Becoming Scientists: Employing Adaptive Interactive Narrative to Guide Discovery Learning

James Matthew THOMAS, R. Michael YOUNG
Liquid Narrative Group
Department of Computer Science
North Carolina State University

Abstract. We propose to develop a domain-independent framework to automatically generate guidance for students in an exploratory scientific learning environment, where this guidance takes the form of interactive narrative. Our hypothesis is that this form of guidance yields statistically significant improvements in learning effectiveness and efficiency over unguided exploratory learning or direct pedagogy.

Keywords. Narrative, interactive, learning, guided, adaptive, discovery, science.

1. Introduction

Problem #1: Science is hard [1]. Scientific knowledge consists of a unimaginable volume of facts and measurements buttressed by ever-evolving theories and models of unlimited complexity. The span of scientific exploration ranges from the infinitesimal to the intergalactic, instantaneous to epochal, from the beginning until the end of time. Yet, many scientific principles are evident in everyday experience and many findings are consistent with the most naïve preconceptions. When we teach science to school children we are asking them to construct an extremely complex knowledge base of multiply-connected concepts and facts. Furthermore, these entities and their interconnections have multiple representations at varying levels of abstraction. And that's not even the hard part. The hard part for the budding student-scientist is to learn to participate in a process of scientific inquiry that appropriately challenges and refines this complex mental web works. Science educators have long struggled with the competing tensions between teaching the facts scientists have collected, teaching how these facts relate to underlying models, and teaching the process by which models and theories are created and evolved.

> "Simply telling students what scientists have discovered, for example, is not sufficient to support change in their existing preconceptions about important scientific phenomena. Similarly, simply asking students to follow the steps of 'the scientific method' is not sufficient to help them develop the knowledge, skills, and attitudes that will enable them to understand what it means to 'do science' and participate in the larger scientific community." [2].

The cognitive scientist Jean Lave has argued that learning should be assessed not merely by a change in minds, but by 'changing participation in changing practices' [3]. For all of the effort that has been put into understanding knowledge construction, at the end of the day what is most important is how knowledge aids a student's functioning. This paper proposes a strategy for scientific instruction where the primary emphasis is on students "becoming scientists". As students develop understanding of scientific processes, they will hopefully become more effective in growing and refining their entire scientific knowledge base. Our hypothesis is that this approach will yield students who are better scientists because they have become aware of the robustness, flexibility, and usefulness of their acquired scientific understanding. Furthermore, we believe these students will be more capable scientists because they have participated in interactive stories in which they have seen themselves succeed as scientists.

2. Related Work

The inspiration for the scientific learning strategies advocated in this paper is derived from four sources. First, the field of scientific discovery learning informs many of the details of the requirements and difficulties in the teaching of scientific process. Second, previous work in designing intelligent tutoring systems describe techniques to adapt instruction to a model of student thought processes. Third, investigations of the active learning principles leveraged by commercial video games present a compelling picture of architecture that promotes high levels of student engagement and efficient knowledge acquisition. Fourth, interactive narrative research describes a domain-independent framework for automated, adaptive generation of learning situations that tie together these other threads.

2.1. Scientific Discovery Learning

In *scientific discovery learning* [4], students construct their understanding of a scientific domain by generating hypotheses and testing them as a scientist would. A prime motivation for discovery learning is to promote learning that is deep and conceptual in contrast to shallow memorization of “inert knowledge” [5, 6]. When students learn with understanding, they remove misconceptions, gain improved knowledge retention, and are able to transfer the knowledge to new situations and explain the knowledge they have acquired in their own terms [7].

Naïve approaches to discovery learning that provide high fidelity computer simulations or real classroom experimental apparatus but lack substantive guidance have repeatedly been found insufficient to produce positive changes in learning outcomes [8, 9]. **Problem #2: Scientific discovery learning requires guidance.**

Situated within a process of scientific discovery learning, each individual experiment is fraught with opportunities for unguided learning to go astray [2, 4]. First, is the **problem of preconception bias**. "With respect to science, everyday experiences often reinforce the very conceptions of phenomena that scientists have shown to be limited or false, and everyday modes of reasoning are often contrary to scientific reasoning" [2]. Second, students have **a problem generating good hypotheses**. One study found only 42% of the hypotheses being generated by college students were properly formed in terms of variables and a relation between them [10]. Others showed that students have difficulty navigating between theoretical variables and the variables that can be manipulated in the simulation [11], and tend to avoid generating hypotheses that have a high chance of being rejected [12].

