1. Introduction

One of the goals of research in science education is to develop and study new images for instruction. Although there is certainly much variety in the positions adopted by researchers, many of the recently proposed innovations share a common orientation: They are driven by the desire to embed K-12 classroom learning in rich, meaningful contexts. The instructional approaches produced by this orientation have taken a number of forms and have been given a number of names. These include:

- "project-based" instruction (Krajcik, Czerniak, & Berger, 1999),
- "anchored" or "problem-based" instruction (Barron et al., 1998; Williams, 1992),
- "learning-by-design" (Harel, 1991; Hmelo, Holton, & Kolodner, 2000; Kolodner, Crismond, Gray, Holbrook, & Puntambekar, 1998).

In work with Daniel Edelson and Matthew Brown I have referred to this class of approaches as *task-structured*, as opposed to traditional *content-structured* curricula, because the sequencing of the resulting curricula are driven by an encompassing task (Sherin, Edelson, & Brown, 2000). Here, I am using the word "task" in a very broad manner, to include all of the "rich, meaningful contexts" in which researchers seek to engage students.

This new wave of curricular innovation shows great promise. However, the field is in dire need of new frameworks for studying and describing the conceptual change that is associated with task-structured curricula. In this proposal, I will argue that, because task-structure curricula often address content across a range of traditional disciplines, task-structured curricula tend to produce learning outcomes that are difficult to encapsulate and characterize. Thus, although we can certainly build on existing theoretical and methodological foundations, substantially new frameworks are needed for research to continue.

For this reason, I propose to conduct a basic research program directed at *devising theoretical frameworks and methodological strategies for characterizing conceptual change in task-structured curricula.* To develop these new frameworks and strategies, I will undertake an empirical program based around student interviews and classroom observations. It is not my intent to engage in new curriculum development or other design work. Instead, I will conduct this research in the context of development efforts already underway at Northwestern University, especially those that are part of the Center for Learning Technologies in Urban Schools (LeTUS).[1]

More specifically, there will be three broad types of research outputs from the proposed work:

(1) **Development of the "Conceptual Dynamics" framework and techniques.** The first product of this research will be an assembly of theoretical frameworks and methodological strategies that I will refer to as the "Conceptual Dynamics" program. To a major extent, the precise nature of the new frameworks and techniques must develop as the project progresses. However, as a starting orientation, I will be adopting a theoretical framework based around what I call a *functional-genetic analysis of knowledge system evolution*. Broadly speaking, this means that for any particular curricular intervention, I will ask two main questions: (a) What cognitive knowledge systems are involved? (b) How do these evolve over the course of the curriculum intervention? The Conceptual Dynamics framework, when developed, will provide researchers with theoretical tools for answering these questions in particular instances.

(2) **Application of the Conceptual Dynamics program.** Second, in parallel to the development of the Conceptual Dynamics program, I will apply the program to describe learning in particular curricula and, more generally, to answer some questions about learning in task-structured curricula. These questions include: What do students learn and fail to learn in task-structured curricula? Are there typical patterns in the difficulties observed? Is the learning that occurs in task-structured curricula, in any sense, fundamentally different than the learning that occurs in content-structured curricula?

(3) **Development of new resources for the teaching of students in the learning sciences.** I believe that, in the years to come, the learning sciences research community should focus more on developing theoretically and methodologically rigorous techniques for understanding the learning that occurs within novel interventions. In order for this to occur, we need researchers with the foundations necessary to develop these more rigorous techniques. Thus, as part of this project, I propose to develop new resources to support the teaching of learning

[1] LeTUS is funded by the NSF under award REC-9720383 (*Center for Learning Technologies in Urban Schools*, Louis Gomez, Ron Marx, Clifton Burgess, Juanita Clay-Chambers, PIs).

sciences students at Northwestern University. This will include the creation of a library of video segments drawn from the video archives generated by this research project.

Overall, I believe that this research project will help lay the foundations for my future work in research and education. I am devoted to developing rigorous theoretical accounts of learning and conceptual change in science. In past work, I have focused on the detailed characterization of particular knowledge systems, particularly within the domain of physics. (Refer to the summary of my past work in Section 6.) This project will allow me to begin work in a new and, I believe, productive direction: I will build a program, suited for a range of scientific domains, that attempts to describe conceptual change at the level of multiple interacting knowledge systems.

2. Why do we need new frameworks?

To this point, I have simply asserted that science education researchers are now advocating for a new type of curricular intervention, and that this type of intervention requires new frameworks for describing learning and conceptual change. In this section, I will argue for this assertion. I will begin by elaborating on some of the features of task-structured curricula, and on why these features pose special challenges for accounts of conceptual change. Then I will discuss the theoretical frameworks that exist for describing science learning, and I will begin to characterize the new type of framework that I believe is required.

2.1 How task-structured curricula are different

During the last two decades, there have been some dramatic ebbs and flows in the ideology of the science education research community, and in the preferred images of instruction that have prevailed among this community. First, beginning in the late 1970s and early 1980s, research concerning students' prior conceptions became prominent, producing dramatic and surprising results. Even students who did well in their coursework still seemed to possess beliefs that were dramatically at odds with fundamental tenets of a discipline. They had "prior conceptions," and these prior conceptions seemed to be extremely resistant to instruction (Eylon & Linn, 1988; Smith, diSessa, & Roschelle, 1993).

In response to these observations, researchers were led to advocate for curricula that were designed with the explicit intention of addressing prior conceptions. In this vein, there were a number of excellent curricula developed by researchers. For example, Barbara White and colleagues developed a curriculum called "ThinkerTools," which was very successful in teaching the physics of motion to sixth-grade students (White, 1993). Similarly, Smith and colleagues (Smith, Maclin, Grosslight, & Davis, 1997) developed a successful curriculum for teaching 8th-graders about matter and density. In their design efforts, the researchers assumed, very explicitly, that it was necessary to address prior conceptions. They thus began by mapping out students' prior conceptions, and articulating precisely what sort of changes were required for students to move to more appropriate understandings of matter and density. In this regard, an important observation was that, prior to instruction, many of the students seemed to have an undifferentiated notion of weight and density. Thus, one goal of instruction, Smith and colleagues surmised, must be to help students acquire differentiated notions of weight and density. As with ThinkerTools, this new matter and density curriculum seemed to produce dramatic positive outcomes.

The prior conceptions-driven approach to science curriculum design is far from dead. In fact, it has recently been codified in several forms, including the benchmarks and procedures developed by AAAS (Project 2061, American Association for the Advancement of Science, 1993). Furthermore, for many practitioners, this approach may still constitute their primary image of reform-based science instruction.

