Part 1: Identification and Significance of the Innovation

Epistemological Engineering proposes to develop web-based simulations of the nature of science for secondary and postsecondary science instruction, and to offer "simulation services" to deliver and moderate these simulations over the web. In these simulations, students take on the roles of scientists, working in groups to advance a (simulated) scientific discipline. Elaborating on an existing prototype, we will develop new simulations, determine how effective they are in accomplishing their goals, and assess how well they can be commercialized in today's educational environment. This innovation facilitates the understanding of science as a way of knowing: the underlying logic with which scientists create knowledge in the interaction between theory and experiment. The project will also foster interactive, collaborative learning.

Part 2: Background and Phase I Technical Objectives

Even if students are not going to become professional scientists, they need to learn about science. As citizens of a technologically-rich democracy, we all need background in science to understand our history and our culture, and to make decisions about critical social issues.

Therefore we teach science in our schools and often require college and university students to take courses from a science department. But what is it exactly that students need to learn? Certainly there is *content*: the facts and theories of a discipline, the body of knowledge that we generally agree is correct. There are also *skills* pertinent to a discipline, including experimental methods and problem-solving techniques. But these are not enough. Many current policy documents (AAAS 1990, 1993; NRC 1996; and State documents such as California 2000) state that it is also important to understand the nature of science. *Benchmarks for Science Literacy* (AAAS 1993) puts it this way:

> When people know how scientists go about their work and reach scientific conclusions, and what the limitations of such conclusions are, they are more likely to react thoughtfully to scientific claims and less likely to reject them out of hand or accept them uncritically…The myths and stereotypes that young people have about science are not dispelled when science teaching focuses narrowly on the laws, concepts, and theories of science. *Hence, the study of science as a way of knowing needs to be made explicit in the curriculum*. (p. 3, our emphasis)

In listing what students should know about the nature of science, these documents include statements such as, "Scientific explanations…must be consistent with experimental and observational evidence about nature, and must make accurate predictions…" (NRC 1996, p. 201), and "scientific knowledge is subject to modification as new information challenges prevailing theories" (AAAS 1993, p. 7). These statements and others together describe a way of knowing: the hypothetico-deductive[1] model of scientific reasoning we are familiar with. Here, scientists create hypotheses to be tested by experiment; no hypothesis can ever be proved fully, only disproved or supported; a *scientific* statement is distinguished by its falsifiability. These documents go further to include what we might call the "practice of science," pointing out that scientists are part of a community; that they value skepticism and peer review; and even that "funding influences the direction of science" (AAAS 1993, p. 20).

Given broad agreement that these topics are important[2], how would we expect students to learn about them? Traditional curriculum materials—say, in a general science text—often address the

[1] In using this term, we're following Popper (1968), Driver (1983), and more recent writers such as Lawson (1995), saying that hypotheses come from the imagination, not directly from the data. But the terminology is not as important as the scope: this is the epistemological, "way-of-knowing" part of the nature of science, as opposed to, for example, experiment design and measurement.
[2] Further evidence for this and many other statements made here about the "nature of science" issues in education appears in Lederman (1992) and in McComas (1998).

nature of science with a section about the scientific method and with boxed vignettes throughout the text describing notable scientists or achievements.

This does not work very well. Exposed to traditional materials, students retain disturbing misconceptions about the nature of science and scientific knowledge. Even carefully-crafted curricula that use history to convey the nature of science show mixed results (Lederman 1992). What effect does that have in later schooling? Ayars (2004) had second-semester university physics students respond to a (fake) theoretical "paper"[3] that proposed an incorrect formula to describe how a cotton ball falls. These students—generally good problem-solvers, and competent in the lab—were to design and perform an experiment to support or reject the hypothesis in the paper. All students took measurements, but in their write-ups, *less than half* actually assessed the validity of the formula, and of that half, only 3/4 showed that it was wrong. So even if the students had remembered the "steps of the scientific method" or had digested a textbook vignette about the Michelson-Morley experiment, most could not use empirical evidence to evaluate a claim. Though they were successful products of science education, they couldn't actually *do* the science.