Beyond hypothesis generation, students have an entire set of **problems with experimental design**. Minimally guided students often engage in experiments from which no meaningful conclusions can be drawn. Students vary too many things in one experiment, or view the process as a challenge to produce a desirable outcome. Upon completion of an experiment, students have the **problem of interpreting experimental data.** **“**Students of all ages show a tendency to uncritically infer cause from correlations” [13]. One study reported that students failed to draw the right conclusions from disconfirming experiments in 56% of observed cases [14]. Finally, when the interpreted data has meaningful results, students have a **problem evolving a conceptual model or theory in light of experimental results**. Donovan and Bransford describe this phenomenon as, “Middle school and high-school students typically think of models as physical copies of reality, not as conceptual representations. They lack the notion that the usefulness of a model can be tested by comparing its implications to actual observations.” [2]

Beyond individual experiments, the full process of discovery learning creates the **problem of students regulating their own learning**. Specifically, "it is frequently reported that successful learners use systematic planning and monitoring, whereas unsuccessful learners work in an unsystematic way [4, 15, 16]". Numerous studies show students who are capable of self-explaining their learning outcomes learn better [17].

Summarizing many of these findings, de Jong and van Joolingen point to computer aided guidance as a possible answer “The promise offered by inquiry learning is tempered by the problems students typically experience when using this approach. Fortunately, integrating supporting cognitive tools with computer simulations may provide a solution [4]”.

2.2. Computer Guided Scientific Discovery Learning

Although guidance has been a feature of automated discovery learning environments going back to Smithtown in [18], nearly two decades ago, there is little established consensus for the best mechanisms and structure for such guidance. In fact, recent studies argue for the less ambitious goal of standardizing the terminology used to describe various forms of guidance [19] and the methods by which systems are assessed [20]. Still, many of these systems aim for the common goal of approximating the guidance offered by human tutors.

One-on-one human tutoring has been shown to be much more effective than traditional classroom instruction. Furthermore, empirical studies have confirmed that even unaccomplished human tutors with little or no training in pedagogical techniques produce a learning effect size of .4 standard deviations (colloquially referred to as "sigma's"), beyond traditional transmission-based classroom. This effect size translates to a performance improvement of approximately a half a letter grade [21]. Well-trained, skillful human tutors produce effect sizes of up to 2.3 "sigma" [22]. Computer based, intelligent tutoring systems (ITS) have shown effect sizes of approximately 1 sigma up to 1.75 sigma. Since the vast majority of human tutors lack formal training in the skills of tutoring [23] it is not unreasonable to claim that Intelligent Tutoring Systems (ITS) generally outperform both classroom instruction and typical tutors.

However, several studies have shown that the learning advantage of both ITS and tutoring disappears if one is only measuring "shallow" or "inert" knowledge. Although a wide variety of teaching methods enable students to pass a test of memorized facts, few provide students a deep understanding of the subject matter that can be applied what they have learned to unfamiliar situations. Therefore, it is reasonable for ITS and scientific discovery learning systems to emphasize learning in terms of deep conceptual understanding.

Obviously, many discovery learning systems do not use intelligent tutors, and many ITS systems do not address discovery learning. One system that is particularly strong in both areas is an "open learning environment" in the domain of mathematical functions called ACE (Adaptive Coach for Exploration) [24]. Through ACE "learners can acquire deeper, more structured understandings of concepts in the domain than they would from a less active style of learning", according to its creators. In addition, "a substantial limitation of open learning environments is that their effectiveness for learning strongly depends on learner-specific features", such as passivity, cognitive ability, academic achievement, generating appropriate hypotheses, interpreting and generalizing results, coverage of domain concepts, etc. Unguided discovery learning environments simply do not work for learners who lack these skills and attributes. Coaching, therefore, is an attempt to broaden the population that can be served through discovery learning. The ACE student model rests heavily on Bayesian Belief Networks tightly bound not only to the domain of mathematical functions but also to the particular manner in which this domain is explored in the ACE system. Although the authors showed that this coaching was helpful, it is difficult to derive a general methodology to extend their student model to other domains.