Although the prior conceptions-based curricula were an important development, there are some respects in which those curricula did not constitute a radical departure from what came before. In particular, the subject matter was still organized and sequenced in a manner tied to the traditional disciplines and content areas. For illustration, note that Smith and colleagues, in developing their new curriculum for matter and density, began with an existing *Introductory Physical Science* (IPS) curriculum (Haber-Schaim, Abegg, Dodge, Kirksey, & Walter, 1987). They then modified this curriculum to better address students' prior conceptions, but retained what they called the "main logic" of the existing curriculum (p. 360). To use my earlier terminology, these curricula were still a type of *content-structured* curriculum; the sequencing of content was driven by some existing understandings of the structure of a traditional discipline.

But this situation has changed with the current wave of task-structured curricula. Rather than being organized around a traditional disciplinary structure, task-structured curricula are organized around tasks, goals, or issues.

There are certainly important similarities between task-structured and prior conceptions-driven curricula. In both cases, the researcher-designers adopt a broadly constructivist orientation; they believe that the knowledge that students bring to the table is important, and that students must be actively engaged in reshaping that knowledge. But, the prevailing image that drives the “logic” of the curricula is different. When setting out to design a content-structured curriculum we ask: “What are the key concepts in this discipline and how can we build them up in a logical manner?” and “What are student misconceptions and how can we address them?” In contrast, when setting out to design a task-structured curriculum we ask: “What would make for an engaging task and would also allow useful content to be addressed?” and “How can we scaffold students as they move from start to finish on this task?”

The point I have been working up to here is this: These fundamental differences in the orientation underlying task-structured curricula lead to real differences in the nature of what is taught and learned. There are two big issues here:

(1) The content required to engage in a meaningful task tends to cut across multiple traditional areas of content. Thus, the slice of content associated with a particular task-structured curriculum unit may well differ from the slice that would be associated with any content-structured curricula. This does not necessarily mean that, over the course of a year or multiple years, different content will be addressed. But it does mean that, at the least, this content will be organized into units and sequenced in a very different manner.

(2) In a task-structured curriculum, the task not only dictates what will be taught, it also dictates the manner—or “depth”—to which the various portions of disciplinary content must be understood by students. For this reason, some issues will be covered in great depth (measured against our traditional conceptions of a discipline). In other places, students will learn just enough to “get by.” In contrast, in a content-structure curriculum, we often attempt to build up content in a manner that we believe reflects the a priori structure of a discipline.

Because of these two issues, there are reasons to believe that producing accounts of student learning within task-structured curricula will be especially difficult. To be clear, it is not just that task-structured curricula teach “process skills” in addition to what has been called content. Rather, the content covered in these new curricula — understood even in a narrow sense that excludes process skills — has different properties when compared to content-structured curricula, and these features may make the description of learning difficult.

In regard to this, one feature of Smith’s intervention is very telling. Note that what Smith and colleagues developed was not just a curricular intervention; it was curricular intervention *and* an account of student learning, all rolled into one. They mapped out student conceptions and the relevant conceptual territory, and then carefully moved students through this territory. This is entirely typical of the prior-conceptions-driven curricula; part-and-parcel of these curricula is an account, from beginning to end, of the conceptual change that the curriculum is intended to engender.

Again, the situation in task-structured curricula is quite different. The central image is one of students moving through a task, not of students moving through a conceptual space. This is not to say that the designers of task-structured curricula don’t worry about drawing out student conceptions, and making sure that students engage with those conceptions. But, unlike the prior wave, task-structured curricula typically do not come equipped with a single broad account of the conceptual change that is to be engendered.

Although the above discussion may seem to favor content-structured over task-structured curricula, it should be noted that a number of strong arguments can be given in favor of task-structured curricula. Central to these arguments is the premise that learning should occur in contexts in which the utility of new knowledge is manifest for students. The various task structures — such as problems, projects, and design tasks — are seen as structures that can provide these contexts. Furthermore, it is argued, these contexts can embody important features of real scientific activity, and thus give students experience with fundamental aspects of the scientific endeavor, such as reasoning from data, and reasoning about complex problems. For these reasons, task-structured curricula merit our attention.

2.2 Current theories of prior conceptions and conceptual change

In the above discussion, I discussed features of task-structured curricula that may make it difficult to describe the learning that occurs in these curricula, and I suggested that a new theoretical apparatus is needed. In constructing this new apparatus, I will be able to build on the substantial work concerning science learning that has already been done. Broadly speaking, there are two related bodies of work that I will draw upon: research concerning *prior conceptions* in science, and research concerning *conceptual change*. Here I briefly discuss each of these bodies of work.

2.2.1 Research on prior conceptions

Research in prior conceptions attempts to find out what students know about various science topics prior to formal instruction. The quantity of this work is truly substantial, and it has been conducted across a great range of scientific domains. These include, for example, portions of physics, chemistry, astronomy, biology, and earth science. (A selection of this work, listed in the bibliography collected by Pfundt and Duit (1994), lists over 4000 citations.) Furthermore, in addition to mapping out the *content* of students conceptions across a great many domains, the prior conceptions researchers have devoted some of their effort to discussing the *form* of students' prior conceptions. At the heart of these discussions are questions such as: How coherent are students' prior conceptions? Do student conceptions take the form of naïve theories? If not, what form do they take? (diSessa, 1993; Samarapungavan & Wiers, 1997; Vosniadou & Brewer, 1992). This existing research on prior conceptions will be helpful for identifying specific prior knowledge resources, and will provide insight into the nature of these prior resources.

2.2.2 Research on conceptual change

Loosely speaking, while the prior conceptions work describes the state of student understanding prior to formal instruction, conceptual change research attempts to characterize the process by which a student's understanding *evolves*. More strictly, the term conceptual change is typically not applied to all sorts of changes to a student's understanding. Rather, researchers generally restrict the use of this term to include only more interesting or "deep" varieties of learning, and exclude less interesting types of learning, such as fact learning or gradual changes to beliefs.

In fact, one of the main goals of conceptual change research seems to be to characterize what differentiates deep from more shallow varieties of change (diSessa & Sherin, 1998). In this regard, an analogy is often made to the history of science; conceptual change is seen as corresponding to a "paradigm shift" in the science understanding of an individual (Posner, Strike, Hewson, & Gerzon, 1982). The idea is that, as in a Kuhnian paradigm shift, conceptual change involves a shift to a new understanding of the world in which the terms employed are incommensurable with the terms of the preceding conceptual system. In other words, conceptual change is understood as involving changes in "the very concepts" at the heart of the conceptual system (diSessa & Sherin, 1998).

This is the sort of change that Smith's matter and density curriculum attempts to engender. The idea is that, to understand the subject matter that is addressed by the curriculum, it is not enough for students to acquire new beliefs concerning their existing concepts of weight and density. Instead, these very concepts must change; the students' concepts of weight and density must be differentiated from a single parent concept.[2]

There are several reasons that the present project requires that I add to this theorizing about conceptual change. First, as I began to argue in the previous section, describing the conceptual change that occurs in a task-structured curriculum may be particularly complex. For curricula that are focused only around such concepts as weight and density, it might be good enough to describe conceptual change in terms of processes such as the differentiation or integration of a few key concepts. But, with interventions such as task-structured curricula, that cut across many traditional areas of content, these monolithic characterizations will probably not be good enough. At the least, any task-structured intervention is likely to involve multiple changes of this sort.

