ANIMATIONS OF THOUGHT:
INTERACTIVITY IN THE TEACHABLE AGENT PARADIGM

Daniel L. Schwartz, Kristen P. Blair Gautam Biswas, Krittaya Leelawong
Stanford University *Vanderbilt University*
Joan Davis
University of Washington, Seattle

To appear in: *Learning with Animation: Research and Implications for Design.*
R. Lowe & W. Schnotz (Eds). UK: Cambridge University Press.

Animations are a versatile media for displaying changes over time. They can show cellular processes, a billion years of continental drift, the assembly of a desk, and even the invisible shifting of political tides. Most animations depict changes to a situation, such as a desk being assembled. In this chapter, we describe a series of software environments, called Teachable Agents (TAs) that use animations in another way. Rather than displaying a situation, the TAs animate the thoughts an individual might use to reason about that situation. For example, using the same well-structured representations as experts, TAs can visually model how to reason through the causal chains of an ecosystem. This is worthwhile, because the goal of learning is often to emulate an expert's reasoning processes, and animations of thought make that reasoning visible. For novices, learning to reason with an expert's knowledge organization is as important as learning the bare facts themselves.

We build TA systems to capitalize on the adage that an effective way to learn something is to teach it, and this framework has allowed us to introduce some uncommon uses of animation. One novelty is that students help build the animation rather than just watch. Students teach their TA by constructing a visible knowledge organization. For example, students can create a concept map that teaches their TA about a river ecosystem. A second novelty is that the TA solves problems based on what it was taught. For example, given a question about a river eco-system, the TA can visually trace its reasoning over the concept map. This constitutes the animation of thought, and depending on the conclusion that the agent reaches, students can revise their TA's knowledge and their own.

TA environments are highly interactive. Students share their ideas with the agents by teaching them. At the same time, the TAs use artificial intelligence techniques to solve problems independently – though still based on what they have been taught. This creates an environment of shared ideas and shared initiative between the student and the agent. The thesis of this chapter is that the shared ideas and initiative of a TA environment help create an interactive "sweet spot" that optimizes motivation and learning. In Sections 2 and 3 we develop the framework of this thesis, and in Sections 4 and 5 we present relevant evidence. First, however, we take a quick detour to provide a concrete example of a Teachable Agent.

1.0 A Quick Tour of a Teachable Agent

Figure 1 shows an example of a TA named Betty's Brain, or Betty for short (for instances of other agents, see <aaalab.stanford.edu>). Students teach Betty by creating a visual network of nodes and links comprised of qualitative causal relations (i.e., increase, decrease, depends-on, is-a, and has-a semantics). Students use a point-and-click graphics editor to create the nodes (e.g., algae, oxygen) and links (e.g., produce). Students use pull-down menus to specify the qualitative relation implied by the link (e.g., algae <u>increase</u> oxygen). The directed graph and the qualitative semantics provide a well-structured representation that is common among experts discussing causal propagation, plus it enables the students to build the thoughts that are animated when Betty reasons.

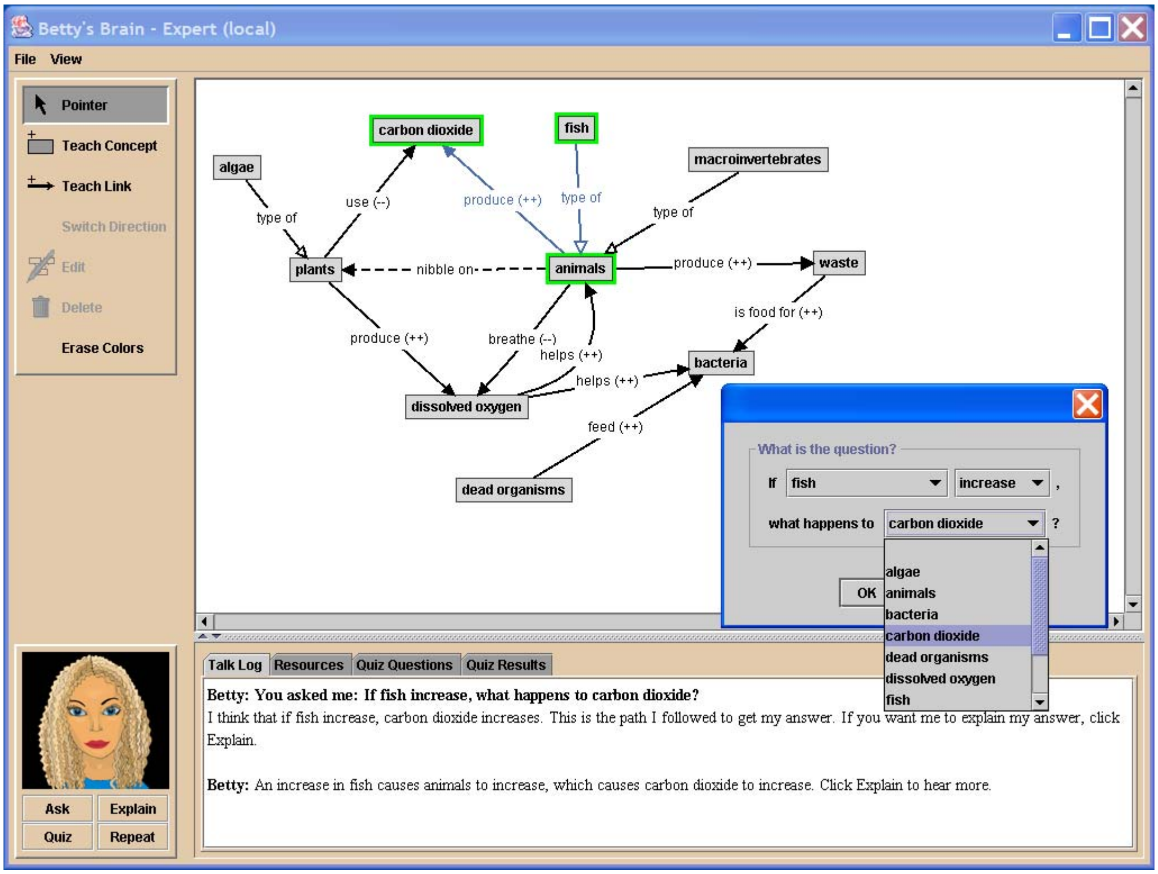

Figure 1. The Teachable Agent Formerly Known as Betty's Brain.

Once taught, Betty can answer questions. Students can ask Betty a question using a floating panel that appears when students click on the "Ask" button. Students can ask questions like "If <fish> <increase> what happens to <carbon dioxide>?" Betty answers the question using the map. This constitutes the animation of thought – she successively highlights nodes and links as she reasons through the network. Changes in color indicate whether she is inferring an increase or decrease for each node while she progresses through the map. The animation is relatively impoverished, and it is probably better described as a dynamic directed graph. However, unlike many animations, where students need to penetrate the surface motion of the animation to learn the underlying principles (Lowe 1999, 2004), Betty makes the critical relations explicit and available for inspection.

To answer questions, Betty uses a simple reasoning engine that adopts generic graph traversal algorithms (e.g., depth and breadth first search) coupled with qualitative inference schemes that can reason through causal and hierarchical chains (Biswas et al., 2005). Betty always reasons "logically," even if the premises she has been taught are incorrect. This helps students learn to emulate Betty's reasoning methods, while also helping students identify gaps in knowledge when Betty reaches a wrong conclusion based on the information they provided.

Figure 1 shows the results of a graphical animation and a text-based response that Betty offers for the question, *"If fish increase, what happens to carbon dioxide?"* Betty inferred that fish are a type of animal. She then reasoned that animals produce carbon dioxide, so an increase in fish will increase the amount of carbon dioxide. Betty can reason forward through much more intricate chains of causes, and she can also reason backward to diagnose what might cause an increase or decrease for a given entity in the map. Students can also ask Betty to explain her answer. Through a multi-step animation, Betty decomposes her chains of reasoning as she provides answers in a sequence of steps. Betty can also unfold her inference through spoken dialog and a text window. Betty's responses help students reflect on the implications of the ideas they taught.

Betty does not learn automatically by using machine learning algorithms. Instead, students must explicitly teach Betty, and this teaching helps students structure their own knowledge. One benefit of the TA paradigm is that it capitalizes on the well-defined teaching schema that includes instruction, assessment, and remediation. This pre-existing schema can help organize otherwise complex student interactions with the computer (cf. Schnotz, Boeckheler, & Grzondziel, 1999), much as people's well-defined schemas for spatial organization inspired the desktop metaphor for computer operating systems. With the TA metaphor, students bring to bear a host of prior ideas about teaching that help them engage the computer in complex learning interactions. Moreover, the TA's use relatively simple visual representations and semantics, and this make it easier for novices to start interacting with their agents, compared for example, to environments that use more general purpose programming constructs (Smith, Cypher, & Spohrer, 1997).

The Betty "kernel" shown in Figure 1 provides basic functionality in a modular, agent architecture (Viswanath, Adebiyi, Biswas, & Leelawong, 2004). This permits us to integrate her into more complex applications and environments. We do not envision Betty as the only means of instruction; students need to learn what to teach Betty from somewhere else. Rather, Betty helps novices abstract and reflect upon important knowledge structures. We provide three quick examples of how Betty can integrate with other environments.