In summary, although guided scientific discovery has shown some good successes, most solutions suffer from one or more of the following problems:

- **Problem 1:** Science is hard. Extensive scaffolding is required.
- **Problem 2:** Scientific discovery learning requires adaptive guidance.
- **Problem 3:** The process of scientific discovery has specific problems:
    - Preconception Bias
    - Hypothesis Generation
    - Experimental Design
    - Model Adaptation
    - Data Interpretation
- **Problem 4:** Regulation of Learning
- **Problem 5:** Overly specific solutions

2.3. The learning structures of games[1]

Meanwhile, the $12.5 billion video game industry is deeply invested in employing effective learning techniques for guiding users through computer-driven exploratory environments.

> "You cannot play a video game if you cannot learn it. If no one plays a game, it does not sell, and the company that makes it goes broke. Of course, designers could keep making the games shorter and simpler to facilitate learning. That's often what schools do. But no, in this case, game designers keep making the games longer and more challenging (and introduce new things in new ones), and still manage to get them learned," [25].

Game writer Hal Barwood has described the ideal learning path in a game as one that alternates periods of intense challenge with periods of recuperation. The challenges must get ever more difficult to avoid boredom,

---

[1] For the purposes of this paper, the phrases "computer game", "video game", and "game" are intended to be synonymous.

but the learning path must not be so steep as to lead to frustration. [26]. James Paul Gee identified thirty-six distinct, but general learning principles exemplified in video games in an in depth investigation [25]. We believe that many of these principles point to general solutions for some of the problems we have outlined with guiding scientific discovery learning.

For example, Gee shows five or more learning principles in games that address the need for scaffolding we articulated as problem #1. In the early phases of the game the learner is sequestered in a sub-domain of the full game where and harmful choices and penalties are introduced gradually and with a lot of specific guidance. Challenges carefully ordered so that the user is presented with a concentrated sample of "the most basic and important actions, artifacts, and interactions that the player will need to deal with throughout the game" early in the learning process. Users learn first in a simplified subset of the full game, where there is usually no time pressure and minimal penalties for mistakes. There are usually either no enemies or only minimally demanding enemies to attack the players. Although aspects of these scaffolding techniques have appeared in some discovery learning environments, few would argue they have been employed as enthusiastically as in the game community.

Games have also addressed the second problem we articulated – the need for guidance. Although the physical devices players use to interact with games are increasingly sophisticated, in terms of both the density of buttons and the need for simultaneous use of multiple buttons and joysticks, few gamers even bother to read the written instructions provided with a game before diving right in. The reason is that game designers have become very adept at integrated instruction into game play. Training levels provide the learner with just the information they need just when they need it. In-game reference materials clarify and remind users of their options. Unfortunately, most of this guidance is one-size-fits-all. There is little ability to expand or contract the amount of training an individual requires in a game, one must step through it all, skill by skill.

Beyond problems #1 and #2, the learning techniques of games are certainly general in a way that solves problem #5, and there are some specific ideas that could be brought to bear on problems #3 and #4. However, the potential gain from adapting these techniques to scientific discovery faces two imposing limitations. First, there is a little in the way of user adaptation in games. Gamers measure themselves by how quickly they can complete a given game; the games are not measuring the users to help them proceed more quickly. Second, although Gee and other students of video games have done a great job of collecting examples of good learning principles in games, these are descriptions rather than prescriptions. They do not lead to any obvious blueprint for **generating** game-inspired learning principles within an instructional environment. Given that ITS and scientific discovery environments are already very time-consuming and expensive to build, it seems that we should endeavor to develop some kind of reusable, general framework for solving these problems.

3. Interactive Narrative as a General Guidance Mechanism

We claim that narrative, specifically, interactive narrative can serve as the basis for a general framework that provides guidance in scientific discovery environments.

3.1. Narrative – Legendary Generality and ~~Reruns,~~ ~~Rehash~~, Reuse

Bruner has argued "the process of science-making is narrative" [27]. He claims that the history of science “can be dramatically recounted as a set of almost heroic narratives in problem-solving”, and advocates that “our instruction in science from the start to the finish should be mindful of the lively processes of science-making, rather than being an account only of 'finished science' as represented in the textbook” [27]. Narrative certainly is a sufficiently general idiom for encompassing the process of scientific discovery. As many have noted, some definitions of narrative have been stretched beyond the point of usefulness, so that “everything” becomes narrative. So the question becomes, “can we develop a set of definitions for narrative constructs that encapsulate the types of guidance we would like to provide in scientific discovery?” On first reading this seems an ambitious goal. Yet commercial games are already using narrative as a mechanism for guiding users.

