

CHAPTER 9 –On the Use of Learner Micromodels as Partial Solutions to Complex Problems in a Multiagent, Conversation-based Intelligent Tutoring System

Xiangen Hu, Donald M. Morrison, Zhiqiang Cai
Institute for Intelligent Systems (IIS)
University of Memphis (UM)

Introduction

More than 40 years since their origins in the early days of computer-assisted instruction in the 1970s (e.g., Carbonell, 1970), considerable progress has been made in developing “intelligent” (computer-based) tutoring systems that successfully scaffold learning in specific domains and for specific purposes (for a recent review, see Graesser, Conley & Olney [2012]). Frequently cited examples include Cognitive Tutor, which supports learning in algebra, geometry, and programming languages (Ritter, Anderson, Koedinger & Corbett, 2007); AutoTutor (Graesser, Olney, Haynes & Chipman, 2005; Graesser, Jeon & Dufty, 2008; Graesser, Lu et al., 2004), which does the same thing for college-level computer literacy, physics, and critical thinking skills; and the so-called “constraint-based” systems developed by Mitrovic’s group in New Zealand (Mitrovic, Martin & Suraweera, 2007), which, among other topics, help students learn to program in SQL.

In spite of these advances at local research and development sites, the field has yet to produce a truly general-purpose system that is capable of supporting rapid development of high-quality applications across a broad range of domains. GIFT, currently under development at the U.S. Army’s Learning in Intelligent Tutoring Environments (LITE) Laboratory, is intended to fill this gap. Citing Picard (2006), the developers claim that the availability and use of ITSs has been limited by the high cost of development, lack of reusability, lack of standards, and “inadequate adaptability to the needs of learners” (Sottolare, Brawner, Goldberg & Holden, 2012:1). These systems, they write, tend to be built as “domain-specific, unique, one-of-a-kind, largely domain-dependent solutions focused on a single pedagogical strategy.” GIFT is presented as a solution to this problem. The modular framework and standards built into the system could “enhance reuse, support authoring and optimization of CBTS³ strategies for learning, and lower the cost and skill set needed for users to adopt CBTS solutions for military training and education” (Sottolare et al., 2012:1).

The development of a general-purpose, domain-specific ITS framework is indeed an important goal, but numerous barriers block the way. These include (but are not restricted to) a lack of agreement in the ITS community about how the different components of an ITS ought to fit together; what the structure and content of the components ought to be; and how knowledge of the world is to be represented, both for experts and learners.

The argument made in this chapter assumes that a general-purpose system will at some point employ an open, multiagent architecture, meaning that core functions are carried out by more or less autonomous software agents united by a common ACL. Some of these agents will perform simple tasks (such as analyzing a learner’s facial expressions), while others will take on more complex ones, such as generating appropriate responses to user questions. As an example, here we describe an autonomous software agent that produces a turn-by-turn analysis of a user’s discourse moves on two dimensions: *relevance* and *novelty* (R-N). In the process, it builds what we call a *micromodel* of the learner’s current state, including

³ Here we use the term ITS to mean the same thing as a computer-based tutoring system (CBTS), a class of adaptive educational system (AES).

a relevance-novelty measure for single turns and for a series of turns.⁴ This micromodel allows the R-N agent to make assertions about the learner’s current and recent contributions to a conversation—assertions which may be broadcast generally or addressed directly to other agents, such as a conversation agent, a pedagogical agent, an agent responsible for constructing aggregate learner models from multiple micromodels, an agent that analyzes the effectiveness of instructional modules, or, where the ITS employs an “open” learner model (Bull, 2004; Bull & Pain, 1995; Kay, 2001; Mitrovic & Martin, 2002), an agent responsible for providing access to the learner model through the user interface.

The Standard Four-Component ITS Model

As noted elsewhere in this book, it is customary to identify an ITS as consisting of four major components, referred to as “models” – the learner model (sometimes called the “student model”); the expert domain model; the tutor model, and the user interface (Elson-Cook, 1993; Graesser et al., 2012; Nkambou, Mizoguchi & Bourdeau, 2010; Psozka, Massey & Mutter, 1988; Sleeman & Brown 1982; VanLehn, 2006; Woolf, 2008). In an ITS where the tutor is capable of mixed-initiative dialog with the user (Carbonell, 1970; Allen, Guinn & Horvitz, 1999; Graesser et. al, 2005), the tutor takes the form of an intelligent conversation agent, backed by a Dialog Advancer Network (Person, Bautista, Kreuz, Graesser & Tutoring Research Group, 2000). Figure 9-1 illustrates the relationship among these four components.




