figure 2

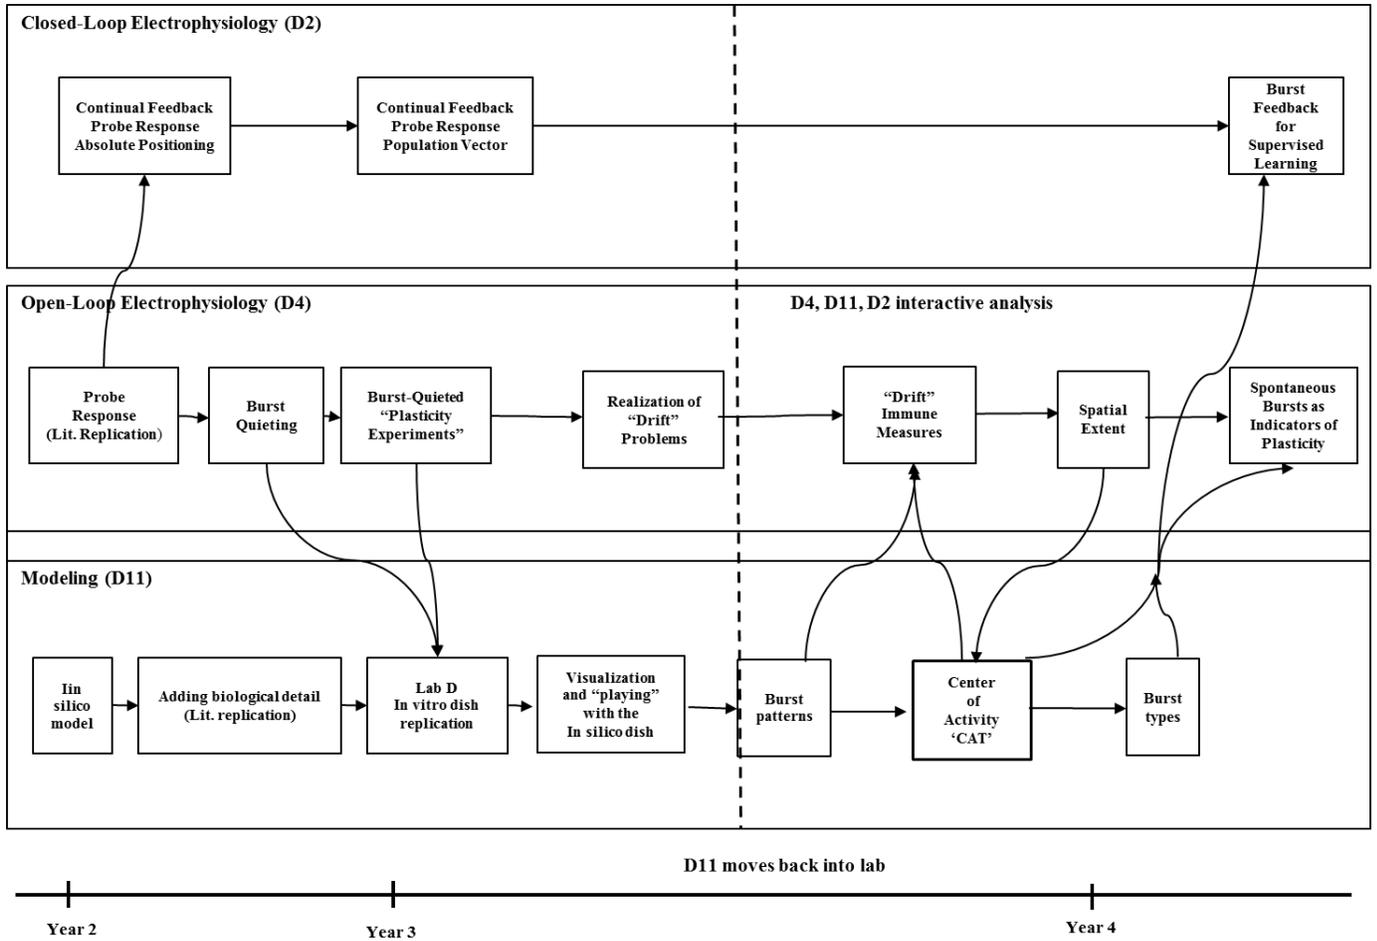


Figure 3

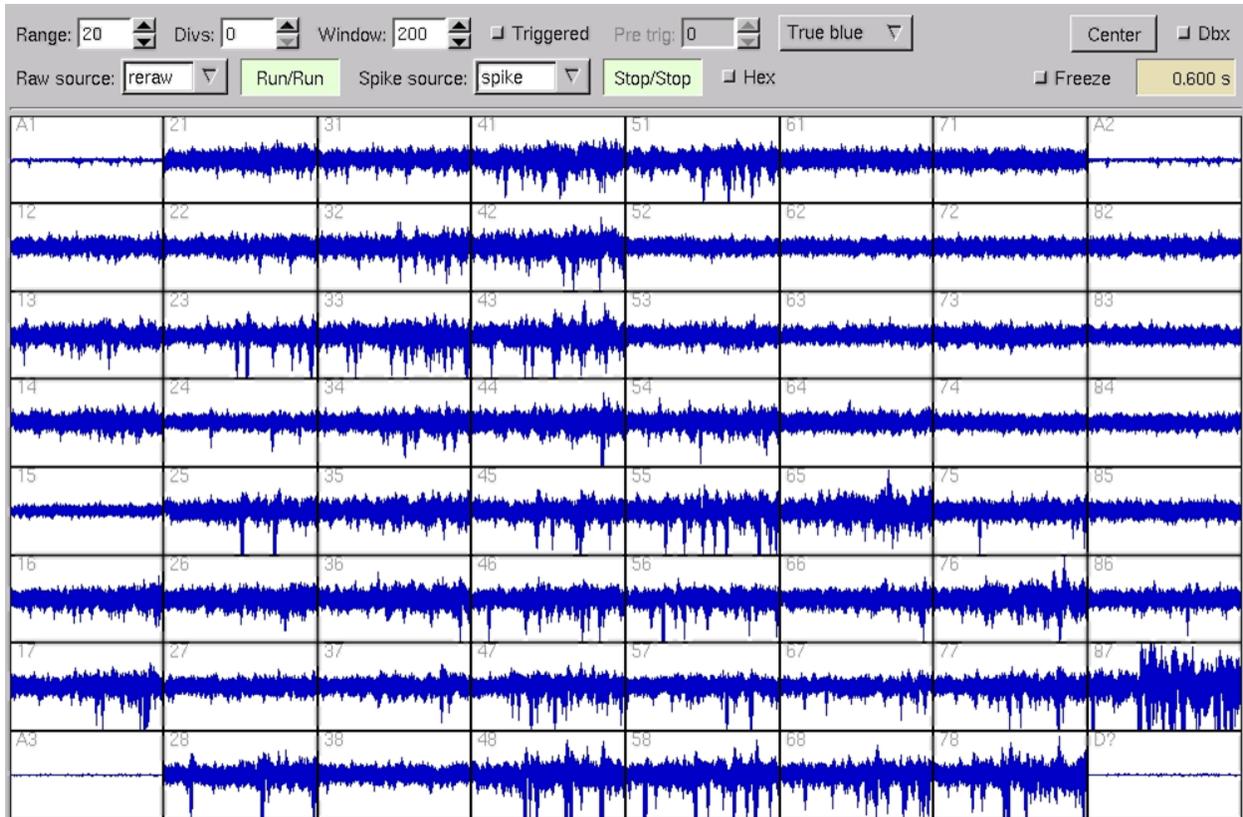