3.2. Example: Narrative Guidance in Half-Life® 2

Released in 2004, Half-Life® 2 is one of the most successful computer games of all time. It was “named Game of the Year by over 35 organizations and sold over four million copies worldwide”[28]. Half-Life® 2 was recognized by the gaming community as especially rich in story. Indeed, the lead writer for Half-Life® 2, Marc Laidlaw, is a passionate champion for increasing the depth of story in games.

Laidlaw has described how Half-Life® 2 considered the options afforded by narrative in the evolution of the design of one particular scene in the game. In this scene, the player's character, Gordon Freeman, comes upon a beach with hidden dangers. The programmers who had constructed the scene told Laidlaw's game writers that they needed to "Tell player to stay off sand" [29]. The programmers set up the scene with unseen "antlions" beneath the beach. If the player walks across the sand rather than jumping from rock to rock, the antlions will awaken and devour him. The first idea for guiding the player was to post a sign "Warning: Do not fall in sand". The next idea was to have a new character sitting at the entrance to the beach, shaking his head and saying: "Poor Tobias, he forgot about the antlions under the sand."

By the time the game was finished, this simple admonition had grown into a short interactive story. Gordon comes up on two characters at the beach, one at the entrance, and another, named Lazlo, laying injured on one of the rocks surrounded by sand. The first character shouts to Lazlo "Don't move! Stay on the rocks! Stepping on the sand makes the antlions crazy". According to Laidlaw, if Gordon (the player) is quick enough, he can intervene and take Lazlo back to safety. For the average player, however, what happens is that Lazlo continues to move, steps on the sand and is eaten. The other character then remarks, "Dear God! Poor Lazlo, the finest mind of his generation come to such an end" [29].

Why go to all the trouble of creating this narrative instead of just posting a sign in the sand? We can see three advantages provided by the story. First, following the prime directive of creative writing, it "shows" rather than "tells". Showing the danger sends a much more powerful message that is more likely to be heeded by the player. Second, it better preserves the player's sense of diegesis, or immersion, in an alternate reality. The sign in the sand is not the kind of thing that normally happens in the Half-Life world and it would immediately signal to the player that it was a message from the real world of the game creators. One of the key goals of the designers was for the player to "become" Gordon Freeman, and for that to work the player must be reminded of the outside world as infrequently as possible. Being "unfairly" devoured by hidden antlions and having to restart the game would obviously be a worst-case instance of extra-diegetic activity. A third benefit of the story is that it gives the player a choice. Shall he be heroic by attempting to save Lazlo, or cautious? Each choice is a chance for the player to answer the question "What would Gordon Freeman do?" Laidlaw makes the point that most people are not heroic in their everyday lives. Each time we offer a player a choice to be heroic, we are actually providing an invitation to define and stay in character.

3.3. Adaptable (Interactive) Narrative

The field of "interactive narrative" studies the automatic generation of stories within virtual worlds in which human users interact with one or more computer controlled agents [30, 31, 32, 33, 34]. The advantage of interactive narrative is that it can adapt to different choices the user may make in the environment. A persistent challenge in these systems is the narrative paradox: "how to reconcile the needs of the user who is now potentially a participant rather than a spectator with the idea of narrative coherence." [35]. More specifically, "two key challenges posed by narrative-centered environments for exploratory learning are (1) supporting the hypothesis-generation-testing cycles that form the basis for exploratory learning, and (2) orchestrating all of the events in the unfolding story to support appropriate levels of student motivation, engagement, and self-efficacy for effective learning" [36].

Our research builds on one approach of balancing narrative and user goals first described in the Mimesis system [37]. Mimesis generates plans for actions of agents in a story world based on hierarchical task decompositions and discrete causal requirements. Mimesis simultaneously solves for plot coherence and character believability, mediating the actions of the human user through a combination of accommodations (when re-planning can preserve narrative goals) and interventions (where threats posed by user action would be fatal to narrative goals). These mediation techniques can be proactive or reactive. In addition, Mimesis allows search space heuristics to be customized [38], to allow the system to judge the relative "goodness" of one story plan over another. The narratives produced by Mimesis have been shown to correspond well to the cognitive models readers use in comprehending story [39].