The first example is a guided-discovery video game called *Pumpkin World* (Blair & Schwartz, 2004; Hartman & Blair, 2005). Figure 2 provides a sample screen shot. Betty takes the form of an embodied agent in a virtual world. Students teach Betty to grow giant pumpkins (so villagers have a place to live). There are other agents with whom Betty interacts (e.g., a store owner, a passer-by), and she can directly affect her world (e.g., she adds nitrogen, if she infers the pumpkins need it).

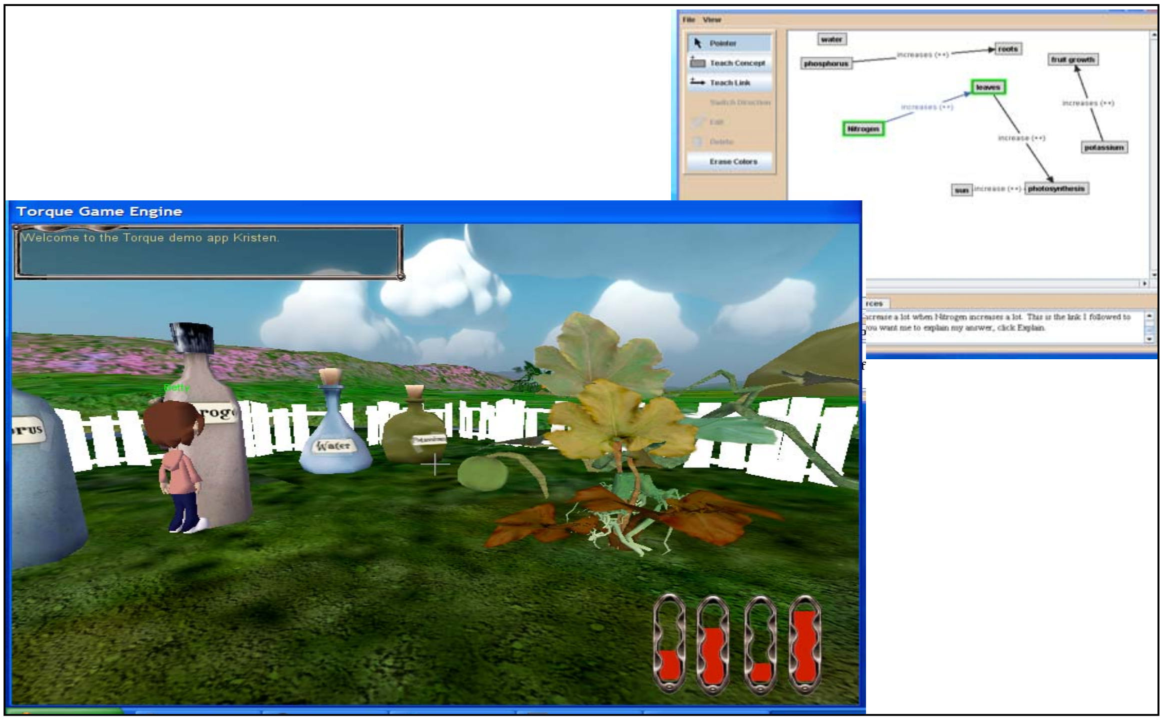

Figure 2. Betty in a guided-discovery videogame application. The window in the corner shows Betty reasoning about a challenge in the game world.

There are many potential benefits of bringing videogame graphics and narratives to learning. They can enhance engagement and a sense of place, plus animated agents offer ways to present non-verbal cues for interaction (Lester, Stone, & Stelling, 1999; Rickel & Johnson, 1998; Tepper, Kopp, & Cassell, 2004). Equally exciting to us, video-worlds provide a felicitous environment for supporting the many ways that people learn. In Pumpkin World, students learn about the role of nitrogen through simulated experiments; they learn about phosphorous by observing a passer-by; they learn about "energy" by listening to the store owner; and so on. Betty serves the role of helping students organize their learning into a well-structured representation of how to reason about pumpkin growth. To nudge students to the different learning resources at the right time, a background planning system makes decisions about what topic to guide the student/agent team to learn about next (e.g., it opens a "game booth" that holds new information when Betty is prepared to learn a new concept). Hence, the name, "guided discovery games."

Our second example is an on-line game show designed to change homework practices. In the Triple-A Game Show, developed with Paula Wellings, a student teaches an agent and customizes its looks. The student and agent then participate in an on-line game show with other students and their agents (Figure 3). Students can log on from home or from school. The game host asks the agents to answer questions and explain their reasoning. The application also includes a chat environment so students can discuss and cheer (or jeer) an agent's performance. Students can also teach their agent a portion of a domain and then "jigsaw" with other agents by merging concept maps to create a Team Betty (Sears & Schwartz, 2004). Our hope is that students will find this socially rich environment both engaging and educative, and it will prepare them for their lessons in school the next day or the next week.

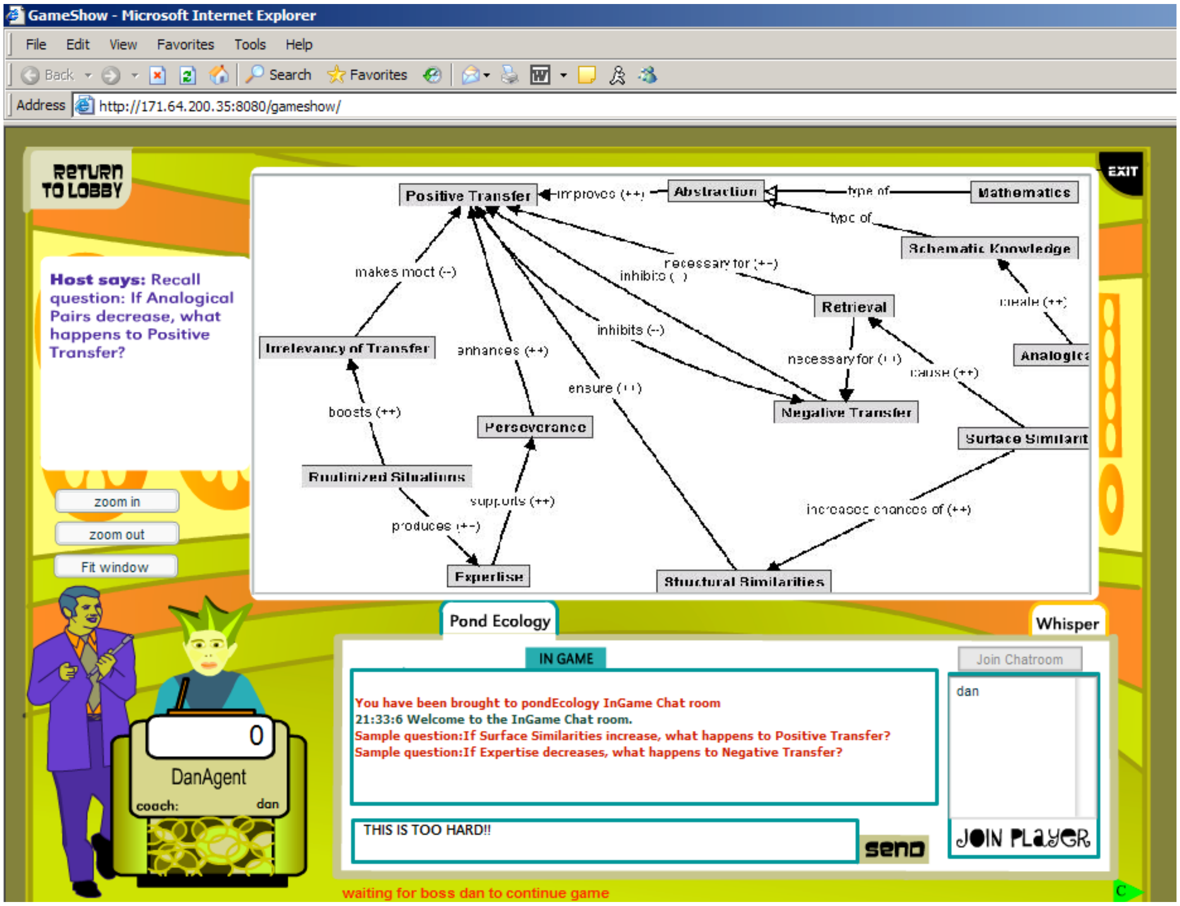

Figure 3. A customized agent performs in an on-line game show with other students and agents.

As a final example of how TAs can be extended, we developed a front-of-the-class assessment environment that can be projected on a large screen. In Figure 4, each panel shows an agent map created by a student. The classroom teacher can ask a question of all the agents simultaneously. A hidden expert map determines the correct answer and compares it to each agent's answer. The results are tabulated and indicated by color coding (red = incorrect; green = correct; yellow = correct for wrong reason). The classroom teacher can zoom in to show why an agent gave the answer it did, and then compare it to another map. If we shed the TA metaphor, one way to think of this system is that students are creating executable models, and then the models get tested. So, rather than the student answering a small subset of questions, the student needs to create a model that can answer any legitimate question in the domain. In a formal study in college classrooms, we found that the front-of-the-class system significantly helped students learn complex relations compared to just seeing the performance of their own agent.