Second, there are many respects in which current characterizations of conceptual change are inadequate, even for describing conceptual change in the regimes that were the focus of the original researchers. This is not because the research that has been done to date is not excellent. The problem is that conceptual change research is very hard to do — it is hard to capture these changes empirically — and conceptual change research is still at a comparatively early stage. The result is that existing accounts of conceptual change are still very limited. For example, with Andrea diSessa I have argued that many of the present accounts are problematic because they describe conceptual change as involving changes in "the very concepts," without ever providing an accompanying definition of "concept" (diSessa and Sherin,1998). Furthermore, the term "concept" is applied to a large variety of notions, including such notions as "dog," "force," and "number." Following these observations, we argue that, for our theories of conceptual change to have any precision, we must replace the informal notion of "concept," with a variety of more precisely defined theoretical constructs.

[2] See Wiser (1988) for a similar discussion of differentiation around the notions of heat and temperature.

2.3 Toward a new theoretical apparatus

In sum, though there is much to build upon, there is still a great deal of work to be done in accounts of science learning. For the present project, some of these accounts must be refined, and they must be extended so that they are appropriate for the new complex interventions that are the concern of this work.

One approach that we could perhaps adopt would be to simply say that the learning and conceptual change in task-structured curricula involves multiple and diverse changes. We would assert that there is nothing new about the type of changes that are involved, there are simply many of these changes or, at least, they are differently distributed. In this case, we might just look to identify a few specific changes that we believe are of particular importance, and we would not attempt to attend to the full scale of any complexity.

However, I think that a better approach would be to not give up on describing the broad scope of the complex changes that occur in task-structured curricula. Instead, I would like to be able to describe the "big picture" of the conceptual changes that occur. To do this, I believe that we need an approach that focuses on a different grain size: the large scale evolution of multiple complex knowledge systems.

To build this new approach, I begin from my previous research. In prior work, on my own and with Andrea diSessa, I have engaged in attempts to characterize complex knowledge systems. That work was done as part of a larger program that we refer to as *Knowledge Analysis* (Sherin, in press). The focus of Knowledge Analysis is on describing the nature and development of extensive knowledge systems. A basic presumption underlying this program is a belief in a complex mental ecology consisting of diverse types of elements. According to this viewpoint, much of science understanding is likely to involve interactions among elements of this complex ecology. Thus, in order to understand scientific learning, we must, to some extent, attend to this complexity.

The Knowledge Analysis program has been exemplified in such projects as diSessa's (1993) work on a portion of intuitive physics knowledge that he calls the sense-of-mechanism. According to diSessa, the sense-of-mechanism is a complex system consisting of many individual elements he calls *p-prims*. In addition, in my own work, I have argued that portions of the sense-of-mechanism evolve into a new system consisting of elements that I call *symbolic forms* (Sherin, in press). In that work, I argue that symbolic forms constitute the conceptual vocabulary in terms of which symbolic expressions are understood by physics experts.

In much of this prior Knowledge Analysis research, we attempted to produce an analysis in terms of elements that constitute the *primitive* constituents of a complete cognitive analysis. (P-prims and symbolic forms are examples of such primitive elements.) However, because learning in task-structured curricula is complex, and involves changes in many knowledge systems, a primitive elements analysis would be extremely difficult, and it might not capture the most important learning phenomena. Thus, as stated above, I instead propose to look at the large-scale dynamics of entire systems, viewed as a whole. As I said in the introduction, the program of research I propose involves asking two big questions: (1) What knowledge systems are involved? (2) How do these evolve over the course of a curricular intervention?

In some respects, this is akin to the analyses performed by meteorologists, who look at the behavior of storm systems and large-scale air currents in order to understand weather phenomena. They study the phenomenology of these large scale systems, and learn to understand their dynamics and make predictions. This does not mean that meteorologists never employ their understanding of more basic, smaller-scale processes, such as the fact that warm air rises, and that condensation occurs under certain conditions. But this knowledge is helpful insofar as it helps them to build some understanding of the large-scale phenomenology.

I propose to engage in a similar program for the knowledge systems that are important to science learning. I will continue, broadly speaking, to work within the Knowledge Analysis paradigm. But I will focus on the large scale behavior and interactions of multiple knowledge systems, rather than on the constitution of individual knowledge systems. For the purposes of convenience, I refer to this new sub-program as the program in "Conceptual Dynamics."

In some respects, this program can be understood to be an extension of the ambitious program conducted by Susan Carey (Carey, 1985, 1988). According to Carey, students' intuitive science knowledge can be understood to be a collection of relatively large and autonomous *theories*. With development and formal instruction, these theories evolve; they can change, combine, and, sometimes, a single theory can split into two. The relationship between what I propose and Carey's work is complex, and I cannot fully discuss this relationship here. But there are some very fundamental differences. For example, I am interested in changes that occur over the course of a curriculum unit, rather than the longer-term developmental changes that are the core of Carey's interests. More profoundly, I believe that my partitioning of student knowledge into knowledge systems will produce a less coarse — and, I hope, more apt — characterization than Carey's partitioning into theories. Here I will say just a little more about how I will identify and partition knowledge systems.

In this section, I have worked from the premise that it would be better to build an analysis around some higher level of abstraction — on systems of knowledge — rather than to attempt to identify the basic constituents of knowledge. But this premise leaves much work to do; in order to make progress, I will need some way to identify and draw boundaries around knowledge systems. To a significant extent, this task must be left as one of the goals of the project; I must work to discover the appropriate type of abstraction for my analysis.

But there is a little more I can say here. First note that, as with weather systems, we can acknowledge that any boundaries must be fuzzy, while still identifying larger systems that are worth treating as identifiable units. Thus, we do not necessarily need to be troubled by the fact that any method for identifying systems will almost certainly leave fuzzy boundaries. Second, the knowledge systems of interest will likely need to exhibit certain properties. We are interested in identifying collections of elements that are moderately large, and that cohere in some manner.

Third, there is one approach for identifying knowledge systems that I believe is promising, what I call a *functional-genetic* approach. What this means is that, in essence, I will look to identify knowledge systems by considering their function and genesis; that is, I will view knowledge systems as arising to serve particular functions within particular classes of activities. Thus, when faced with a some new curricular activities I will ask: What previous activities might students have participated in that would have produced useful knowledge for this task? How must this knowledge be adapted for its new purposes? In some preliminary studies of how students learn to use external representations, I have employed a similar approach to good effect (Sherin, under review).