4. A Model For Applying Narrative Guidance to Scientific Discovery Learning

4.1. Vision/Goal

Our claim is that narrative provides a framework with sufficient power and generality to encapsulate scientific discovery learning. Note that we do not claim that a given experiment or scientific concept predetermines a particular narrative representation. It is not our intent to automate the divination of stories based on some intrinsic structure of scientific discovery. Our intent is to provide a useful apparatus of reusable devices to allow human curriculum designers to define narrative elements and goals. The system then automatically generates narrative that can be executed in a scientific discovery environment situated within a virtual world. Furthermore, both the story world both informs and is informed by a learner model at run-time so that the narrative adapts to the needs of the user. We believe this can be accomplished on our existing Mimesis architecture, through set of conventions we impose on the literals/conditions, operators/actions, and objects that compose our planning libraries. However, we are at a very preliminary phase in considering implementation strategies for these goals and we retain the caution that we need more substantive changes to our planning architecture to accomplish these goals.

4.2. Narrative "devices" for problems of scientific discovery

Narrative is the organizing principle of the guidance our system offers to students. In order to build narrative plans that take account of actions occurring in the story world, these actions must be articulated in a form amenable to automated planning. Any action in the world we wish to intelligently control must be specified in terms our automated planner understands: preconditions and effects. This "action library" is a restricted case of a traditional planning operator library, where the restrictions take the form of conventions that are added on to definitions of the planning operators and literals. Note that any conventions that are applied to the library elements may be enforced by a system external to the planner (e.g., a GUI used to construct the planning library). An example of a likely convention is to require that each operator define whether it may be invoked by the user, by automated agents (a.k.a. NPCs, or non-player-controlled characters) or by either an NPC or the user. This requires in turn that the objects in the story world be classified so that we know which objects pertain to agents, and which agents are NPCs or human users. All of this can be accomplished within the descriptive power of existing of the STRIPS-like [40] language on which our Mimesis planning algorithm operates. Note also that these actions may be abstract and thus decompose into other actions. This allows us to support hierarchical definitions of stories.

4.3. Concept paths

Curriculum designers can use "Concept Paths" to describe a progression of related learning tasks designed to achieve a particular learning outcome. These can be defined more powerfully and abstractly than many previous descriptions of learning tasks, where simply "exploring" or being shown a node is equated with understanding a concept. Instead, the "prerequisites" and effects of a learning task can be expressed as literals of propositional or modal logic (e.g., "believes(?user, rainbows-are-water-vapor)). These propositions can in turn drive Bayesian models of student beliefs.

We would like to allow a curriculum designer to specify a progression of classes of progressively more complex learning tasks that allow the student to iteratively deepen their understanding of a concept. These paths can describe one or more templates that capture the minimal and maximal variations within each class as well as evaluation criteria for the student to establish mastery at each level. We can make use of the multiple levels of abstraction afforded by definitions of operators in our system to create hierarchical concept paths.

Ideally, each concept path describes a narrative of discovery, a story in which the protagonist is the user in the role of making a scientific discovery. In the best case, learning becomes in Gee's words "a cycle of probing the world (doing something); reflecting in and on this action and, on this basis, forming a hypothesis; reprobing the world to test this hypothesis; and then accepting or rethinking the hypothesis". The understanding students gain through this process is situated in their experience and can best be evaluated in terms relevant to their experience. For this reason, the criteria for concept mastery should be established with and through the set of learning tasks for each particular concept.

Additionally, there are some scientific phenomena that cannot be understood as a simple sequence of facts and experiments. Many of the more difficult and deeper conceptual learning that systems like ours purport to be their main focus are actually sets of interrelated phenomena. Concept paths can provide a way of describing these interrelationships so that they can be taught and assessed. Consistent with Frank Herbert's "First Law of Mentat": "A process cannot be understood by stopping it. Understanding must move with the flow of the process, must join it and flow with it."

4.4. Learning Tasks

Concept paths are decomposed into one or more learning tasks. A learning task is an abstract task consisting actions that eventually ground out into "world state actions" but have additional pedagogical preconditions (prerequisite knowledge) and effects (expected learning outcomes). These operators and literals may be tagged as "pedagogical" and mixed in with the world state library, or maintained in a separate library from those of the "world state".