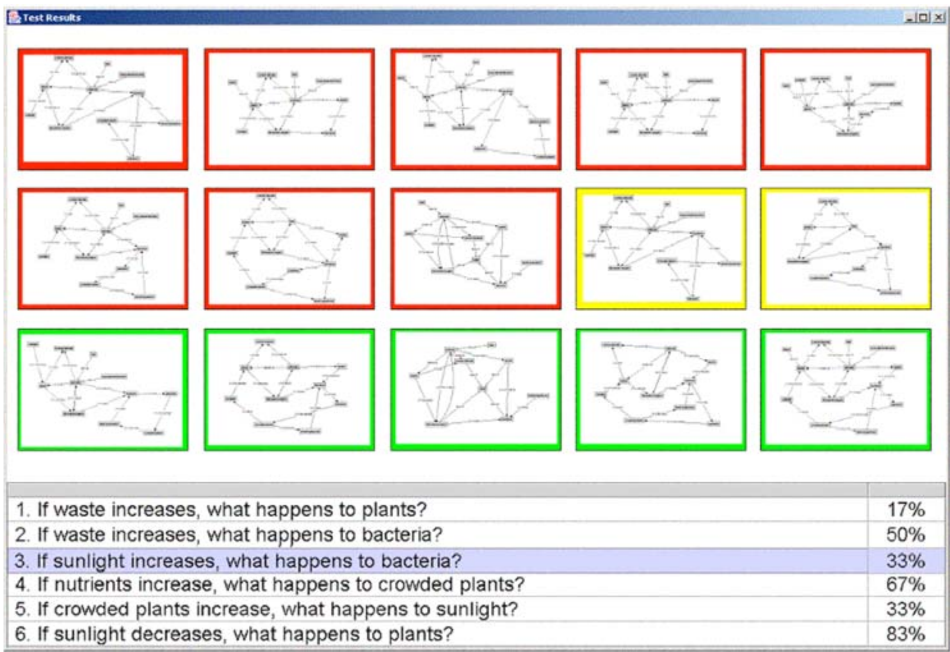

Figure 4. Front of class quiz system for showing agents perform.

2.0 A Framework for Achieving a Learning Sweet Spot in Interactivity

Our thesis is that there is an interactive "sweet spot" for learning that applies to TAs and beyond. TAs operate under a social model of interactive learning rather a physical one. Often times, discussions of interactivity tacitly borrow from models of physical interaction, where people probe a stable environment to help induce its underlying rules or causes. These models lead designers to focus on the contingency of the system; for example, is timely feedback more useful than delayed feedback (e.g., Mantosh & Koedinger, 2003)? Social interactivity presents a different root metaphor that supports a host of new interactive possibilities (e.g., Moreno, Mayer, Spires, & Lester, 2001). For example, Deutsch (1973) describes the minimal criteria for cooperative interactions, "A cooperative process is characterized by open and honest communication of relevant information among participants. Each is interested in informing, and being informed by, the other" (p. 29). One does not think of a physical environment as having interests or communicative intents, but with a social model of interaction it is possible. For example, we have found that children take responsibility for their TA's interests. Children willingly study more to revise their agents so they can pass a test they failed, something students are not always willing to do for themselves (Biswas et al., 2001).

We propose that there is a sweet spot of social interaction that generates both high motivation and high learning. The defining quality of the sweet spot is that participants can produce, share, and see their ideas reflected, transformed, and acted upon by another person (or agent). Elsewhere, we review the motivational basis of the sweet spot, which we termed *productive agency* (Schwartz, 1999). For example, the most satisfying academic conversations occur when people acknowledge and give credit to one another's ideas, and then build on them in a way that new ideas emerge. The least satisfying interactions occur when the participants do not listen to each other, or one has no agency to produce or share ideas at all.

For learning, there are two key dimensions to sweet spot social interactions: initiative of action and inclusion of ideas. Figure 5 presents a qualitative portrayal of the two dimensions. The circle in the center reflects our proposal that optimal novice learning occurs when initiative is shared and ideas are merged among participants.

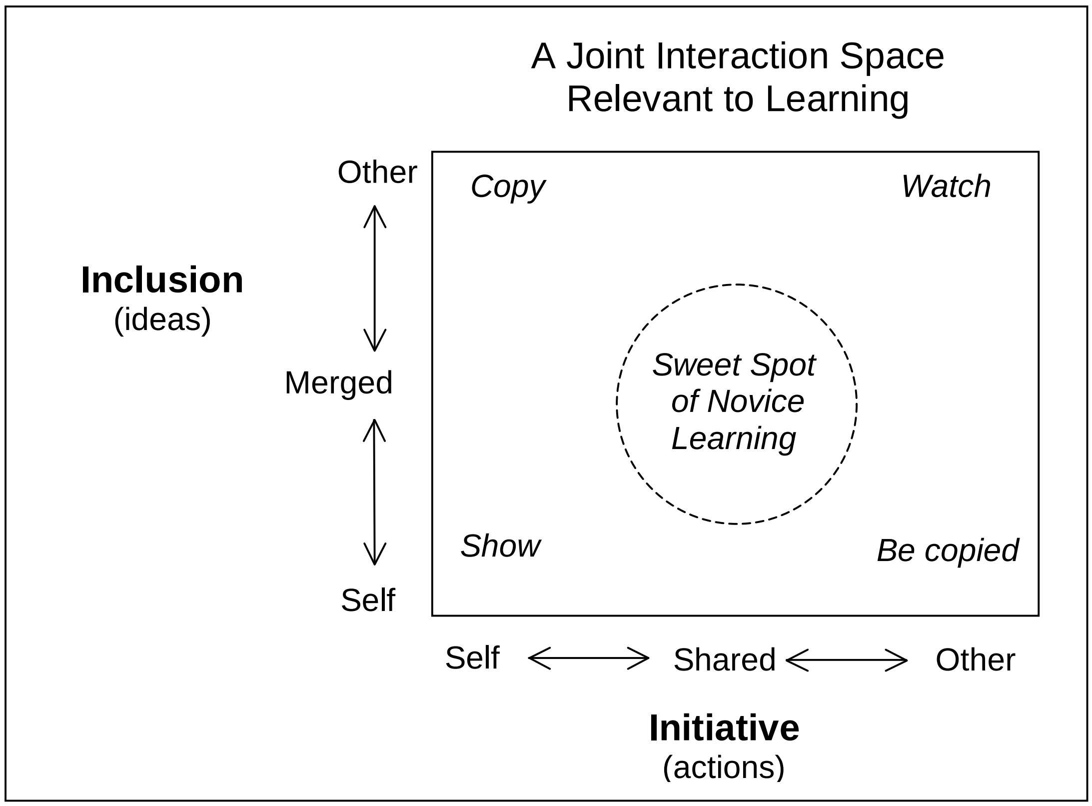

Figure 5. Two dimensions of interactivity relevant to novice learning. Each dimension captures the degree to which the self and another person or technology engages the interaction. The examples in each corner represent non-interactive experiences from the perspective of the learner.

The vertical dimension of interaction – inclusion – captures the degree to which participants incorporate and merge their ideas with one another. People often adopt one another's words and gestures (Bernieri & Rosenthal, 1991; Brennan, 1996). This can draw them closer together (Bailenson & Yee, in press), and the opportunity to see how people modify one's idea can be highly informative. Steiner (1972) reviewed the literature on small group interactions and found that tasks that permit the cumulative building of ideas improve group performance. Moreover, when participants merge their ideas, they can co-create ideas that neither would have developed alone (Schwartz, 1995). However, if participants cannot include their ideas or see another person's ideas, these benefits will not occur.

The second dimension of interactivity – initiative – captures the degree to which each participant can guide the interaction and take independent action. Cassell (2004) demonstrated a story-telling system where young children and an embodied conversational agent took turns creating a narrative, and this improved the children's linguistic skills. Permitting other people to take the initiative offers alternatives to one's own inertia. It also generates "projective" feedback by revealing the variations another person applies to one's ideas. Coaches, for example, may learn a great deal by watching their players take the initiative to adjust a play during a game. In contrast, when an interaction is characterized by an imbalance of initiative there is less learning. Barron (2004) found that small groups that blocked the initiative of one of its members often failed to capitalize on the correct ideas that the individual provided, and therefore, the group members did not perform or learn very well.

Inclusion and initiative are especially important for early learning. Experts who have well-structured knowledge and a wealth of prior experiences can learn by quietly watching or listening. Novices do not have equal knowledge, and they may need more interactive opportunities. An example of an infant-mother dialog can further clarify the sweet spot:

Son: Ball.

Mother: You want me to get the ball?

Mother: That is an apple.

Son: Apple?

The child initiates the exchange and includes an idea (ball) into the joint space. In turn, the mother builds upon the child's ideas, so the child can both recognize his original idea and what is new in what the mother says (apple). The mother's initiative is relevant to the child's own, and this helps the child see the implications of his initiative (e.g., If you want an object, then you have to give it the right name.) The recursive structure of the exchange, shown in Figure 6, makes new information from the mother more comprehensible to the child and leads to better learning. Tomasello and Farrar (1986), for example, demonstrated that infants learn object names more effectively if the mother labels an object the child is handling compared to a situation where the mother labels an object she is handling.