3. A task-structured curriculum: The Global Warming Project

In the preceding sections, I argued for the need for a new variety of theoretical framework, and I began to say how I propose to fill that need. But the above discussions were all carried out at a high level, and more concrete examples are needed. For that reason, in this section and the one that follows, I will present extended examples in order to illustrate the nature of the investigation I have in mind. Here, I will draw on work done with Daniel Edelson, Matthew Brown, and other collaborators at Northwestern University (Sherin, Edelson, & Brown, 2000).

With these collaborators, I have conducted some pilot work in the context of a particular project-based science curriculum, the Global Warming Project (Edelson, Gordin, & Pea, 1999). The Global Warming Project (GWP) is an 8-10 week middle school science unit created by the Center for Learning Technologies in Urban Schools (LeTUS), in collaboration with Chicago Public Schools. In this curriculum, the students adopt the role of scientific advisors to heads of state for various countries, and they are given the task of preparing briefings for the leaders of their respective countries. More specifically, the curriculum is organized around three briefings that students must prepare, each of which pertains to some aspect of global climate change: (1) How could we tell if the Earth was getting warmer? (2) What might be causing global warming? (3) What are the predicted implications of global warming for individual countries and what solution strategies should they pursue?

The curriculum gives the students a significant amount of support in acquiring the background that they need in order to prepare these briefings. In the weeks prior to the writing of each briefing, the students perform laboratory experiments, engage in computer-based activities, and participate in classroom discussions, all of which are directed at helping them to understand the core content issues.

3.1 The nature of content in the Global Warming Project

As a first step in examining the Global Warming Project curriculum, I want to discuss how the GWP exemplifies the properties that I earlier attributed to task-structure curricula. To begin, note that, in order to prepare the briefings specified in the GWP, students need to understand content that is usually taught separately in a number of distinct disciplines, such as biology, chemistry, physics, and the earth sciences. For example, the topics that are addressed in the Global Warming Project include:

- radiative energy transfer, reflectivity, and absorption;
- respiration, photosynthesis, decomposition, and the carbon cycle; and
- the Earth's energy balance, the hydrological cycle, and the greenhouse effect.

It is absolutely necessary that aspects of all of these topics are included in the curriculum; students would not be able to complete their briefings, as desired, if these topics were not addressed. For this reason, the GWP curriculum materials include specific activities that are directed at these topics. This is the first feature that I have argued is typical of task-structured curricula: the task creates a need for content that cuts across traditional disciplines.

The second property of task-structured curricula — the issue of "depth" — follows quickly from the first property. Clearly, it would not be possible for the Global Warming curriculum to address all of the above topics to a high level of scientific detail. If we were teaching a content-structured curriculum focused around certain aspects of plant biology, for example, we could choose to spend a significant amount of time to teach students about the process of photosynthesis. In a task-structured curriculum, however, we are not free to make such decisions. Instead, where we spend time and effort is largely dictated by the task. In the case of the GWP, for example, extensive attention to the processes underlying photosynthesis would probably not be sensible.

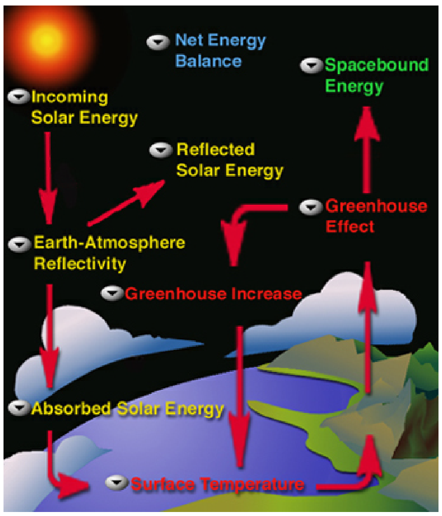

Figure 1. Diagram of the Earth's energy balance.

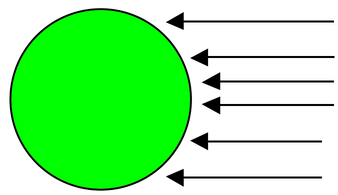

Figure 2. Incoming solar energy striking the earth. Notice that the arrows near the equator are perpendicular to the surface. In contrast, the arrows nearer the poles strike the Earth at an angle that deviates more greatly from the perpendicular.

Because the issues here are very important, I want to illustrate them further by working through one part of the Global Warming curriculum in some detail. In the Global Warming Project, students must learn what happens to the solar energy that is incident upon the Earth. Some of the relevant information is shown in the diagram in Figure 1, which illustrates the Earth's energy balance. In the teacher's manual that accompanies the Global Warming curriculum, the teachers are given precisely the following description:

> "Sun's rays reach the earth's atmosphere.
> Some sunlight is reflected by the earth's atmosphere and surface.
> Some sunlight is absorbed by the earth's atmosphere and surface.
> The energy that is absorbed contributes to the warmth we feel. … "

Thus, part of what students are intended to learn in the GWP is the story that is encapsulated in these few lines. However, in order for students to understand the climate phenomena that underlie global warming, some aspects of this story must be significantly elaborated; that is, more time must be devoted to unpacking some aspects of the relevant phenomena. For example, one place that the curriculum elaborates is around the factors that influence earth-atmosphere reflectivity. In particular, students learn that the fraction of light that is reflected by the Earth's surface and atmosphere varies over the Earth. Where the surface and atmosphere are lighter, more sunlight bounces off, and the Earth is cooler. In contrast, where the surface is darker, less sunlight bounces off, and the Earth is warmer.

Another place that the curriculum elaborates is in how incoming solar energy varies over the surface of the earth. The students learn that, because sunlight strikes the Earth more directly near the equator, the amount of incoming solar energy is greater nearer the equator than near the poles (refer to Figure 2). The GWP devotes a significant amount of time to this issue. For example, the students perform a lab in which a flashlight is shone on pieces of paper, with the flashlight inclined at various angles. In addition, they engage in computer-based analyses of scientific visualizations that show how incoming solar energy varies over the surface of the Earth.

Once again, the point here is that where the curriculum elaborates depends in a sensitive manner on the requirements of the task. This is made somewhat clear when, as above, we note where the curriculum chooses to focus attention. However, this sensitive dependence on task can be made even more striking if we look closely at what is *not* addressed by the curriculum. For illustration, consider again issues surrounding reflectivity and incoming solar energy. We can ask: How much do students really need to know to understand the relevant issues? For example, do they need to know *why* lighter colors reflect more than darker colors, or simply that they do? More fundamentally, do they need to know what light is? Do they need to know what energy is?

Once again, it would not be possible for the curriculum to answer all of these questions to a high level of scientific precision. The GWP does work hard to teach students the simple fact that lighter colors reflect more than darker colors, but it does not attempt to teach them *why*. More profoundly, the curriculum does not ever directly attempt to help students learn about the basic nature of light. Instead, the GWP chooses to rely on the intuitive understandings of light that students possess prior to the curriculum. This is important, and I will say more in a moment.