This is not surprising. After all,

> Scientific inquiry is more complex than popular conceptions would have it. It is, for instance, a more subtle and demanding process than the naïve idea of "making a great many careful observations and then organizing them." It is far more flexible than the rigid sequence of steps commonly depicted in textbooks as "the scientific method." It is much more than just "doing experiments," and it is not confined to laboratories. (AAAS 1993, p. 9)

So how should students learn about the nature of science? We know that active learning is more effective for basic concepts (e.g., Cooley 2004; Hake 1998; Redish et al., 1997; Lawson 1995; Karplus and Their, 1967). Why not adopt a more active approach if we want to teach about the nature of science?[4]

One way to learn actively about the nature of science is to become a scientist. This has traditionally happened in graduate school, and one could argue that you can't really understand science without the conceptual and procedural armamentarium from sixteen years of schooling. But we believe that this is a "false prerequisite"—that younger students can do it; and furthermore, that getting the big picture of science may motivate students to better understanding of content and make them more likely to persist in taking science classes. (Etkina et al., 2003)

But precisely how can they "become scientists?" In this project, students will do it in *simulation*. The class will become a community of scientists investigating a phenomenon or pattern. We mean this very broadly: students will participate in the development of an entire discipline; they will uncover a body of knowledge. They will see hypotheses come and go, supported or shattered by experiments they design; their theories will rise from accumulated research; and they will know first-hand that science is not a solitary enterprise, but depends on interaction and communication. `

At this point, it would be best to show an example, and connect these ideas to a classroom activity. To that end, we will describe a prototype we have developed[5] and tried with a group of ten physics teachers and that Adkins (2004) has tested with several classes of high-school

---

[3] Developed under our current award, SBIR Phase II DMI-0216656, a project to develop high-school physics activities that are rich in data analysis.

[4] A push for *inquiry* in science education has successfully taken this more active approach. It has spawned much research and curriculum materials (e.g., McDermott 1996; Etkina and Van Heuvelen 2001—see http://paer.rutgers.edu/pt3). Inquiry in this context usually means that students "describe objects and events, ask questions, construct explanations, test those explanations against current scientific knowledge, and communicate their ideas to others" (NRC 1996, p 12). Our notion of the nature of science includes this but is somewhat broader, encompassing the epistemological layer as well, following Lederman (1998).

[5] This prototype is itself an adaptation of earlier, pre-computer work (Erickson 1986).

students. Bear in mind that this prototype is on the “abstract” end of what we have in mind; we will discuss later how we can make it more concrete. We also defer some obvious questions (e.g., Why simulation? What happened?). These materials are adaptable to classes from middle grades through college; but this vignette is set in a high-school general science class.

Vignette: A Four-Color Universe

Your teacher tells you that your group is going to be a team of scientists at a university. Your team chooses a name (you pick “University of Lenny”) and uses your web browser to go to a particular URL. There you enter the name and click through introductory screens until you come to this one:

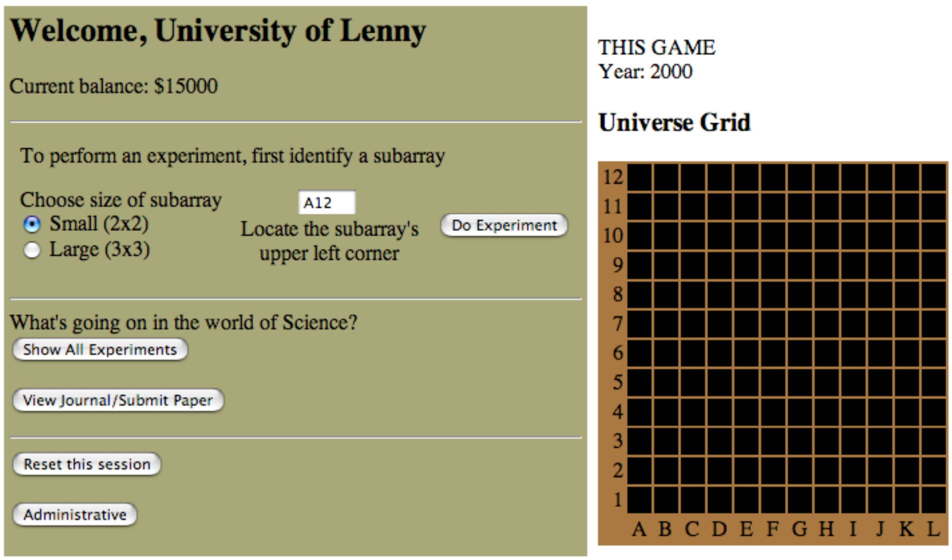

Here, the teacher explains that the universe is a 12-by-12 array of cells. Each cell holds exactly one of the following four colors: blue, red, green, or orange.

You can find out about the universe by performing experiments. There are two types of experiment: 3-by-3 and 2-by-2. When you perform that experiment—by specifying a 3x3 or 2x2 sub-array of the universe—you learn the *distribution* of colors within that small square, but not the colors’ actual positions.

You and your group, a little mystified, decide to look at the 3x3 sub-array in the lower left. You specify it—in the prototype, you enter the coordinates of the upper-left corner (A3)—and press **Do Experiment**. You learn that the 3x3 at A3 contains two red, two green, and five blue cells.

You go back to the game page and explore the other buttons. Under “What’s going on in the world of science?” the **Show All Experiments** displays the results you just got. The **View Journal/Submit Paper** button shows you this:

Read Your Journal!

Journal of Universe Studies			
1 articles		expts	cited
Snerd & Doogin 2000	The universe is half green and half orange.	1	-

You know Snerd and Doogin; they’re in another group. You click on the link and see their two-sentence “paper” and their data (the 2x2 at A12 is two oranges and two greens). But from your data, you know that their conclusion was wrong. You go back to the journal and click the link to **Compose a Paper for Publication**. This leads you to a screen where you can compose a very,

very short paper. You write, *Snerd and Doogin are wrong. In fact, blue dominates*. The referee (the instructor in this case) rejects your first attempt because you have neglected to reference Snerd & Doogin and your own experimental results. But you fix that (as shown) and resubmit.

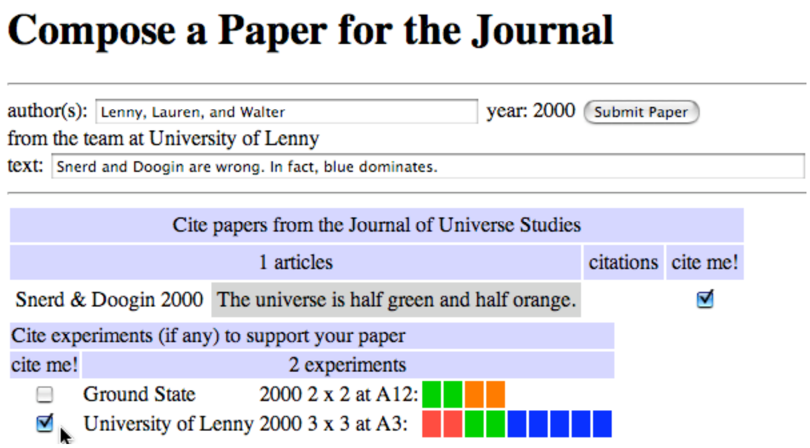

Now the journal shows your paper as well.

Flushed with success, you and your group perform more experiments (deciding together, with some discussion, what experiments to do), using overlapping 2x2s to try and figure out which cells in your corner have which colors. Lauren thinks the five blues are probably in the shape of an "L." Lenny thinks they're a "plus." Personally, you think they're an "X." But on submitting your fourth experiment—with only the Lenny Conjecture disproved—the computer informs you that you cannot afford the experiment. Talking with the other groups, you discover that experiments cost money—and that 2x2 experiments are more expensive than the 3x3s.