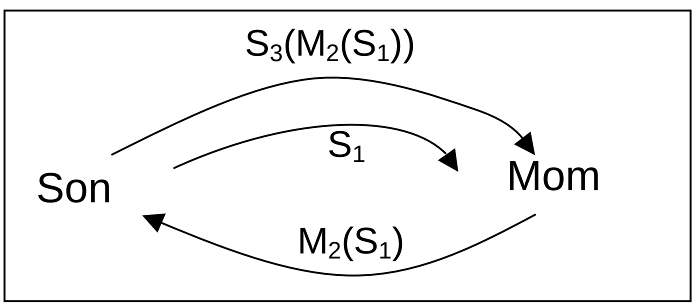

Figure 6. A schematic notation of the mother-child interaction. Arc S_1: The son initiates the exchange and includes an idea into the joint space. Arc $M_2(S_1)$: The mother incorporates the child's intent and takes the initiative to turn the conversation into an object naming lesson. She introduces her name for the object. Arc $S_3(M_2(S_1))$: The child picks up the mother's meaning and tentatively renames the object.

Teaching involves interactions that can create a sweet spot (for reviews, see Biswas et al., 2001, 2005; Renkl, 1995). Learning through peer-tutoring (Cohen, Kulik, & Kulik, 1982) and reciprocal teaching (Palincsar & Brown, 1984) can be highly effective. Many graduate students have stated that they never really understood statistics until they had to teach it. Teachers naturally have the initiative and include their ideas into the interaction. At the same time, good teachers reflect on how students build on those ideas and how students take the conversation in new and unexpected directions. This can help teachers to see new implications and connections, to seek out new sources of examples and evidence, and to abstract out more fundamental structures in their own knowledge. Teachers may find that they need to rethink how they describe a domain, and this can result in a search for more effective explanatory structures.

Of course, not all teachers and students find the sweet spot. Teachers may be overly didactic, and as a consequence, they will not learn much in the act of teaching. Students may also fail to contribute to the interaction. Moreover, if inexperienced children are asked to teach, they may not have sufficient skills to keep teaching interactions in the sweet spot. This is one reason why TAs can be valuable. TAs help to avoid a potential problem of peer-teaching, where students are put at risk if the child-teacher, or the child-pupil, are not very good.

3.0 Putting the Sweet Spot into Technology

When creating interactive technologies, different considerations become important depending on one's guiding model of interaction. Technologies that simulate conversational interactions, for example, need to tune the timing and complexity of the responses they generate. In our work, we do not try to generate sophisticated natural language interactions or realistic agents embodied with human traits. Our goal is to design and implement interactive environments that are sufficient to elicit social schemas that can engage the sweet spot for learning and leverage animations of thought.

Interactive technologies and their animations can optimize the two dimensions of interactivity to varying degrees. Simply watching an animation would be low on student initiative and ideas. Though students must take some initiative to watch and interpret the animation, they cannot substantively alter the course of the interaction, and their primary task is to extract information from the animation rather than contribute to it. For novices, it is not clear that passively viewing an animation yields superior learning compared to studying a well-crafted still image (Tversky, Morrison, & Betrancourt, 2002). Enabling novices to slow down or replay segments of an animation improves the balance of initiative and should help (Lowe, 2004), but the animation still does not include the students' ideas.

Intelligent computer tutors are another important interactive technology for learning (e.g., Koedinger & Anderson, 1998). The student and the system share the initiative, because the student has some latitude in how to solve problems, and the system has the latitude to redirect the student and introduce problems. However, the computer tutor does not merge ideas with the student. The explicit goal of the program is to entrain the student into its way of thinking. Computer chess programs and other learning games are similar in that they share initiative with the player. However, whereas a tutor program explicitly enforces its ideas, a chess program explicitly hides its ideas. A novice, who cannot infer the chess program's underlying strategy, may learn less than if there were a shared representation of the program's strategy.

In designing TA's, our goal is to foster the sweet spot. We believe this goal is implicit in other learning systems that use agents that are neither completely ignorant nor all-knowing (e.g., Learning Companions, Chan, 1995; Peoplepower, Dillenbourg & Self, 1992). TA's, however, are explicitly designed to support the merging of ideas. Students provide the facts of the matter, and the TA provides conventional knowledge structures and reasoning mechanisms in the form of dynamic visual representations. For example, students introduce concepts to Betty, but they use Betty's directed graph structure to organize those concepts. Betty then animates how she reasons with their shared representation. Our assumption is that this visible merging of student ideas and agent reasoning helps students adopt the TA's knowledge structures to organize and reason with their own concepts. The two studies in Section 4 test this assumption.

The TA's also have provisions for shared initiative. Each TA has the ability to take independent actions based on how it has been taught. For example, Betty can answer questions. This permits students to reflect on Betty's reasoning, and we suppose this helps them learn more deeply. It is also possible to enhance a TA's initiative beyond answering questions, and the two studies in Section 5 examine the value of enhanced shared initiative.

4.0 Empirical Studies on the Dimension of Inclusion

Two studies explored whether merging ideas and representations with Betty leads students to adopt her knowledge organization. Both studies used the basic Betty kernel to teach about biological systems. Biological systems are well-suited to Betty's qualitative causal reasoning.

4.1 Study 1: Adopting the structure of the agent's thoughts

The first study examined whether learners incorporate Betty's knowledge structure into their own. The study complements research on the positive benefits of concept mapping (e.g., Kinchin & Hay, 2000; Novak, 1998). However, Betty differs from most concept-mapping activities, because she enforces semantic relations that students might otherwise violate in paper and pencil activities. For example, students often use concept map links to indicate a vague notion of "related to," whereas Betty requires students to use semantics that enable her to reason about causality (e.g., increase, decrease). Plus, Betty shows the implications of those relationships by answering questions, something concept maps, by themselves, cannot do.

Sixteen undergraduates read a four-page passage on metabolism. Half of the students were assigned to the Summary condition. They wrote a summary of cell metabolism. They were told to write about things like the relation between ATP resynthesis and lactic acid. Students in the Betty condition taught Betty about cell metabolism. They were shown how to teach and query Betty using the ATP – lactic acid example. We videotaped the sessions and asked the participants to think aloud. All Betty students worked to the cutoff point of 40 minutes. The Summary students averaged 32 minutes of work.

During the session, the Betty students were much more attentive to issues of causality than the Summary students. Every Betty student, but only one Summary student, recognized that they had been thinking in terms of correlations rather than causation. For example, one Betty student realized that he did not know whether mitochondria increase ATP resynthesis or vice versa. Three-fourths of the Betty students considered the size of a causal effect, whereas none of the Summary students considered amounts of change. For example, one Betty student taught Betty that (a) oxygen increases ATP resynthesis, and (b) lactic acid inhibits resynthesis. This student wondered whether oxygen and lactic acid cancelled each other out.

As a simple posttest, we removed all the materials and gave students a paper with five metabolism terms (e.g., mitochondria, lactic acid, etc.). For each term, students had to "list relations to other entities and processes in cellular metabolism." Students in the Summary condition tended to assert single relations; for example, "mitochondria increase ATP resynthesis." Students in the Betty condition tended to assert chains of two or more relations; for example, "mitochondria with glycogen or free fatty acid increase ATP resynthesis." On average, the Betty students produced 3.75 chains of two or more relations, compared to 1.0 for the Summary students ($p < .05$). It is possible that the Summary students also learned about complex causal pathways, and they simply did not think it was important to include them in their lists. Even so, we can reiterate that the Betty students incorporated Betty's way of thinking and took that incorporation as an important task demand.

In summary, Betty influenced the students' own knowledge. The Betty students became aware that they had not sufficiently differentiated between causation and correlation when reading the passage. In contrast, writing a summary drew attention to issues of topic sentences and paragraphing. Teaching Betty also influenced how students structured their own knowledge. When asked to list relations, the Betty students tended to list complex causal pathways as compared to the Summary students. These results make sense because developing chains of causal relations is exactly what Betty illustrates and requires. The results demonstrate that merging ideas with Betty can shape students' domain knowledge.

4.2 Study 2: The benefit of animated thought

The first study demonstrated that students adopt Betty's representations, but it did not isolate the value of animating Betty's thoughts. In the study, we compared teaching Betty to writing a summary. Maybe simply asking students to draw a concept map would work as well as teaching Betty. Therefore, we conducted a second study to see if Betty's animations make a difference for whether students' incorporate her reasoning structures.

Twenty-five 5th-grade students taught Betty about river ecosystems across three one-hour sessions. Students had online resources to help them learn the relevant content needed to teach Betty (see Biswas et al., 2004). In the Animation condition (n = 13), students could ask Betty questions and activate her animations. In the No Animation condition (n = 12) students simply made Betty's map without ever asking questions and seeing the map animate. Thus, both conditions used Betty's formalism, but only the students in the animate condition saw how Betty incorporated their ideas into her reasoning.