Thus, the manner in which the GWP builds up content is somewhat complex. There is no illusion of building up material from prerequisites in a systematic manner. Instead, the curriculum must, in various places, “make do” with students’ existing understanding.

In summary, in this section I have looked at the nature of content in the GWP curriculum unit, and we have seen how the GWP exemplifies the properties that I ascribed to task-structured curricula: The content covered within this single unit cuts across multiple traditional topics and disciplines, and the manner in which specific issues are addressed depends in a sensitive manner on the task.

4. Conceptual Dynamics in the Global Warming Project

In this section, I continue my attempt to ground my high level discussions in analyses of the GWP. In particular, I will try to give a sense for what a Conceptual Dynamics account would look like, and how it can be grounded in empirical work. Recall that a Conceptual Dynamics analysis involves asking two broad questions: (1) What knowledge systems are involved? (2) How do these knowledge systems evolve during the time of the curricular intervention. In this section I will describe how, in pilot work, I have begun to answer these questions for the Global Warming Project. As part of this work, my collaborators and I have conducted clinical interviews with approximately 40 middle-school students who were currently engaged in the GWP. These interviews were conducted in the context of 5 separate enactments of the curriculum.

4.1 What knowledge systems are involved?

As a partial answer to the first question, I want to discuss two broad categories of knowledge resources that may be relevant to the learning that occurs in the GWP curriculum:

(1) *Resources for reasoning about climate phenomena*. One category of knowledge resources are resources developed expressly for the purpose of reasoning about climate phenomena. It is an open question as to whether students possess any substantial resources of this sort prior to entering the GWP curriculum.

(2) *Resources for reasoning about other physical phenomena*. Students will possess relevant resources for reasoning about physical phenomena that were not expressly developed for the purpose of reasoning about climate. For example, in everyday life, students have opportunities to reason about phenomena involving light, heat, and temperature.

Before proceeding, a couple of caveats are merited here. Note, first, that the analysis embodied in the above two items is at a level of above the level of knowledge systems. This list describes *categories* of knowledge systems. In contrast, in my Conceptual Dynamics analyses, the intent is to perform analyses in terms of specific knowledge systems. Second, even though these two categories are potentially broad, they still omit a substantial portion of the relevant knowledge. For example, they omit specific geographic knowledge, and resources pertaining to the use and understanding of external representations. (See, for example, diSessa, Hammer, Sherin, & Kolpakowski, 1991; Sherin, under review). Nonetheless, these two categories can serve as a starting point for me to illustrate the sort of analysis I have in mind. To start, I will discuss observations from interviews conducted with students *prior* to the GWP curriculum, which were intended to explore the resources in these categories as they existed before instruction.

4.1.1 Resources for reasoning about climate phenomena

In daily life, there are occasions in which students might reason about climate phenomena. They might wonder, for example: Why is it colder in mountainous regions? Why is it colder near Lake Michigan? I'm traveling to Australia, what is the climate like there at this time of year? In some of our interviews, we attempted to get at the knowledge underlying students' abilities to reason about these phenomena; we asked questions designed to see whether and how students reason about the temperature variations that exist over the surface of the Earth.[3] For example, in some of these questions, students were shown pictures of various landscapes, including a desert and forest. They were then asked to discuss the climate that would be found in these regions, and explain why this region would have this climate. In addition, students were asked to color a map of the world in order to indicate the temperature of different areas of the map. Then they were asked to explain their colorings.

Perhaps not surprisingly, it turns out that, prior to the GWP, the middle school students we interviewed possessed very limited resources for reasoning about climate phenomena. In particular, much of students' knowledge here seemed to take the form of *nominal facts*.[4] For example, most of the students stated that temperatures were higher near the equator, and they believed that it was always hot in deserts. In addition, many of the students had experiences that permitted them to make judgements about particular locales. For example, one student had been to Australia, and stated that it was hot there. However, most of the students had great difficulty in producing explanations that went beyond these nominal facts. For example, they did not have ready answers concerning *why* it is hotter near the equator. When pressed by the interviewer, the students could sometimes produce explanations, but we believe that these explanations were usually invented on the spot, and students seemed to have little commitment to them.

Thus, to the extent that middle school students have a knowledge system for reasoning about climate phenomena, this seems to be of a relatively weak sort. It is sparsely populated with a few elements, many in the form of nominal facts, and these elements are probably not richly interconnected. In fact, we may ultimately not want to call this collection a "system;" prior to instruction, there is little knowledge here and even less coherence.

For comparison, contrast this to the knowledge that people have for reasoning about physical phenomena, such as the motion of objects in the world. Although there are certainly problems with this knowledge from the point of view of formal physics, this knowledge is nonetheless rich. We know a great deal about such things as what happens when two objects collide, or when a liquid is spilled across a surface. Thus, learning about the physics of motion may involve very different cognitive learning processes than learning about climate phenomena. This is the sort of issue that a Cognitive Dynamics analysis can capture

4.1.2 Resources for reasoning about other physical phenomena

Since the existing resources for reasoning about climate are weak, the GWP curriculum must do a lot to co-opt other knowledge resources, particularly other resources for reasoning about physical phenomena. These include resources developed for reasoning about heat and temperature in everyday phenomena, as well as about light and energy. The GWP counts on the existence of these resources; as we will see the GWP curriculum makes significant assumptions about what resources are available, and how much work must be done to refine and integrate these resources.

In the pilot work with my collaborators, I set out to look empirically at the nature of some of this other knowledge. For example, we asked students a variety of questions concerning light-related phenomena, centered around a few imaginary situations. In one of these situations, we asked students to imagine that they, with the experimenter, were standing near to a wall, in a very large dark room. The experimenter, who is holding a lamp, begins to walk backward from the wall. The student is then asked: How would what you see on the wall change as the lamp moves backward from the wall? Why?

In answer to this question, we hoped students would say that the light on the wall appears dimmer, and that the central circle of light on the wall grows larger. However, there was actually a great deal of variety in students' answers, revealing quite an array of modes of reasoning about light-related phenomena. For example, strikingly, some students did not seem to reason about light as something that *travels*. Instead they seemed to believe that, at

[3] These interviews were primarily designed and conducted by Elena Kyza.

[4] Here, I am using the term "nominal fact" as it is used by diSessa (1996). A nominal fact is a knowledge form that typically has a simple formulation as a sentence. "Everything is relative" is an example.

the instant that a light bulb is turned on, an illuminated area is instantly created around the bulb (refer to Figure 3). In some other cases, students seemed to reason about light as something having fluid-like properties, including the ability to flow around surfaces.