Twenty minutes have elapsed, and the instructor calls time. She leads a brief discussion of what has been discovered so far (a third group believes that the universe has a line of mirror symmetry on the main diagonal; a fourth has discovered that the center 2x2 is all red) and explains the schedule and criteria for new funds being added to your group's account. The next discussion will be on Friday; for now, she returns to the regular course content. At the end of class, she reminds you all to stay in touch with your group, and log back in—from home if necessary, but you can use the media center—to perform more experiments, read the Journal, and submit more papers.

Two weeks later everyone has gone though several funding cycles and *J.Univ.Stud.* has 33 papers. There is substantial support for that diagonal symmetry theory, and one group seems to think that the universe also has point symmetry about the center (though they don't use that terminology). Between various papers, groups claim to know twelve particular cells for sure, though class members disagree about the logic. In the half-hour "symposium" that day—the longest amount of class time devoted to the simulation since Day One—a debate rages over whether it is even possible ever to know the exact contents of all the cells.

What Happens When Students Use This Simulation?

In a high-school class (Adkins 2004), the teacher reported that even in its prototype state, the activity (followed by discussion) was able to establish the following concepts within a single class period:

- Sometimes scientists must work with incomplete information.
- A planned, systematic experiment reveals more than trial and error.
- Collaborative efforts can capitalize on investments already made elsewhere.

- A hypothesis is testable if an experiment can be designed to probe its consequences (this is particularly important as most high school students do not move beyond ‘a hypothesis is a testable educational guess’)
- A theory can guide the development of hypotheses.
- More detailed experiments are more expensive.
- There is not an unlimited amount of money available for research.
- Sometimes all you can do is look for patterns in other people’s data.

In the vignette, and in the real classroom, the students began to understand what it means to do science. They explored a new field of knowledge. They wrote papers and cited the work of others. They decided which experiments to perform. Some theories rose to prominence and were later discarded as simplistic and naïve. This is different from learning concepts and skills. In itself, the simulation does not teach the nature of science—it will take some attention from the teacher to turn the raw experience into learning, and we will discuss that. But the simulation gives students the experiences they will need to refer to; that is, now they have *been* scientists, and this is more powerful than reading about them.

Adkins’s classes are typical of those with which we work in that they are diverse ethnically, economically, and in intention. These are not all college-bound students; *all* students need this.

Why simulation? Why not be even more hands-on?

Three reasons: first, time. To get at the epistemological issues, students need to create hypotheses and have them shot down; they have to design experiments to distinguish between rival theories; and they need to do this repeatedly in order to understand how it all fits together. To do this with real equipment and materials is usually prohibitive in time, cost, and student distraction.

The second reason for simulation is freedom of topic. In a simulation, you can study anything whether it really exists or not; and whether or not it is impractical—or unsafe—to study using real materials. It is hard to use real materials and have students actually create hypotheses that make sense. [6] The balance between what the students know, the actual science, and the vagaries of equipment has to be perfect.

The third reason to simulate is that simulation focuses on particular features of the system. Simulation *distills*. It does not include everything. For example, in the vignette, research groups do not have to apply for funding; maybe they should. We can test such features in this project. Our early experience, however, suggests that even the very simple activity is extremely rich.

The Importance of Curriculum

Simply having the experience of being a scientist in a simulation does not guarantee that you learn about the nature of science. The learning generally gets “cemented” during reflection, often a debriefing session or sessions, where the leader—the teacher—facilitates a discussion or assigns writing to help the participants recognize *explicitly* what they have been doing. While Adkins, the field-test teacher, is experienced and artful, we need a way to help more teachers make that reflection happen. This help will take two forms: a printed or video guide of advice and examples, and structural features of the simulation itself (such as symposia or “seed” papers in the journal) to ensure that the “right” things happen.

[6] See Lederman (1998) and Etkina and Van Heuvelen (2001) for different thoughtful approaches to learning about the nature of science that use real materials.