These operators will likely be specified at a finer grain size than is traditionally the case in curriculum design. As with concept paths preconditions and effects are used to describe operator features in terms of cognitive processes applicable to the user

4.4.1. Example: Learning Task Realized As Experimental Design

Each experiment the student performs is a step in a process of scientific discovery. As shown earlier, students face varied challenges in navigating this process, categorized as "Preconception Bias", "Hypothesis Generation", "Experimental Design", "Model Adaptation", and "Data Interpretation".

An individual experiment can be specified directly in a format conducive for automated planning of interactive narrative. A simple experiment could be defined as an abstract operator with a series of actions like the following:

(Abstract action): Verify that preconception bias is resolved
(Abstract action): Solicit hypothesis from user in terms of expected patterns of observations, underlying model
(Abstract action): Obtain and set up equipment
(Concrete action): Set input on experimental device A on value A1
(Concrete action): Observe result with observational device O as observation O1
(Concrete action): Set input on experimental device A on value A2
(Concrete action): Observe result with observational device O as observation O2
. . . Continue cycles of observations according to specified hypothesis
(Abstract action): Solicit from user explanation or interpretation of results
(Abstract action): Solicit from user next step – update model and/or experiment

Notice that many of these actions are "abstract" meaning they are candidates for further decomposition. Such decompositions may even involve actions that require action on the part of other agents, or NPCs.

4.5. Additional Narrative Devices

4.5.1. Equipment Intermediaries

It may be useful to "wrap" intermediary narrative constructs around pieces of observational and experimental equipment. Instead of allowing the user to manipulate a device directly, each device could have an NPC caretaker with whom the user must communicate. In other words, rather than giving the user direct access to a virtual microscope, assume the microscope is "owned" by Microscope Max. This could offer multiple benefits. First, for scaffolding purposes, Microscope Max may restrict the choices available to the user. Second, Max can force the user to self-explain results (e.g., Max asks "is that what you expected?"). Third, Microscope Max can solicit information from the user to better update the system's model of the user's understanding. Fourth, Max can prod an inactive user or a user who is exploring ineffectually with hints. In a given narrative of experimentation, each instrument could have its own NPC, or one caretaker could interface all or some subset of the instruments.

4.5.2. Peanut Gallery / Greek Chorus

Additional NPCs may be added to a given scenario, playing different roles with the user from patient mentor, to anxious apprentice, to unimpressed cynic. We support an unlimited number of guiding NPCs because this is a key advantage of narrative we would like to exploit. For a traditional ITS there are many non-obvious conventions that constrain the variety of guidance a designer can provide, largely because there is just a single agent interacting with the human user. Previously, ITS researchers have worried that violating these conventions and politeness maxims "would produce a tutor that is abrasive, rude, boring, pedantic, or incoherent. But there may be specific conditions in which it is advantageous, pedagogically or stylistically, to violate specific maxims [40]." With an unlimited number of characters, designers have the freedom to portray agents who do not share the learner's goals, may offer bad advice, or be used to provide negative examples.

4.5.3. Preconception Biases

Common intuitions that learners have which may effect (positively or negatively) their ability to effectively perform the experiment can be captured in the language of automated planning. These biases may be expressed in terms of modal beliefs that must be established or removed prior to a particular action can be added to the plan. Templates can be provided that pair preconception operators with learning operators in a principled manner.

4.5.4. Hints

Hint may be expressed as expository actions available to agents who act as explicit "mentors" in the world, or they may be expressed as 3rd party conversations to be overheard by the student, "notes" left behind by more accomplished scientists, or simply visuals or sounds that appear at opportune moments to draw the user's attention in a particular direction. Negative hints could be expressed by NPCs we expect the student to regard as unreliable (e.g., "Robby Rat said to turn up the temperature, but the Robby is usually wrong, so I will turn it down instead"). For a given learning task, the set of hints may be quite large. For each hint, we may want to specify how "strong" the hint is, either on some fixed range, or perhaps in terms of knowledge prerequisites inferred of the student.