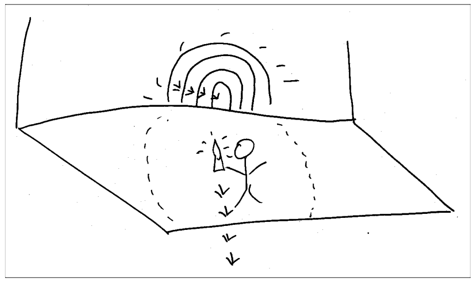

Figure 3. One student's drawing of the experimenter and lamp moving backward from the wall. This student believed that the lamp simply causes a sphere of light to come into existence around the bulb. In the this drawing, note that the illuminated area on the wall gets smaller as the bulb is moved backward.

Some of these observations are possibly worrisome; it seems that some of students' existing resources might be of a problematic sort, and they might not be sufficient to support the learning that is the goal of the GWP. For example, I just mentioned that some students do not even reason about light as something that travels. For these students, light simply fills the volume around a light source. But, if a student doesn't even think of light as something that travels, how can they understand that light travels from the Sun to the Earth? Clearly, this is a possible problem, and one that should be investigated empirically.

4.2 How do the knowledge systems evolve during the curricular intervention?

How do the above knowledge systems evolve during the GWP? How are they refined and integrated? What new knowledge systems are formed? In the above discussion, I already suggested some of the changes that must occur if the curriculum is to be successful. Specifically, I suggested that resources for reasoning about other physical phenomena must be co-opted and integrated with the sparse knowledge that students possess for reasoning about climate phenomena, in order to form an enriched knowledge system for reasoning about climate.

However, any straightforward account in which some existing resources are simply co-opted is likely to be over-simplistic. In the above section, we saw that some of students' prior knowledge might be problematic; in particular, some students seemed to possess alternative models of light that might not be appropriate for the understanding that the GWP hopes to engender. Thus, at least in some cases, these resources for reasoning about other phenomena may have to be refined before they can integrated into a new, useful knowledge system for reasoning about climate phenomena.

Furthermore, there is reason to be worried that this refinement might not occur in the GWP. As I stated earlier, the students are never taught about the nature of light; thus, none of the alternative models of light are explicitly addressed by the curriculum, and more appropriate models were never explicitly taught. Thus, it is clearly worth wondering whether this co-opting of existing resources can progress as desired in the GWP.

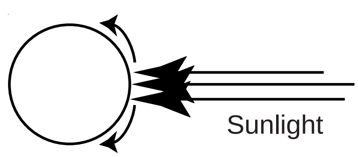

Figure 4. Sunlight flows toward the Earth, striking near the equator and then spreading out toward the poles. As the light spreads further from the equator it weakens, thus leading to lower temperatures.

It is not necessarily the case that there are significant difficulties here. Although the fundamental nature of light is not directly addressed in the curriculum, there is certainly a great deal of relatively direct instruction, and many opportunities for students to "pick up" more appropriate ways of reasoning about light. Indeed, when some students with problematic models were re-interviewed following the curriculum, their inappropriate models seemed to be replaced in favor of more appropriate models, even though they were not explicitly addressed.

However, in other cases, inappropriate knowledge seemed to remain, and this knowledge was incorporated into students' final understanding in undesirable ways. For example, several students had adapted fluid-like models to explain why the poles are cooler than the equator. In these explanations, light was seen to strike the equator, and then spread out over the surrounding parts of the earth (refer to Figure 4).

4.3 Summary of the Conceptual Dynamics analysis of the GWP

What I have described in the preceding sections is fundamentally a program in Conceptual Dynamics; I looked at the processes involved in refining and integrating multiple knowledge systems within a particular curriculum. More specifically, I expect that several features of the above analysis will generalize to future work:

What knowledge systems are involved?

- *Preliminary identification of resources*. I began with a preliminary attempt to identify the relevant knowledge resources that students will possess prior to entry in the curriculum.
- *Empirical exploration of resources.* I performed an empirical exploration of these resources prior to the instructional intervention. In future work, this empirical exploration will lead to the refinement of hypotheses concerning the nature of the prior knowledge resources.

How do these knowledge systems evolve during the intervention?

- *Empirical exploration of resources at several points during the intervention.* I performed an empirical exploration at later times during the curriculum unit.
- *Identification of difficulties.* The empirical work allowed for the identification of possible difficulties, places where resources would not be available or might not change as desired.

In the case of the GWP, I noted the existence of sparse knowledge for reasoning about climate phenomena, as well as the existence of possibly richer knowledge systems that were developed for the purpose of reasoning about other physical phenomena. Furthermore, I speculated about a pattern of evolution involving these two knowledge systems; aspects of richer knowledge systems were co-opted for reasoning about the phenomena that were the target of the sparse knowledge system. This story needs elaboration, and I certainly do not believe that conceptual change in task-structured curricula always proceeds in this manner. Nonetheless, we may well be seeing a pattern of learning that is relatively common.

In addition, we have seen suggestions of difficulties that might turn out to be typical of task-structured curricula. A task-structured curriculum will, both implicitly and explicitly, make assumptions about the nature of existing knowledge resources, and these assumptions might not prove to be correct. For example, the GWP assumes a certain critical level of intuitive understanding about the nature of light. This sort of assumption is not unique to task-structured curricula, but it might turn out that task-structured curricula tend to rely particularly heavily on a range of previously existing knowledge systems that are not directly addressed by the curriculum.

5. The plan of work

The goal of the proposed work is to devise new *theoretical frameworks and methodological strategies for characterizing conceptual change in task-structured curricula.* As I hope to have suggested in the previous sections, this program has the potential to contribute significantly to basic research concerning science learning. In addition, if we can, as I propose, improve our understanding of what students learn in task-structured curricula, we should be able to refine these curricula. Moreover, this improved understanding could ultimately help in broader reform efforts. If we want our reforms to gain wide acceptance, we must be able to document what students learn in our new curricula. As a first step toward getting this documentation, it is necessary that we understand, in relatively explicit terms, what it is that students learn.

In line with the goals of the CAREER program, I have in mind long-term effort. Thus, to a certain extent, the plan of work must develop over the period of the grant. Nonetheless, the broad outlines of the work are clear; I know the context for the work, and can take off from the promising pilot work discussed above. In this last part of the proposal, I will lay out the proposed activities.

5.1 Research context

As stated in the introduction, I will work in the context of existing curricular efforts, rather than engaging in any new design work. This, I am convinced, is a very advantageous feature of the proposed work. Although I believe strongly in the integration of design with more traditional research activities, design and development work often take more than their share of limited resources in an integrated effort. Thus, I believe that the dedicated effort I propose is merited.