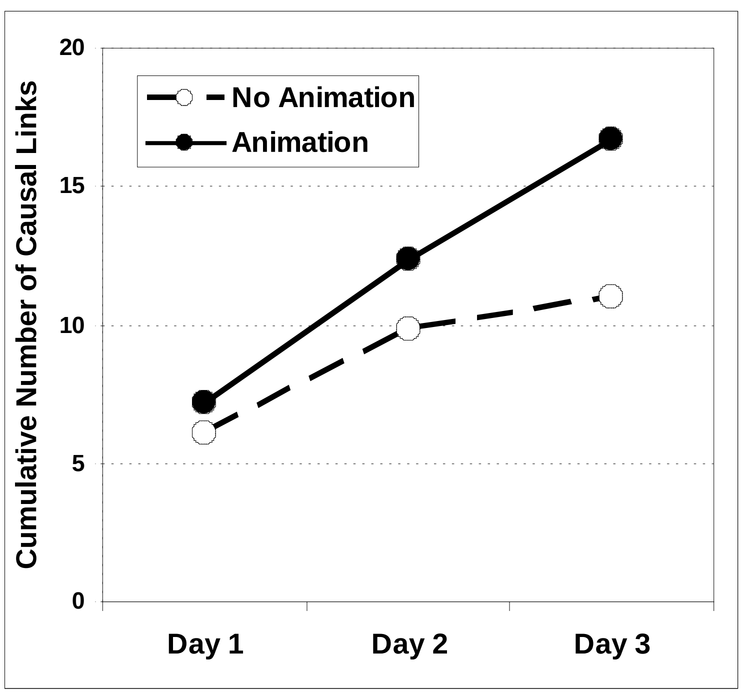

Figure 7. Students who see Betty animate her reasoning add more causal links.

Figure 7 shows that students in the Animation condition generated maps that accumulated more causal links; $F(1, 23) = 4.24$, $MSe = 15.7$, $p < .05$. This makes sense because Betty animates causal chains in her reasoning. Evidently, animations of thought make a difference. This is a useful finding, because Tversky, Morrison, and Betrancourt (2002) found that static diagrams can be as effective as animations. However, their work did not examine the interactive potential of animations, which in the current study was more useful than static drawings. Merging ideas into an animation can help students learn to think with the structures that drive that animation.

5.0 Empirical Studies on the Dimension of Initiative

The previous studies examined the joint inclusion of ideas. These studies used the basic Betty kernel, which is relatively low on initiative. It only re-acts and responds to the student when asked to do so. For novices, this relatively simple initiative can appear to yield choice-filled behavior in the TA, because the performance setting provides complexity (Simon, 1996). Even so, the TA is still reactive. In the next two examples, we capitalize more fully on the TA metaphor for enhancing shared initiative. The studies also raise the standard for demonstrating that students learn. In the previous studies, we largely inferred what students learned based on the maps they created or the reasoning they applied while creating the maps. In the following studies, we measure student learning once they have left the original environment of teaching.

5.1 Study 3: Shared initiative to promote metacognition

In exit interviews from our previous Betty studies, the students often emphasized that they would have liked Betty to take more initiative and exhibit characteristics of a good student during the teaching phase. One student, for instance, wanted Betty to "react to what she was being taught and ask questions on her own." Our assumption is that people learn more by teaching when their pupils can introduce their own issues. Tutors, for example, gain a deeper understanding when they answer tutee questions, explain materials, and uncover misconceptions (Chi et al., 2001; Graesser, Person, & Magliano, 1995; Uretsi, 2000).

For children, an agent that initiates a discussion about its learning can be especially valuable, because children have fewer metacognitive skills for managing learning interactions. Hegarty (2004) argues, "Since it is clear that not all students have the necessary metacognitive skills to learn effectively from interactive media, teaching students to use interactive media effectively may lead to greater improvements in learning outcomes than changing the medium of instruction" (p. 348). If an agent can take the initiative to model how it monitors its own thinking, this may help students learn to perform this type of metacognition (Lin & Lehman, 1999; Palinscar & Brown, 1984). Shimoda, White, and Frederiksen (2002), for example, developed metacognitive coaches that students can modify to give hints for managing their scientific investigations, and Baylor (2002) found that agents can improve pre-service teachers' meta-cognitive awareness.

To address these issues, we gave Betty's initiative over how she is being taught, and we put her into a larger performance environment (Biswas et al., 2005). Under specific conditions, Betty spontaneously offers a meta-cognitive strategy or concern. For example, as students build Betty's map, she occasionally starts animating her map to draw inferences. She then remarks (right or wrong) that the answer she is deriving does not seem to make sense. These spontaneous prompts help students reflect on what they are teaching, and hopefully like a good teacher, check on their tutee's learning progress (and their own). Table 1 provides some examples of the actions Betty initiates and under what conditions.

Table 1. Examples of Betty Taking Initiative.

Student Action	State of System	Samples of Betty's Dialog
Tell Betty to take a quiz	Student has not asked Betty a question since the last quiz.	"I still do not feel prepared to take a quiz. I don't understand enough about the causal relationships in the river. Please ask me some causal questions to see if I understand. Mr. Davis can help you learn more about being a good teacher."
	Betty's answers have changed since the last quiz	"What you have taught me has changed my thinking. I had some questions right on the quiz, but now I think I would answer them wrong."
Add a link to Betty's concept map	First causal link in the session	"Hey! Let me see if I understand this." (Betty reasons with link, and explains her reasoning)
	First causal path with two or more links	"OK. I think I know how this works." (Betty reasons with path, and explains her reasoning)

The meta-cognitive Betty environment comes with a number of other assets, including online resources for learning content, quiz sets, and a mentor agent named, Mr. Davis. Mr. Davis helps to complete the teaching narrative, because he administers and grades Betty's quizzes. Mr. Davis also provides meta-cognitive tips when students ask for help. Mr. Davis does not give factual answers, but rather, he suggests strategies. As instances, he can suggest which of the online resources is helpful for a particular concept; how to be a good teacher (e.g., "test Betty and examine her answers closely"); and, how to be a good learner (e.g., "set goals").

To evaluate the benefits of shared-initiative, 54 5th-grade students worked for five 45-minute sessions on river ecosystem concepts. The study included three conditions. In the Shared-Initiative Teaching condition, students worked with the enhanced Betty system. Students taught Betty so she could pass a test to become a member of a school science club. In the Basic Teaching condition, students also prepared Betty for the club test, but Betty was similar to the prior studies where she simply took a quiz and answered student questions when asked. Mr. Davis did not provide tips on teaching and learning. Instead, he provided feedback to Betty on each quiz question. Finally, in the Being-Taught condition, there was no cover story of teaching an agent, and Betty was not present in this environment. The students simply had to construct a concept map. Mr. Davis told the students to construct concept maps to demonstrate their learning. They were told to examine the quiz questions as a guide for what to learn. Students could ask Mr. Davis if their concept map was correct for a given quiz question. Mr. Davis would tell the student the correct answer and provide directive feedback for how to correct the map. Thus, the Being-Taught condition replicated standard computer-based instruction, where the computer has the initiative to teach and test the student.

After students completed the five sessions, they drew their maps from memory. The maps from the three conditions looked about the same. This result was not a surprise, because the students had worked for a long time developing their maps in each condition. This was not where we expected to find the difference between the conditions.

Our hypothesis was that the Shared-Initiative condition would show its benefits later. We thought the meta-cognitive emphasis would prepare students to learn about a new, related topic. Because the Shared-Initiative students interacted with Betty's metacognitive strategies, we thought they might incorporate those strategies when learning new content.

A month after the instructional intervention, students had two new sessions on the land-based nitrogen cycle. Students had not been taught about the nitrogen cycle, so they would have to learn from resources. Students from all three conditions completed this "learning at transfer" task (Bransford & Schwartz, 1999) in the same environment, so any differences would be due to what had happened a month before. They all used a modified Being-Taught environment that did not include directive feedback. Mr. Davis simply told the students whether their map was right or wrong when they asked a question or took a quiz.

Table 2 shows that students from the Shared-Initiative condition learned significantly more about the nitrogen cycle as reflected in their concept maps. Their performances were still relatively low, but learning about the nitrogen cycle on their own in two sessions is a difficult task for 5th-graders.

Table 2: Quality of Student Maps when Learning about Nitrogen Cycle at Transfer.

Student Maps Included:	Shared-Initiative M (SD)	Teaching M (SD)	Being Taught M (SD)
Expert Concepts	6.1[a] (0.6)	5.2 (0.5)	4.1 (0.6)
Expert Causal Links	1.1[ab] (0.3)	0.1 (0.3)	0.2 (0.3)

Note: [a] Significantly greater than Being Taught; [b] Significantly greater than Teaching.

The log files help explain the advantage of the Shared Initiative students on the learning posttest. Across the five river ecosystem sessions, the Shared Initiative students increasingly asked Betty to answer questions about chains of causes. By the fifth session, theses students had asked three times as many causal questions as the Being Taught condition. During the learning posttest on the nitrogen cycle, the Shared Initiative students asked twice as many causal questions as the Being Taught condition, even though Betty was no longer mixed initiative. In other words, sharing the initiative does not reduce the students' overall initiative, but rather it increases it. Meta-cognitive Betty led the students to increase their initiative to ask her questions, and this carried over, so that the students exhibited more learning initiative a month later when they had to do it on their own.