Figure 9-1. Standard ITS components

In general terms, the learner model represents what the tutor has established to be the learner’s current level of knowledge, skill, and affective state, while the expert domain model represents the knowledge and skills the learner is supposed to acquire – and has therefore been called the “ideal student model” (Corbett, Koedinger & Anderson, 1997). Through interactions with the learner via the user interface, the conversation agent, playing the role of the tutor, seeks in some way to bring the learner model in line with the expert model. In this sense, the learner model is said to be an “overlay” of the expert model (Wenger, 1987)

Although the various ITS research and development communities seem to agree that these are the main components, the model is really more of a conceptual framework than a working blueprint. In practice, different systems employ quite different architectures, data structures, and strategies, reflecting different instructional philosophies and purposes (e.g., Nkambou et al., 2010, Schatz & Folsom-Kovarik, 2011). While this makes sense locally, the lack of a standard overall system architecture and way of constructing the different system components is problematic for a number of reasons.

For one thing, it means that components that have proven to be effective in one system are not easily imported into another, thus limiting progress that might be made through the collective efforts of the rapidly expanding network of ITS research and development groups around the world. Also, the lack of a standard method of structuring the learner model means that when a learner moves from one system to the

⁴ Although we refer here to “user” as a “learner,” in fact the agent we describe here is capable of evaluating the discourse moves of any interlocutor, including those of another agent.

Design Recommendations for Intelligent Tutoring Systems - Volume 1: Learner Modeling

next, the new system must start from scratch in establishing the learner's history and current state of knowledge. To use an analogy from medicine, it is as if different doctors had different languages for describing a patient's health, so, when dealing with a new patient, each doctor would have to reassemble a patient's history from scratch.

What then, in the absence of a standard architecture and data structure for representing learner and expert domain knowledge, is one do to? The 60s-era slogan "If you're not part of the solution, you're part of the problem" seems relevant here. Specific solutions to general problems of learner modeling ought to be crafted in such a way that they are generally useful no matter what environment they are asked to work in. To return to the medical records analogy, if we are developing a new procedure for say, measuring pupil dilation or "knee-jerk" reflex response, so long as we have a standard way of reporting our results, we don't have to worry about how our "microrecord" fits into the overall structure of the patient's medical record, which can be assembled by someone else. In other words, we can be part of a solution without knowing exactly what that solution is.

The Argument for a Multiagent Architecture

The notion of an ITS as a multiagent system is not new. Of course, any ITS is a multiagent system in the gross sense that there are two autonomous agents at work: the user and intelligent tutor. However, recent years have seen an increasing emphasis on development of ITSs with multiagent architectures in the more interesting sense that overall system functionality emerges from the collective work of individual software agents (Bittencourt et al., 2007; Chen & Mizoguchi, 2004; El Mokhtar En-Naimi, Amami, Boukachour, Person & Bertelle, 2012; Lavendelis & Grundspenkis, 2009; Zouhair et al., 2012). Through the use of a shared, speech-act-based agent communication framework such as Knowledge Query and Manipulation Language (KQML) (Finin, Fritzson, McKay & McEntire, 1994), Foundation for Intelligent Physical Agents (FIPA)-ACL (O'Brien & Nicol, 1998), or Java Agent Development Framework (JADE) (Bellifemine, Caire, Poggi & Rimassa, 2008), combined with a set of domain-specific ontologies (concepts and their relations specific to the system domain), the agents in the system assert beliefs, make requests of other agents, deny requests, and so forth, much as human workers in a large collective enterprise do (see Chaib-draa & Dignum, 2002; Kone, Shimazu & Nakajima, 2000).

For example, in a multiagent ITS, different agents can take on the different tasks of user registration and authentication, interfacing with learning management systems, building learner models dynamically, monitoring learner affect through the use of various sensing systems, and managing conversations among users and other agents. Further, agents with especially complex tasks, such as a conversation agent, may be supported by a network of specialized agents dedicated to specific subtasks. Each of these agents can have its own internal algorithms, data structures, and local methods of obtaining data. Importantly, as long as the agent "knows" the system's ACL, i.e., can post, send, and read messages in a shared language, it doesn't matter how it is organized internally, in the same way that different sort functions can take the same input and produce the same output using different internal algorithms.

In the remainder of the chapter, we give an example: an agent that is capable of evaluating a learner's discourse moves on two important dimensions: *relevance* and *novelty*. Instantiated as a highly specialized agent within a multiagent, conversation-based ITS architecture, the R-N agent is capable of making assertions about a small but important piece of the learner model (a learner micromodel), information that may be of some value to other agents, such as conversation agents and learning model agents, for their own purposes.

The Problem of Conversational Relevance and Novelty

Designing a computer program that can carry on a conversation with a human, one of the oldest and arguably the hardest challenge in AI, is exactly the sort of problem that lends itself to multiagent treatments. For example, different agents may be responsible for converting speech to text, parsing text into its grammatical elements, classifying utterances as different kinds of speech acts, and extracting (or estimating) meaning through some form of semantic analysis. Other agents, or clusters of agents, can be responsible for generating responses in the form of text strings, while still other agents convert the text strings into speech, and have them spoken by an animated avatar.