4.5.5. Assessments

Following completion of a learning task, are there assessment tools that could be employed through the narrative? Example: an NPC asks "Hey Joe, did that experiment tell you anything about how that nasty looking green goo behaves?" In this manner, narrative provide a less-intrusive vector for updating our student model as well as a convenient vehicle for prompting for self-explanations.

5. Summary

This paper describes an approach to enhancing scientific discovery environments with guidance in the form of a structured interactive narrative. We are confident that interactive narrative provides a powerful platform for integrating discovery learning, adaptive student modeling, and a variety of learning principles evidenced in modern video games. As our work proceeds, we expect to broaden the variety of narrative devices available to describe the process of scientific discovery.

6. Acknowledgements

This is material is based upon work supported under a National Science Foundation Graduate Research Fellowship and National Science CAREER award 0092586.

7. References

[1] National Science Foundation: Science Hard, 2002, June 5. The Onion, America's Finest News Source. Issue 38 – 21. Available from URL: http://www.theonion.com/content/node/38575

[2] Donovan, M.S., and Bransford, J., eds. (2005). How Students Learn: History, Mathematics, and Science. National Research Council. The National Academies press.

[3] Lave, Jean. (1988). Cognition in practice. Boston: Cambridge University Press.

[4] de Jong, T. and van Joolingen, W. 1998. Scientific Discovery Learning with Computer Simulations of Conceptual Domains. *Review of Educational Research* 68(2):179-201.

[5] Bereiter, C., & Scardamalia, M. (1985). Cognitive coping strategies and the problem of "inert knowledge." In S. Chipman, J. Segal, & R. Glaser (Eds.), *Thinking and Learning Skills: Research and Open Question* (vol. 2, pp. 65–80). Hillsdale, NJ: Erlbaum.

[6] Cognition and Technology Group at Vanderbilt. (1990). Anchored instruction and its relationship to situated cognition. *Educational Researcher*, 19 (6), 2–10.

[7] Aleven V. & Koedinger, K. R. (2002). An effective metacognitive strategy: Learning by doing and explaining with a computer-based Cognitive Tutor. *Cognitive Science*, 26, 147-179.

[8] Mayer, R. (2004). Should there be a three-strikes rule against pure discovery learning? The case for guided methods of instruction. American Psychologist, 59(1), 14-19.

[9] Kirschner, P. A., Sweller, J. & Clark, R. E. (2006). Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. Educational Psychologist, 41(2), 75-86.

[10] Njoo, M., & de Jong, T. (1993). Exploratory learning with a computer simulation for control theory: Learning processes and instructional support. Journal of Research in Science Teaching, 30, 821–844.

[11] Glaser, R., Schauble, L., Raghavan, K., & Zeitz, C. (1992). Scientific reasoning across different domains. In E. de Corte, M. Linn, H. Mandl & L. Verschaffel (Eds.), Computer-based learning environments and problem solving (pp. 345373). Berlin, Germany: Springer-Verlag.

[12] van Joolingen, W.R., & de Jong, T. (1993). Exploring a domain through a computer simulation: traversing variable and relation space with the help of a hypothesis scratchpad. In D. Towne, T. de Jong & H. Spada (Eds.), Simulation-based experiential learning (pp. 191-206). Berlin, Germany: Springer-Verlag.

[13] Kuhn, D., Amsel, E., and O'Loughlin, M. (1998). *The development of scientific thinking skills*. San Diego, CA: Academic Press.

[14] Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. Cognitive Science, 12, 1-48.

[15] Lavoie, D.R., & Good, R. (1988). The nature and use of predictions skills in a biological computer simulation. Journal of Research in Science Teaching, 25, 335–360.

[16] Simmons, P.E., & Lunetta, V.N. (1993). Problem-solving behaviors during a genetics computer simulation: Beyond the expert/novice dichotomy. Journal of Research in Science Teaching, 30, 153–173.

[17] Chi, M. T. H., Siler, S., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from human tutoring. Cognitive Science, 25, 471-533.

[18] Shute, V.J., & Glaser, R. (1990). Alarge-scale evaluation of an intelligent discovery world: Smithtown. Interactive Learning Environments, 1, 51–77.

[19] Quintana, C, Reiser, B. J, Davis, A, Krajcik, J., Fretz, E., Duncan, Kyza, E., Edelson, D., Soloway, E., A Scaffolding Design Framework for Software to Support Science Inquiry, Journal of Learning Sciences, 13 (2004), pp. 337-386.