5.2 Study 4: The importance of independent performance

In the final study, we carefully isolated the value of agent initiative. In this case, the agent does not initiate interactions with the student, but instead, it interacts independently in another context after it has been taught. The description of the study takes some extra prose, but we think it is worthwhile. The study demonstrates that when learners solve problems themselves and construct representations, they do <u>not</u> learn as well as when they construct the exact same representations (to teach) and then see the agent perform based on those representations. In other words, "watching" within the sweet spot can be more effective than doing it oneself outside the sweet spot.

To motivate this finding, consider dissertation defenses in Sweden. The doctoral candidate does not answer questions directly. Rather, there is an advocate who answers questions based on what the candidate wrote in the thesis. The advocate's ability to defend the thesis depends on the candidate, but the advocate has the initiative to decide what to say. We assume this situation leads to much more careful thinking and writing by the candidate. When people teach somebody else who has to perform, they cannot count on their "situational smarts" to generate answers on the fly or sidestep challenges as they arise. They need to formalize their knowledge in a clear and unambiguous way. Moreover, by seeing how their students perform, for example, as the students answer questions from an outsider, teachers receive projective feedback, and this provides an opportunity for additional reflection (plus a sense of responsibility).

5.2.1 Moby: A Hypothetico-Deductive TA

To explain how we implemented the experimental contrast, we need to take a second detour to describe the TA named Moby. Moby helps students learn science through a hypothetico-deductive process. Scientific reasoning is notoriously problematic (e.g., Kuhn, 1995), so we decided to see if the TA paradigm could help students learn about deduction and induction.

With Moby, students construct visual representations of empirical hypotheses, and Moby makes predictions based on these representations. This is consistent with our emphasis on making thinking visible. Thus, Moby should be characterized as a hypothesis visualization tool rather than a data visualization tool (e.g., Gordin & Pea, 1995).

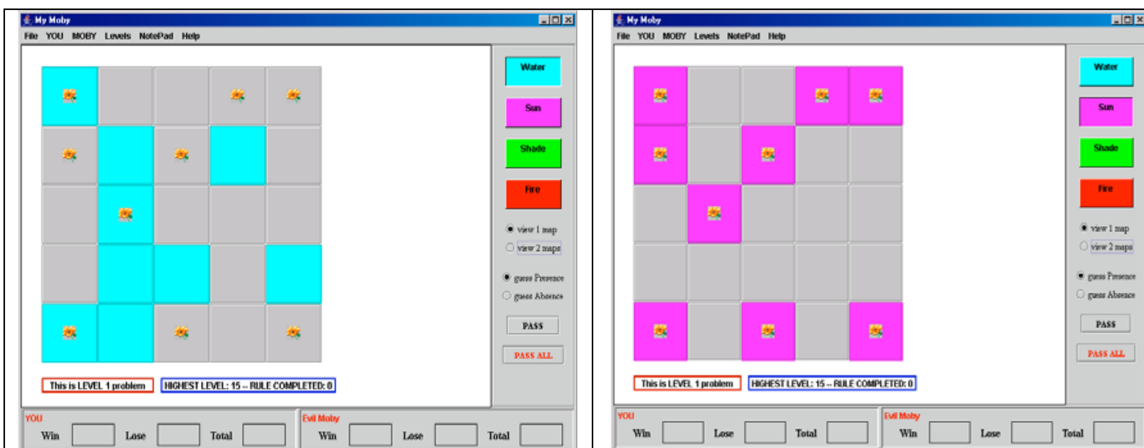

Overlay Water | Overlay Sun

Figure 8. Two examples of a student revealing a factor to induce the conditions that cause an outcome.

Moby resides in a game environment. Students need to complete rounds so they can progress to the next level of difficulty. Each round has four phases: Induce →

Predict → Teach → Observe. Each round begins with Induce. Students receive a grid with a target outcome appearing in various cells. For example, the appearance of flowers is the target outcome in figure 8. There are four factors that might be responsible for the outcome. Students click through the various factors to see where they appear in the grid. In the left panel of Figure 8, a student has revealed a factor (water) that is not responsible for the outcome (flowers). In the right panel, the student has revealed the correct factor (sun). The underlying rule for the grid in Figure 8 is, "Sun is necessary and sufficient for a flower to appear." A rule, however, can be much more complicated; for example, "Fire or shade is sufficient but not necessary for the absence of a flower." To help students find multi-factor rules, they can reveal the locations of two factors at a time.

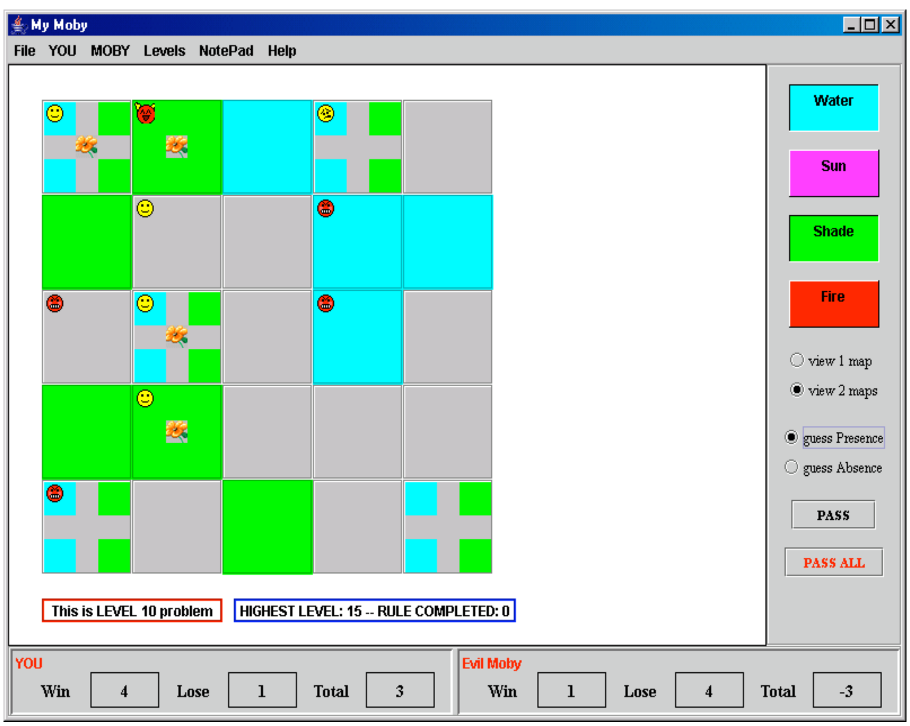

Figure 9. Students play against Joe to see who can better predict outcomes. Students overlay factors to guide their predictions. In this example, the student hypothesizes that Water and Shade are involved and reveals the location of both.

After students believe they have induced the rule for a given round, they move to the Predict phase. In the Predict phase, they play against another agent named Joe. The game generates a new grid (based on the same rule). The grid, however, does not show the outcomes (e.g., flowers). Students need to predict which cells have the target outcome based on their hypothesis. To do this, they reveal where the relevant factors appear on the grid. Wherever they see the factor or factor combinations that they believe cause the outcome, they can click on that cell to see if the outcome is present. If they are correct, a smiling face appears. If they are wrong, a frowning face appears. A scoreboard keeps a tally. After the student takes a turn, Joe takes a turn, then the student, and so on, until all the outcomes have been found. Figure 9 shows a student and Joe in mid-game. Joe has the wrong rule, and most of his predictions are incorrect. The student has revealed two factors to guide her predictions, but she has the wrong rule as well. Based on the student's performance playing the prediction game against Joe, she can return to the Induce phase to refine her rule or move on to the Teach phase.

In the Teach phase, the students' visual understanding of a hypothesis gets merged with Moby's formal representations. Figure 10 shows there are two representations that students use to teach Moby. In the top of the figure, the students merge their visual rule into Moby's propositional representation. Using pull down menus, students choose factor(s), their combination, and the qualifier that governs their relation to the outcome (Necessary, Sufficient, Necessary and Sufficient). In the bottom of the figure, the students teach using a matrix representation. After choosing the factor(s), the students fill the cells with (A)lways, (S)ometimes, (N)ever, indicating when flowers appear in these cells. If students are inconsistent between the two representations, Moby says he is confused and asks them to try again.

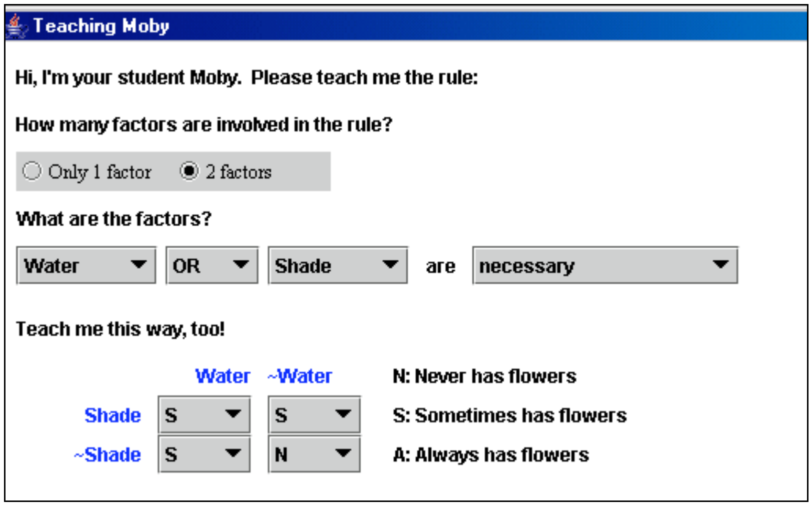

Figure 10. Students teach Moby using propositional and matrix formats.