Underlying Structure of the Simulation

Our simulation has four indispensable features:

- As a participant, you do what scientists do: observe; hypothesize patterns or models that could give rise to the data; and test predictions that derive from the models.
- Resources are scarce. Each action is expensive. Consequently you must choose your experiments carefully, e.g., to distinguish between rival hypotheses.
- It is social. You are part of a community of investigators all working on the same broad problem, so sharing data and insights is useful.
- It is time-efficient. Students perform many observations or experiments in very little time, so the teacher can include nature-of-science lessons without sacrificing the existing syllabus.

What does our simulation look like "under the hood?"

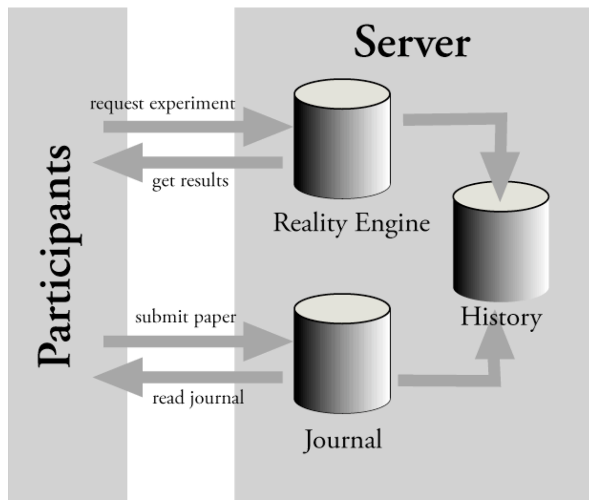

To create a simulation, an instructor chooses a scenario and specifies its parameters. Each individual or group connects to our server to interact with the simulation. Each instance of a simulation has its own "reality" that participants probe by performing experiments. A "reality engine" on the server receives these experiment requests and returns results to the participants. Likewise, each simulation has a moderated journal through which groups can communicate with one another. As they do with experiments, participants submit papers to the journal though the web interface, and read the journal online.[7] A "history" file (plain text in the prototype, but probably a database as we develop it) records all of the events for review by the instructor for student assessment, by project staff as part of our internal evaluation, and by other investigators interested in studying the process of students' learning about the nature of science.

Instructors or institutions pay to have us create, host, and maintain the simulations; this could be by-the-game or involve a subscription, for example, to run unlimited games for a year.

The Reality Engine

In our vignette, the reality engine is simple. It stores the 12-by-12 pattern, and responds to experiments by giving participants the distribution of colors that fall within a sub-array of the pattern. This is appropriate for such an abstract "reality," and creates surprising richness in our prototype. But this design is flexible and modular: for example, we can make the simulation address discipline-specific content.[8]

Suppose we're doing a unit on the Copernican revolution, one of the seminal events in intellectual history. The engine could maintain a model of a simple solar system with circular, coplanar orbits. When students request an observation, they receive simple pictures of a night sky, and

---

[7] Note that—just as in real life—there is no direct connection between the scientific journal and objective reality. This is an important lesson in itself, and reminds us that students should be assessed not on the *correctness* of their papers, but only on whether they make sense given the data that is available.

[8] A benefit of abstraction, however, is that it starts all students at an equal level: nobody already knows about the pattern of colors.

tools to measure the positions of objects. Students identify planets (because they move relative to the fixed stars), and eventually figure out which are closer to the sun or farther away, observe retrograde motion, and construct their own model of the solar system.

That in itself would be challenging for many students, and has clear connections to content. But the simulation could let the instructor specify an alternative solar system. Perhaps observers are on a moon (so planets would wobble in the sky), or the sun might not be the central object of the system. The students—the community of scientists—would have to come up with a model to explain their observations; and being already steeped in our Sol-centered Copernican scheme, would first struggle to shed it. Students could experience a "revolution" in their own views—a paradigm shift (as in Kuhn 1962).

So far, we have used astronomical examples for convenience and clarity[9]. But imagine, briefly, taking core samples in search of oil; receiving plague victims in a hospital and figuring out how to treat them; discovering an alternate-universe periodic table; or trying to find good dosages for drugs in combination therapy. The Phase I work includes sketching these possibilities more fully.