[20] Ainsworth, S. (2003). Evaluation Methods for Learning Environments. Workshop held in conjunction with the 11 th International Conference on Artificial Intelligence in Education (AIED'2003), Sydney, Australia.

[21] Cohen, P. A., Kulik, J. A., & Kulik, C. L. C. (1982). Educational outcomes of tutoring: a meta-analysis of findings. American Educational Research Journal, 19, 237–248.

[22] Bloom, B. S. (1984). The 2 sigma problem: the search for methods of group instruction as effective as one-to-one tutoring. Educational Researcher, 13, 4–16.

[23] Fitz-Gibbon, C. T. (1977). An analysis of the literature of cross-age tutoring. Washington, DC: National Institute of Education, (ERIC Document Reproduction Service No. ED 148 807).

[24] Bunt, A., & Conati, C. (2003). Probabilistic Student Modelling to Improve Exploratory Behaviour. User Modelling and User-Adapted Interaction, Kluwer Academic Publishers, Printed in the Netherlands, 13: pp. 269-309.

[25] Gee, J. P. 2003. *What video games have to teach us about learning and literacy.* New Your; Palgrave Macmillan.

[26] Barwood, Hal. "The Language of Games". Austin Game Writers Conference. 27 October 2005.

[27] Bruner, J. 2004. Narratives of Science. In E. Scanlon, P. Murphy, J. Thomas, and E. Whitelegg (Eds.), *Reconsidering Science Learning,* pp. 90-98. London; Routledge.

[28] Steam News: About Half-Life2. Available at URL: http://www.steampowered.com/v/index.php?area=news&id=648

[29] Laidlaw, Marc. "Gaming The Narrative." Austin Game Writers Conference. 27 October 2005.

[30] Cavazza, M., Charles, F., and Mead, S. 2002. Interacting with Virtual Characters in Interactive Storytelling. In Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems, 318-325. Bologna, Italy.

[31] Magerko, B., Laird, J., Assanie, M., Kerfoot, A., and Stokes, D. 2004. AI Characters and Directors for Interactive Computer Games. In Proceedings of the 2004 Innovative Applications of Artificial Intelligence Conference, 877-883. San Jose, CA.

[32] Mateas, M. and Stern, A. 2005. Structuring Content in the Fac□ade Interactive Drama Architecture. In Proceedings of Artificial Intelligence and Interactive Digital Entertainment, 93-98. Marina del Rey, CA.

[33] Swartout, W., Hill, R., Gratch, J., Johnson, L., Kyriakakis, C., Labore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J., Rickel, J., Thiébaux, M., Tuch, L., Whitney, R., and Douglas, J. 2001. Towards the Holodeck: Integrating Graphics, Sound, Character and Story. In Proceedings of the Fifth International Conference on Autonomous Agents, 409-416. Montreal, Canada.

[34] Young, R. M. & Riedl, M. (2003) Towards an Architecture for Intelligent Control of Narrative in Interactive Virtual Worlds, in the Proceedings of IUI, January, 2003.

[35] Aylett, R.S. 2000. Emergent Narrative, Social Immersion And "Storification" *Proceedings, Narrative Interaction for Learning Environments*, Edinburgh.

[36] Mott, B. and Lester, J. 2006. U-DIRECTOR: A Decision-Theoretic Narrative Planning Architecture for Storytelling Environments. In Proceedings of the Fifth International Conference on Autonomous Agents and Multi-Agent Systems, Hakodate, Japan. Forthcoming.

[37] Riedl, M., Saretto, C., and Young, R. M. *Managing interaction between users and agents in a multiagent storytelling environment*. In Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-03) (2003).

[38] Riedl, M., and Young, R. M. *An intent-driven planner for multi-agent story generation.* In Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-04) (2004).

[39] Christian, D. and Young, M. 2004. Comparing Cognitive and Computational Models of Narrative Structure. In Proceedings of the Nineteenth National Conference on Artificial Intelligence, 385-390. San Jose, CA.

[40] Fikes, R., and Nilsson, N. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intellligence,* 2, pp. 189-208.

[41] Graesser, A. C., Person, N., & Magliano, J. (1995). Collaborative dialog patterns in naturalistic one-on-one tutoring. Applied Cognitive Psychology, 9, 359–387.