After students teach Moby, they move to the Observe phase. They watch Moby play the prediction game against Joe. This constitutes the major initiative and independent performance of the TA. Moby's animation consists of choosing factors and making predictions based on the rule he was taught. If Moby consistently loses to Joe, students need to re-teach, and if necessary, return to the Induce phase to develop a new hypothesis. If Moby wins twice in a row, students get to move to the next round and a more complex rule.

5.2.2 Isolating the significance of agent initiative

Moby permitted a unique test that parceled out the value of the TA's capacity for solving problems on its own from other aspects of the TA's (e.g., visual representation). Ninety-four high school students were assigned to four conditions. In the Control condition, students never used the software and simply took a posttest. In the other three conditions, students played the game for about 90 minutes progressing through the rounds of increasing difficulty. In the Teach condition, students completed the full cycle described above. They merged their ideas into the formal representations, and they saw Moby take the initiative to play against Joe. In the Represent condition, students completed the Induce and Predict phases. When they beat Joe twice in a row, they filled in the representations from the Teach phase. However, they did <u>not</u> complete the representations in the context of teaching, and there was no independent initiative of the agent for them to observe. They were simply expressing the rule they had learned, and then they moved to the next level. Thus, these students were merging their knowledge into a formal representation, but they were not sharing initiative. Finally, in the Explain condition, students also completed the Induce and Predict phases. After beating Joe twice in a row, a text window asked them to explain the rule they had used, and they progressed to the next level. They did not see the formal representations and simply had to find a way to express the rule in words. Thus, these students neither merged their knowledge into a formal representation nor did they see any agent initiative.

If the initiative of the agent to perform independently is a valuable aspect of animations of thought, then we should expect the students in the Teach condition to do the best on the posttest, even though students in the Represent and Explain conditions had to induce, use, and formulate rules too. The students in all three software conditions reached the same game level in the same amount of time, so we can be sure there is not a time on task or relative exposure confound.

A few days after using the system, students took an 18 question posttest that included three classes of questions. Induce questions asked students to infer a rule given a combination of factors and outcomes. Imply questions provided a rule, and students had to deduce the implications. Translate questions asked students to convert between tabular and verbal expressions of a rule.

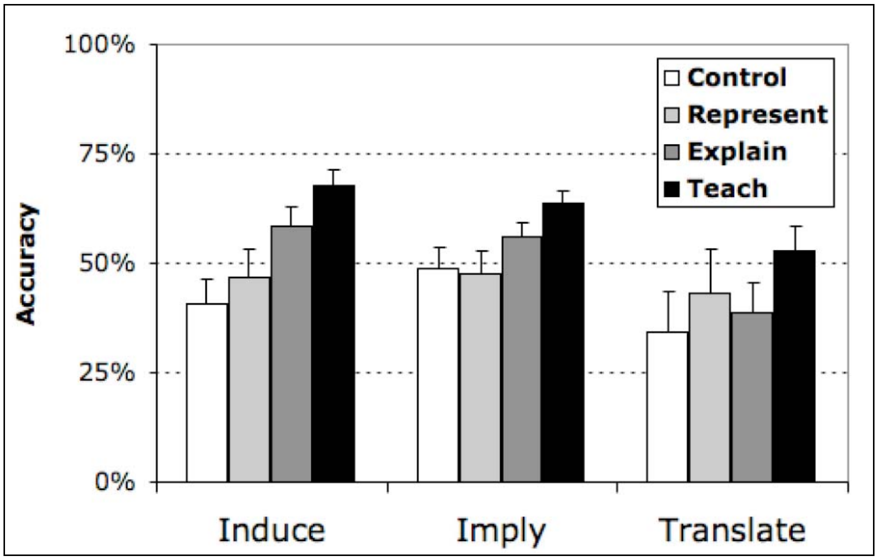

Figure 11. Students in the Teach condition were significantly more accurate on all question types.

Figure 11 shows the posttest performance. The Teach condition produced significantly better performance across all three types of measures compared to each of the other conditions, $\underline{F}(3, 90) = 4.9$, $\underline{MSe}=.04$, $\underline{p} < .01$. None of the other conditions were significantly different from one another, despite the apparent descriptive differences.

The fact that the Teach students did better than the students who played and passed similar levels, but without teaching, is an important result. It shows that seeing a TA perform with feedback is more valuable than just working on problems by oneself and receiving feedback. Moreover, it showed that agent initiative is important for cashing in the value of merged representations. Students in the Represent condition merged their ideas with the same representations as the Teach students did. Even so, the Represent students did not reap the benefits of this merging. The Express students, who never saw these representations, did about the same as the Represent students. This fits our story that the sweet spot of interactive learning occurs when there is both the merged inclusion of ideas and a shared initiative in action.

More broadly, these results suggest one possible learning benefit from programming simulations in general, and not just TAs. Simulations require a full specification of one's knowledge. The "run" of the simulation then provides independent performance and feedback that is not prone to subtle, situation-specific reasoning that can yield successful results but for the wrong reasons. To our knowledge, there has not been a direct test that compares the benefit of a simulation versus an otherwise equivalent activity of solving problems and formalizing one's knowledge, so the current results may have some value beyond the specific demonstration of the value of TA's.

6.0 Conclusions

Learning brute empirical facts is important, but for novices, learning to think with the expert's organization of those facts is equally important. Making thinking visible through animations can help. We assume that learning from an "animations of thought" can be enhanced by making the animations interactive. So, rather than having students only watch animations that demonstrate canonical forms of reasoning, we implemented the Teachable Agents where students help to create and query those animations. Students teach a computer agent, and then see the agent animate its thinking based on how it has been taught. Our TAs are particularly thin in human-like appearance and behavior, but it never ceases to surprise us how readily children (and adults) are willing to adopt and are motivated by the fiction of teaching another person (cf. Reeves & Nass, 1996). However, even if they do not buy into the fiction, they can still draw upon the well-known teaching schema to guide their learning interactions.

The TA's rely on a schema of social interaction, and, therefore, we developed a framework for guiding decisions about the design of TA interactions. We argued that there is a sweet spot of learning interactions that is characterized by (a) how much participants include and reflect one another's ideas, and (b) how much there is shared initiative for taking actions. In four studies, we examined whether interactive animations that targeted a sweet spot led to superior learning.

By design, TAs support a merging of ideas in an interactive context. Students provide domain specific content, while the agent provides canonical knowledge representations and reasoning. Each agent is designed to model a specific form of reasoning, and each agent implements a specific reasoning algorithm and associated representation. TAs are not as complete or powerful as general-purpose programming languages. While this limits a TA's expressiveness, it also helps novices quickly engage a model of reasoning that is suited to initial domain learning. The first study showed that students adopted the representation of Betty, a qualitative reasoning agent. The second study showed that, given a chance to see Betty reason, students incorporated her causal reasoning compared to students who simply built a static causal map. Animations of thought, at least in an interactive context, help students learn.

TA's afford different levels of shared initiative depending on the particular implementation. The basic architecture of a TA always includes the capacity for (a) being taught so the student has initiative, and (b) acting upon what it has been taught so it also has initiative. This bit of shared initiative is helpful. (Students who saw Betty answer questions used more causal reasoning than those who did not). Nevertheless, it is possible to enhance the shared-initiative to improve learning even more. So, in the third study, we compared three levels of shared initiative: Being Taught (lowest), Teaching (middle), Shared-Initiative with a meta-cognitive Betty (highest). The systems looked equally effective in the short run, but when we tested the students after a month, the Shared-Initiative students showed the greatest readiness to learn. The posttest in this study is of particular note. A goal of most conceptual instruction is to prepare students to learn in the future (Bransford & Schwartz, 1999). Except for the narrowest of training, no amount of schooling can provide students all they need to know, and they will need to learn. Therefore, it is important to use assessments of students' preparation for future learning, lest we miss the true value of an instructional approach. For example, had we not examined the students' abilities to learn new content, the value of sharing initiative with a metacognitive agent would have been missed.

In the final study using a hypothetico-deductive agent, we found strong evidence for the added-value of an agent that can initiate actions. Students who saw their agent perform did better on a posttest of hypothetico-deductive reasoning compared to students who otherwise completed the exact same activities of inducing, testing, and formalizing rules. Notably, the opportunity to merge ideas with the formal structures provided by the agent did not benefit students if they did not also get to see the implications of those rules played out in the behavior of an agent. The sweet spot of interaction requires both the merged inclusion of ideas <u>and</u> the sharing of initiative.

In conclusion, much of the research on animation has focused on animations that portray continuous changes to a referent domain. This application of animation has naturally led to comparisons between media that indicate changes over time; for example, videos versus animations, slow versus fast animations, static drawings versus animations, and texts versus (or with) animations. In this chapter, we took a turn that led us to a different set of questions. We created agents that students teach and that animate their thinking. This generated new issues for guiding design and research. In particular, it led us to explore the hypothesis that students learn better from animations when the learner and animation share ideas and initiative. We believe this hypothesis extends to interactions without agents and to human-human interaction as well. The initial empirical results are promising, but of course, there are many studies and design possibilities that we have not explored. Hopefully, our initial work can suggest some fertile new directions.