Let us look briefly at three more possibilities for the reality engine:

- The vignette's observational results require logic and good recordkeeping, but not quantitative data analysis. We will also create models and reality engines in which the results of an experiment are numerical; students would use an external application such as Fathom (Finzer 1999) to analyze the data. Fathom is especially easy to use here because of its ability to seamlessly import data from the web.
- The reality engine can simulate instrumentation issues such as measurement error and the capabilities of different equipment (e.g., the resolution of a spectrometer or telescope).
- The astute reader may object to the "patterns" vignette because students study the pattern, not the reasons behind it. Indeed, a theory is often an *invisible* set of reasons (e.g., universal gravitation) that manifest themselves only in their visible consequences (elliptical orbits). We intend to devise such systems; making them tractable but not transparent is the challenge.

The Journal

In the tests we ran, the journal became the nexus of the community: it was where you went to find out what was going on—even when the other groups were in the same room. The journal is a record of the history of the scientific enterprise, and (as in professional science) an obvious place to turn when we're looking for how to assess student understanding. Do their papers propose genuine hypotheses? How do students respond—using the journal—when data contradicts a current theory?

Teacher-participants mentioned the journal as the most important component of their experience using this simulation. Also, the possibility of running a simulation with groups from other classes, other schools, or other countries is intriguing—and the journal makes that practical.

Technical Objectives

The project's long-term goal is to create commercially viable, web-based "simulation services" and related materials that science educators can use in lessons that help students learn about the nature of science by letting them take on the roles of scientists. Meeting the following objectives during Phase I will position us for an effective Phase II proposal and eventual successful commercialization.

1) Enumerate a wide variety of possible scenarios—both in terms of content and of class logistics—in which we could plausibly implement our simulation ideas. We would do

[9] The examples have also been *observational*; genuine experiments should be possible as well.

this with the collaboration of our existing network of teacher-consultants and with the help of advisors.

2) Prototype a small number (about 3) of these scenarios as a proof of concept. Each scenario should have
    - A “reality engine” and interface that gives students experimental results, designed for as much modularity and re-usability of code as possible.
    - A journal as the primary communications medium among students
    - A journal referee feature (missing in the current prototype)
    - A password system for students and teachers (also missing).
    - Administrative functions so that the teacher can specify simulation parameters; add and delete users and journal articles; award new money, etc.

    To meet this objective, we will also need to overcome various technical challenges as they arise, including deciding on a development environment; incorporating a database (e.g., mySQL) for storage and access control; and so forth.

3) Field-test at least one of these in real classrooms in a series of six or so small, less formal field tests. This will let us “iterate” and try new versions. The reasons to iterate are both technical (how the software worked in the classroom) and pedagogical (whether the features of the reality engine, the journal, and the curriculum induced students to demonstrate nature-of-science understanding).

Research Questions

Here are questions that will guide our work; the numbers in parentheses indicate which technical objectives address each question.

1) What are some examples of science content that are amenable to this treatment? (1)
2) What attributes of these simulations make them practical for real classrooms? What do teachers need to know in order to use such simulations effectively? (2, 3)
3) What can we learn about students’ understanding of the nature of science from their work in the simulation? (2, 3)
4) What role do the network/cooperative/community aspects of the simulations play in the student experience, particularly as facilitated by the journal? (2, 3)
5) What features of the simulation and its accompanying materials best and most efficiently elicit student work that shows understanding of the nature of science? (3)
6) What development and server environment(s) are most appropriate for creating, running, and maintaining these simulations? (2)

We cannot answer these questions fully in Phase I, but we can at least map out the territory and refine the questions. For example, for questions (2–5) above, we will observe students and teachers as they use our prototypes; conduct informal interviews of teachers and students; and study the “artifacts” (the journals) from those lessons. We will also consider what other measures might be appropriate (e.g., questionnaires).

Part 3: Phase I Research Plan

The timeline (next page) shows our initial plan for accomplishing the three technical objectives, and thereby address the research questions. Much of the technical discussion of each objective has already appeared in Part 2. But how, specifically, do we plan to carry them out?