Figure 4

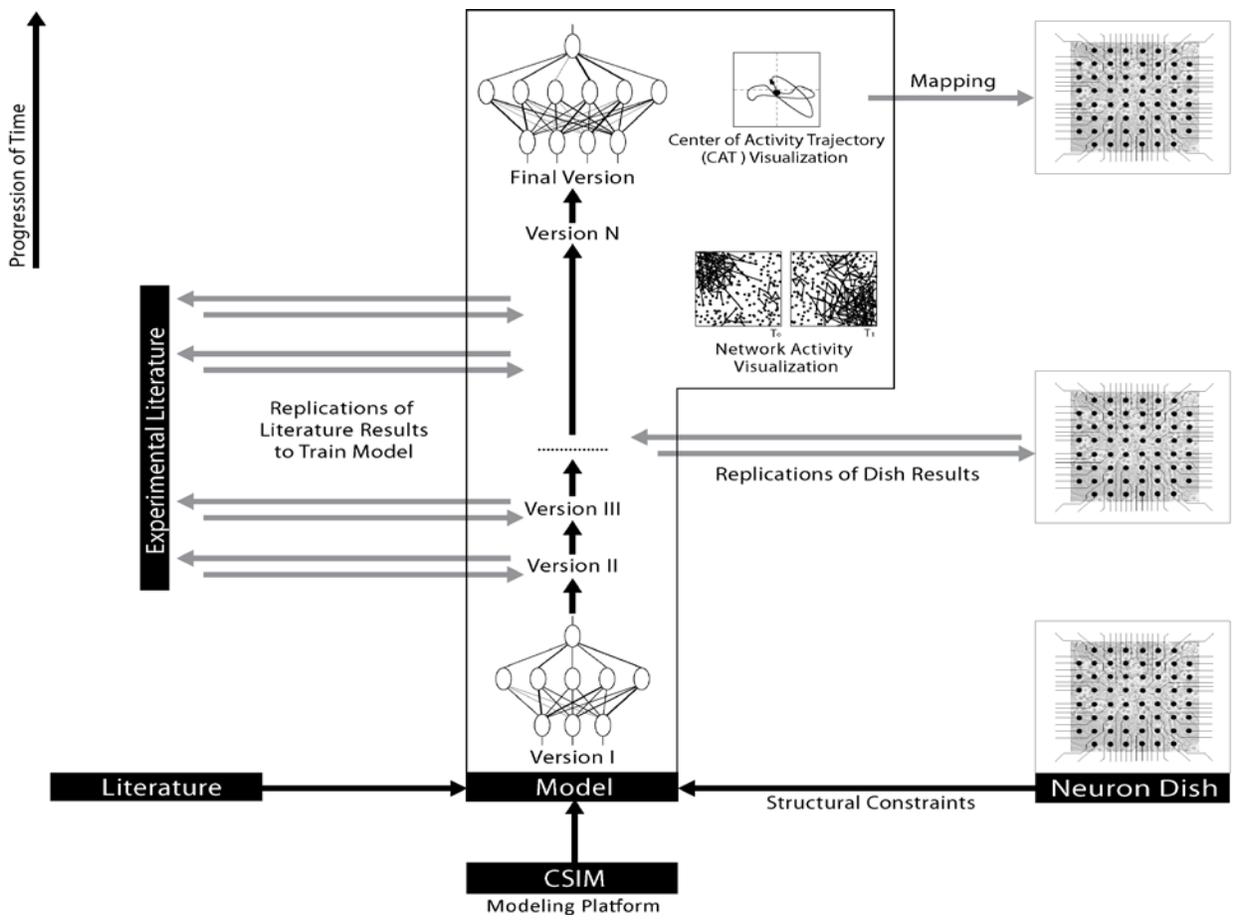


Figure 5