7.0 Acknowledgments

Additional contributors to this work include Nancy Vye and John Bransford. Justine Cassell provided invaluable advice. This material is based upon work supported by the National Science Foundation under Grants REC-0231946, REC-0231771. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

8.0 References

Bailenson, J., & Yee, N. (in press). Digital Chameleons: Automatic assimilation of nonverbal gestures in immersive virtual environments. Psychological Science.

Barron, B.J. (2004). When smart groups fail. Journal of the Learning Sciences,12, 307-359.

Baylor, A. L. (2002). Expanding preservice teachers' metacognitive awareness of instructional planning through pedagogical agents. Educational Technology Research & Design, 50(2), 5-22.

Bernieri, F., & Rosenthal, R. (1991). Interpersonal coordination: Behavior matching and interactional synchrony. In R. S. Feldman & B. Rime (Eds.), Fundamentals of nonverbal behavior (pp. 401-432). Cambridge, UK: Cambridge University Press.

Biswas, G., Leelawong, K., Belynne, K., Viswanath, K., Vye, N. J., Schwartz, D. L., & Davis, J. (2004). Incorporating Self Regulated Learning Techniques into Learning by

Teaching Environments. *26th Annual Meeting of the Cognitive Science Society*. Chicago, Illinois, pp. 120-125.

Biswas, G., Schwartz, D. L., Bransford, J. D., & TAG-V. (2001). Technology support for complex problem solving: From SAD Environments to AI. In K. Forbus & P. Feltovich (Eds.), Smart machines in education (pp. 71-98). Menlo Park, CA: AAAI/MIT Press.

Biswas, G., Schwartz, D. L., Leelawong, K., Vye, N., & TAG-V. (in press). Learning by teaching: A new paradigm for educational software. Applied Artificial Intelligence, 19(3).

Blair, K. P., and Schwartz, D. L. (2004) Milo and J-Mole: Computers as constructivist teachable agents. In Y. Kafai et al. (Eds.) Embracing Diversity in the Learning Sciences: The proceedings of the Sixth International Conference of the Learning Sciences (p. 588). Mahwah, NJ: Erlbaum

Bodemer, D., Ploetzner, R., Feuerlein, I., & Spada, H. (2004) The active integration of information during learning with dynamic and interactive visualizations. Learning and Instruction 14(3), 325-341

Bransford, J. D., & Schwartz, D. L. (1999). Rethinking transfer: A simple proposal with multiple implications. In A. Iran-Nejad & P. D. Pearson (Eds.), Review of Research in Education , 24, 61-101. Washington DC: American Educational Research Association.

Brennan, S. E. (1966). Lexical entrainment in spontaneous dialog. Proceedings, 1996 International Symposium on Spoken Dialogue, 41-44.

Cassell, J., (2004). Towards a model of technology and literacy development: Story listening systems. Journal of Applied Developmental Psychology, 25, 75-105.

Chan, T-W. (1995). Social learning systems: an overview. In Collis, B., Davies, G. (eds.), Innovative adult learning with innovative technologies. IFIP Transactions A-61, North-Holland, 101-122.

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

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

Deutsch, M. (1973). The resolution of conflict: Constructive and deductive processes. New Haven, CT: Yale University Press.

Dillenbourg, P. & Self, J. (1992). People Power: a human-computer collaborative learning system. In Gauthier, G. & McCalla, G. (eds.), Intelligent Tutoring Systems. Lecture Notes in Computer Science 608, Springer Verlag, 651-660.

Gordin, D. N., & Pea, R. D. (1995). Prospects for scientific visualization as an educational technology. Journal of the Learning Sciences, 4, 249-280.

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

Hartman, K. & Blair, K. P. (2005). The unless switch: adding conditional logic to concept mapping for middle school students. CHI '05 Extended Abstracts on Human Factors in Computing Systems (pp. 1439-1442). ACM Press, NY, NY.

Hegarty, M. (2004). Dynamic visualizations and learning: getting to the difficult questions. Learning and Instruction, 14, 343-351.

Kafai, Y., & Harel, I. (1999). Learning through design and teaching: Exploring social and collaborative aspects of constructionism. In I. Harel & S. Papert (Eds.), Constructionism. Norwood, NJ: Ablex.

Kinchin, I.M. and D.B. Hay. (2000). How a qualitative approach to concept map analysis can be used to aid learning by illustrating patterns of conceptual development. Educational Research, 42(1), 43-57.

Koedinger, K. R., & Anderson, J. R. (1998). Illustrating principled design: The early evolution of a cognitive tutor for algebra symbolization. Interactive Learning Environments, 5, 161-180.

Kuhn, D. (1995). Scientific thinking and knowledge acquisition. Monographs of the Society for Research in Child Development, 60(4), 152-157.

Lin, X.D., & Lehman, J. (1999). Supporting learning of variable control in a computer based biology environment: effects of prompting college students to reflect on their own thinking. Journal of Research In Science Teaching, 36, 837-858.

Lester, J., Stone, B., & Stelling, G. (1999). Lifelike pedagogical agents for mixed-initiative problem solving in constructivist learning environments. User Modeling and User-Adapted Interaction, *9*(1-2), 1-44.

Lowe, R. (1999). Extraction information from an animation during complex visual learning. European Journal of Psychology in Education, 6(2), 225-244.

Lowe, R. (2004). Interrogation of a dynamic visualization during learning. Learning and Instruction, 14, 257-274.

Mathon, S., & Koedinger, K. R. (2003). Recasting the debate: Benefits of tutoring error detection and correction skills. In U. Hoppe, F. Verdejo, & J. Kay (Eds.), Artificial Intelligence in Education: Shaping the Future of Learning through Intelligent Technologies, Proceedings of AI-ED 2003 (Vol 97, pp. 13-18). Amsterdam, IOS Press.

Moreno, R., Mayer, R.E., Spires, H.A., & Lester, J.C. (2001). The case for social agency in computer-based teaching. Do students learn more deeply when they interact with animated pedagogical agents? *Cognition and Instruction, 19*(2), 177-213.

Novak, J.D. (1996). Concept mapping as a tool for improving science teaching and learning. In D.F. Treagust, R. Duit, and B.J. Fraser, (Eds.) Improving Teaching and Learning in Science and Mathematics (pp. 32-43). London: Teachers College Press.

Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension monitoring activities. Cognition and Instruction, 1, 117-175.

Reeves, B. and Nass C. (1966). The Media Equation: How people treat computers, televisons and new media like real people and places. Cambridge, UK: Cambridge University Press.

Renkl, A. (1995). Learning for later teaching: An exploration of mediational links between teaching expectancy and learning results. Learning and Instruction, 5, 21-36.

Rickel, J. and Johnson, W.L., (1998). STEVE: A Pedagogical Agent for Virtual Reality (video). In Proceedings of the Second International Conference on Autonomous Agents, Minneapolis/St. Paul: ACM Press.

Schwartz, D. L. (1995). The emergence of abstraction in dyad problem solving. Journal of the Learning Sciences.

Schwartz, D. L. (1999). The productive agency that drives collaborative learning. In P. Dillenbourg (Ed.), Collaborative learning: Cognitive and computational approaches (pp. 197-218). New York: Elsevier Science.

Sears, D., & Schwartz, D. L. (April, 2004). Risks and Rewards of Jigsawing Teachable Agents. Paper presentation at AERA. San Diego.

Shimoda, T. A., White, B., Frederisken, J. (2002). Student goal orientation in learning inquiry skills with modifiable software advisors. Science Education, 86, 244-263.

Schnotz, W., Boeckheler, J., & Grzondziel, H. (1999). Individual and co-operative learning with interactive animated pictures. European Journal of Psychology in Education, 14(2), 245-265.

Simon, H. A. (1996). Sciences of the Artificial, 3rd Edition. Cambridge, MA: MIT Press.

Smith, D.C., Cypher, A., & Spohrer, J. (1997). Programming agents without a programming language. In J. M. Bradshaw (Ed.) Software Agents (pp. 165-190). Menlo Park, CA: AAAI/MIT Press.

Steiner, I. (1972). Group process and productivity. NY, NY: Academic Press.

Tomasello, M. & Farrar, M. J. (1986). Joint attention and early language. Child Development, 57, 1454-1463.

Tepper, P., Kopp, S., Cassell, J. (2004). Content in Context: Generating Language and Iconic Gesture without a Gestionary. Proceedings of the Workshop on Balanced Perception and Action in ECAs at AAMAS.

Tversky, B., Morrison, J. B., & Betrancourt, M. (2002). Animation: Can it facilitate? International Journal of Human-Computer Studies, 57, 247-262.

Viswanath, K., Adebiyi, B., Biswas, G., & Leelawong, K. (2004). A Multi-Agent Architecture Implementation of Learning by Teaching Systems. 4th IEEE Intl. Conference on Advanced Learning Technologies (pp. 61-65), Joensuu, Finland.

Uretsi, J. A. R. (2000). Should I teach my computer peer? Some issues in teaching a learning companion. In G. Gautheir, C. Frasson, & K. VanLehn (Eds.). Intelligent Tutoring Systems (pp. 103-112). Berlin: Springer-Verlag.