The School of Education and Social Policy at Northwestern University provides a unique site for undertaking the work I propose. There are a variety of efforts currently underway at Northwestern to develop project-based science curricula. (Recall that a project-based curriculum is a type of task-structured curriculum.) Much of this work is being done under the aegis of the Center for Learning Technologies in Urban Schools (LeTUS). LeTUS is a large, four-way collaboration, among two research universities (Northwestern and the University of Michigan) and two large school districts (Chicago and Detroit). The curriculum development efforts of LeTUS are primarily targeted at middle-school students.

The LeTUS center provides an absolute wealth of school contexts in which to conduct this work. The curricula developed by LeTUS have already been used in about 40 middle-schools in Chicago. Working in collaboration with my LeTUS colleagues, I will identify classrooms in these schools that are appropriate sites for this work.

5.2 The core empirical work

My core empirical efforts will initially be based on techniques similar to those employed in the pilot work. In the pilot research, my collaborators and I began with some initial hypotheses concerning the knowledge system analysis. Then, in order to explore this knowledge, we designed clinical interviews to be conducted at a few points during the curricular intervention. In addition, although it was not emphasized above, classroom observations were carried out in tandem with the interviews. These classroom observations were particularly helpful in light of the fact that there were large variations, across classrooms, in how the curriculum was enacted.

	Research Activities	**Education Activities**
Phase I (Years 1-2)	• Selection and preliminary analysis of 2 LeTUS curricula. • Empirical studies of these curricula. Includes classroom observations, clinical interviews, and written tests. • Analysis of data and iteration in Year 2.	• Preliminary development and trials of video library starting in Year 2. • Start of project as research site for university students in learning sciences classes.
Phase II (Year 3)	• Empirical studies of 4 enactments of content-structured curricula.	• Iteration and continuation of above activities.
Phase III (Years 4-5)	• Studies of 3 additional LeTUS curricula. • Final analysis and paper writing.	• Iteration and continuation of above activities.

Table 1. Timeline for the proposed project.

As in the pilot studies, my empirical work will be primarily based in clinical interviews, supplemented by classroom observations. The five years of the project will be broken into three phases. During Phase I, which will last through the first two years of the project, my efforts will be focused more narrowly, and I will work to hone my empirical and analytical techniques. Then in Phases II and III, which span the remaining years of the project, I will attempt to apply these techniques more broadly. (Refer to Table 1 for an overview of the project timeline.)

5.2.1 Phase I (Years 1 & 2): Honing of techniques

During Phase I, I will focus my efforts around two LeTUS curricula, and I will conduct empirical work within enactments by two different teachers for each of these curricula. (One of the two curricula will be the GWP curriculum.) Approximately seven students in each classroom enactment will be selected for participation in the clinical interviews, and these seven students will be interviewed at several points during the curricular interventions. These interviews will be videotaped and transcribed. In addition, all of the associated classroom sessions will be observed by a member of the research team, and selected classroom sessions will be videotaped.

In addition to the interview data, all students will be given written tests at the beginning and end of each curriculum unit. The results of the tests will be used to supplement the interview data. In addition, through LeTUS it will be possible to administer test items to students at a wide range of sites. The LeTUS center already has a working infrastructure for administering pre- and post-tests to students at all participating schools.

5.2.2 Phase II (Year 3): Studies of content-structured interventions

In Phase II, I will begin the first of several attempts to apply the framework more broadly. During the third year, I will use the techniques developed in Phase I to study the learning that occurs in some more traditional content-structured curricula. Over the course of Year 3, I expect to look at approximately four separate classroom enactments. If possible, this work will occur with the same sites and teachers that were studied in Phase I.

5.2.3 Phase III (Years 4 & 5): Further extension; final analyses

In the fourth year, the primary empirical goal will be to further test the generalizability of the techniques and conclusions generated in the first two phases. During this year, I will look at enactments of three task-structured curricula that were not studied in Phase I, again selected from the LeTUS curricula. Though I will use the same techniques employed in the earlier phases, this analysis well be less-detailed, looking for breadth rather than depth.

Finally, in the fifth year, I will complete any remaining data collection and analysis. Significant time will also be devoted to the writing of papers and the planning of future endeavors.

5.3 Analytic techniques

As with other aspects of this work, the analytic techniques must largely be developed as the work proceeds. However, as a starting point, I can begin with the methods that have provided the foundation of the Knowledge Analysis program. These analytical methods involved the fine-grained examination of transcript data, with triangulation on the constituents of knowledge across multiple parts of the data corpus. As discussed in Sherin (in press) this is still a highly heuristic process, and poses particular methodological difficulties. In fact, diSessa (1993) devotes several pages to discussing the methodological difficulties associated with the Knowledge Analysis program. However, both Sherin (1996, in press) and diSessa (1993) provide guiding heuristics which can help.

At a somewhat concrete level, analysis of data generally proceeds as follows. I first look through a corpus of transcripts, selecting episodes for focus. Then, based on these focus episodes and the guiding heuristics, I develop preliminary hypotheses concerning the knowledge systems of interest. Finally, I iteratively view and review the entire corpus of transcripts, interpreting the transcripts in terms of the working hypotheses. After each iteration, the working hypotheses are refined. Ultimately, when it is presented in publications, the plausibility of any final account depends on a process of competitive argumentation, between the final hypothesis and alternatives, drawing on examples from the data corpus.

As I said, I will supplement the qualitative work with quantitative analyses of written test data. But, for number of reasons, the qualitative analyses will form the heart of the analytic work. First, because the research I am proposing is at a formative stage, a first step should be to get a sense, through qualitative work, of the learning phenomena that are the focus of this endeavor. More fundamentally, the goal of this work is ultimately to produce *descriptions* of types of systems and processes. In order to get a handle on these descriptions, at least initially, I must perform a qualitative exploration of the structure of the systems and the processes of evolution.

5.4 Educational activities

I see the learning sciences as a very young field. For this reason, it makes sense for me to constantly draw on the evolving understanding from my research — and the evolving understanding of the field — in my teaching. This perspective has led to a tight coupling between my teaching and research. In this section, I will first discuss my teaching activities at Northwestern, then I will discuss how the proposed work will provide new resources to support that teaching.

In my role as Assistant Professor, I mentor and teach students in four populations: Ph.D. candidates in our Leaning Sciences Program, masters candidates in the Learning Sciences, masters students who are obtaining their teaching credential, and undergraduates in various majors. Since I have arrived at Northwestern, I have taught the following four courses, which span this range of populations:

(1) **Learning and understanding: A cognitive science approach.** This course is an introduction, for undergraduates, to cognitive science and theories of learning.

(2) **New approaches in science teaching: Theory and practice.** This course is primarily targeted at masters students who are obtaining their teaching credential. It covers current theories of science learning, as well as new instructional practices.

(3) **Intuitive knowledge, conceptual change, and science education.** This is a seminar course, for Ph.D. and masters students, that covers current theorizing concerning the nature of prior conceptions in science, and how these conceptions change during instruction.

(4) **Thinking, learning, and acting with external representations.** This is a seminar course, for Ph.D. and masters students, that looks across a range of issues relating to external representations and education.

In addition, I am currently designing a new course to be taught, for the first time, in the fall of 2000:

(5) **Knowledge representation for the learning sciences.** This course will be a required course for all graduate students entering our masters and Ph.D. Learning Sciences program.

Because learning sciences research is inherently interdisciplinary, preparing graduate students poses some particular challenges. The graduates of the Learning Science programs at Northwestern are expected to be facile across a range of endeavors including: designing learning environments, designing and using educational technology, and studying student learning. Given the magnitude of this challenge, the current programs at Northwestern — and elsewhere — do an excellent job. We successfully train researchers who can design and study rich learning environments.

Nonetheless, based on my (admittedly limited) experience to date, there are some broad ways in which I would like to improve the learning of my students, especially graduate students. I believe that, at Northwestern and elsewhere, it is time to do better in training students to produce deep and theoretically precise accounts of the *learning* that occurs in the environments they design. For this reason, I would like to help train a generation of researchers that will work together to produce a shared, rigorous, and scientific framework.

My goals in this regard are modest. Within the confines of this project, I am only proposing to have an impact on students with whom I have direct contact. With these modest ends in mind, there are a number of ways in which the current project can help:

- **Student apprenticeships**. Where possible, I will involve students as research apprentices in the proposed project. These apprenticeships will include the Ph.D. students in positions funded by this grant, as well as undergraduates in various research apprentice positions.

- **Creation of a video library for use in teaching**. More significantly, it is my intent to draw aspects of this research project into my classroom teaching. One way that I will do this is to create a library of medium-length (10-15 min.) digitized video excerpts, drawn primarily from the clinical interviews. To start, this collection does not need to be especially large (in number) or technologically elaborate. A simple collection of digitized video files, organized by topics, can do the jobs I have in mind. With these video segments, I will be able to more concretely illustrate ideas that are addressed in my courses. In addition, the video library can provide the raw material for students to hone their own analytical skills.

  Our facilities at Northwestern University are particularly suited to this endeavor. Currently, the great majority of our video analysis is performed using digitized video. All of our videotapes are digitized and stored on a central server that is accessible from anywhere in the building. Thus, once I have created the library of excerpts, they will be easily accessible to all students during and outside of class. Beginning in the second year of the project, I will make use of the video library in courses (3) and (5) above.

- **This project as research site for my classes.** In the spring of 2000, the students in one of my courses (the "intuitive knowledge" course above) worked to design a clinical interview to be conducted in the context of my pilot studies. Several students conducted the interviews; then, as a class, we worked to analyze the resulting videotapes. From the point of view of my research, this effort did not turn out to be especially profitable. A significant amount of time was required for what, in the end, was a very tiny fraction of the research. However, the students found the exercise to be extremely valuable. Thus, as part of this project, I propose to continue and expand this type of integration with my courses.

It is very difficult to evaluate the success of graduate instruction of this sort (as I have been asked to do). I do not expect to perform formal evaluations, beyond a survey of students at the end of my courses. The best measure of success for activities of this sort is, I believe, the extent to which the ideas and issues addressed show up in the research activities of students. Do they cite the papers that we read in the course? Do they use the methodological and analytical techniques we examined? Over the years that this project unfolds, I will look, informally, at the research products of students. Given these observations, I will adapt my courses in the years that follow.

5.5 Role of the advisory board

For this project, I have assembled an advisory board with five members: Andrea diSessa (University of California, Berkeley), David Hammer (University of Maryland, College Park), Janet Kolodner (Georgia Institute of Technology), Richard Lehrer (University of Wisconsin, Madison), and Ricardo Nemirovsky (TERC). I will conduct meetings of the advisory board during the summers between each academic year of the grant, totaling 4 meetings. This will provide an opportunity for me to describe what has been done, and to gather the feedback and suggestions of the board. Letters of commitment from the board members are included as part of this proposal.

5.6 Project organization and management

I will have responsibility for overseeing all components of the project. The staff will include one graduate student during Years 1 and 2, and two graduate students during Years 3-5. These graduate students will aid in all aspects of the research, including the design and analysis of the studies. In addition, a research analyst will be employed to aid with the quantitative analysis of written test data. This is a quarter-time position.

6. Summary of prior research accomplishments

My prior research has involved both the design of novel interventions, and fine-grained studies of student learning. Here, I will discuss just a few aspects of the latter, particularly my studies of student learning within the "knowledge analysis" paradigm described above, since that work is most relevant to the current project.

My dissertation was a study of the use of symbol systems in physics instruction. In part of that work, I looked at how algebraic equations are used and understood by moderately advanced physics students. My results contradicted, in some very significant respects, accepted theories of physics problem solving, such as those given by Larkin (Larkin, 1983; Larkin, McDermott, Simon, & Simon, 1980) and Chi and colleagues (Chi, Feltovich, and Glaser, 1981). Although there is certainly some variety in the models espoused by researchers, I believe that all earlier models have, to a great extent, shared an important feature: There is no sense in which the details of an equation – the exact arrangement of the symbols that comprise it – are taken as expressing a meaning for the problem solver. According to these earlier researchers, equations are generally only invoked from memory by the problem solver, and then manipulated.

According to my observations, however, this image of problem solving is only accurate in some very limited regimes. Beginning at least with moderately advanced students, equations are not just written from memory, they are composed in order to express particular ideas. In short, there is a conceptual vocabulary in terms of which equations can be composed and understood. As part of my dissertation work, I mapped out this conceptual vocabulary. Furthermore, I found that, in large part, this vocabulary seems to build directly upon the intuitive conceptions of naïve physics. This has important implications for how we must understand conceptual change in physics; it appears that naïve physics knowledge evolves so as to play a fundamental role in physics expertise.

This dissertation work was, in part, supported by a Dissertation Year Fellowship from the Spencer Foundation. A large paper on this work was recently accepted to the journal Cognition and Instruction (Sherin, in press).

Currently, I am engaged in extending this work in a variety of ways. First, I am performing some related analyses within elementary mathematics. As part of my research with Karen Fuson and the Children's Math Words Project[5], I am looking at how students come to understand the world mathematically, and how they embody this understanding in mathematical expressions. A core result of this work is that, very early on, children develop a generative capacity for producing a particular type of mathematical abstraction of situations. The first papers on this work are currently being written.

In addition, I recently began work with Ken Forbus to develop a new computational tool for visual modeling. As part of this project, I am extending my prior work to develop accounts of how the world is parsed and schematized for the purposes of scientific modeling.

[5] This research is supported under award REC-9806020 (*Building algebraic, multiplication/division/fraction, and measurement understandings in urban classrooms in English and Spanish,* Karen Fuson and Bruce Sherin, PI's).