The agent responsible for generating responses to a user's utterances (discourse moves) arguably has the most complex task and is most likely to benefit from the assistance of simpler agents that can take on pieces of it. The task is hard because it is really a special form of mind-reading, requiring an ability to continuously create and test theories about an interlocutor's present state of mind. Bakhtin's distinction between monologic and dialogic discourse (Bakhtin, 1981; Wells, 2007) is a useful way of framing the problem. In a conversation that is primarily monologic, the speaker's purpose is to convey information in such a way as to "duplicate one's own idea in someone else's mind" (Bakhtin, 1986:69), without the need to be concerned about what is already in the listener's mind, what the listener may be thinking at the present moment or in expectation of a particular response. This is the discourse stance of a lecturer. Dialogic discourse, on the other hand, is inherently a "social form of thinking" (Wells, 2007:256), a much harder form of discourse, at least for machines, in which interlocutors must work to understand each other's "present state," and thereby arrive at a shared understanding.⁵ Consider, for example, the following exchange:

A: So, what did you do today?

B: Attended a physics lecture.

A: What did you learn?

B: It was really hot.

Given such a response, if A is a human speaker, A will assume, on the Gricean principle, that B is a cooperative interlocutor; that B's response must be in some way *relevant* (Grice, 1975); and that "it" refers either to the lecture or the lecture hall. In the former case, "hot" would be an attribute of the lecture, representing a positive reaction, implying that speaker B had liked the lecture and possibly learned a lot from it. In the latter case, "hot" would be an attribute of the lecture hall, implying an uncomfortable temperature. Because the word "hot" is more commonly associated with rooms than lectures, A would probably test the latter case first:

...

A: So, are you saying the lecture hall was uncomfortable?

B: Yes.

A: I'm sorry to hear that. It must have been hard to concentrate. What were you able to learn?

B: We learned about physics.

A: Okay, but what about physics?

⁵ In a paper titled "Why is conversation so easy?" Garrod & Pickering (2004) argue that, for humans, dialog is easier than monologue because interlocutors automatically align linguistic representations at various levels (phonological, syntactic, semantic, and situational), thus building up a shared workspace. This joint construction of meaning and purpose has the effect of distributing the processing load, thus making conversation relatively easy. Humans, as they put it, are "designed for dialog."

Design Recommendations for Intelligent Tutoring Systems - Volume 1: Learner Modeling

In order for an intelligent conversation agent to carry on such a conversation, it would need, among many other things, to have some knowledge of the world, of the following type:

1. Lectures are a method of teaching.
2. Lectures are about something.
3. A person may learn something from a lecture.
4. Lectures take place in rooms called lecture halls.
5. Rooms may be comfortable or uncomfortable.
6. Humans are sensitive to temperatures that are outside their comfort range.
7. “Hot” refers to a temperature that is outside a human’s comfort range.
8. Learning requires concentration.
9. When a person is uncomfortable, it is hard to learn.
10. “Hot” is a slang word for something that a human finds attractive...

and so forth. In other words, A’s ability to judge the *relevance* of B’s utterance “It was very hot” depends very much on a complex set of concepts and relationships.

However, relevance is not the only important measure of the degree to which a discourse move is felt to be cooperative. In addition to the principle of relevance is that of *quality* (Grice, 1975), which includes the assumption that an interlocutor’s move will add something *new* to the conversation. In the imagined conversation considered here...

Learner (B): We learned about physics.

Tutor (A): Okay, but what *about* physics?

A’s “Okay, but what *about* physics?” is exactly the right thing to say because B has already reported that it was a *physics* lecture, so of course it was about physics. In other words, B has violated the Gricean maxim of quality by failing to make a truly novel contribution.

So, if it is to possess anything remotely like the intelligence of a human speaker, it seems that an intelligent conversation agent must have a way of evaluating both the relevance and novelty of an interlocutor’s discourse moves. Admittedly, this is just one part of the problem of natural language understanding by a machine, perhaps even a minor part, and so a solution, which itself is likely to be only partial, will be worth little if not combined with many others. That said, within a multiagent environment, it might be useful to have an agent that is capable of sending messages like:

```
json={"target":"conversation manager","learner":"5021","time":"2013-02-01  
10:45:22 UTC", "novelty":" .3","relevance":".5"}
```

or

Design Recommendations for Intelligent Tutoring Systems - Volume 1: Learner Modeling

```
json={"target":"learner model manager","learner":"1259",  
"session":"5","date": "2013-02-01","average novelty":"medium", "average  
relevance":"high"}
```

where the language is an ACL, the decimals represent Bayesian estimates of likelihood, and the assertions are addressed to other agents with a “for your information” (FYI) intention, i.e., the originating agent is agnostic as to how other agents will use the information (beliefs) it asserts.⁶

An Example from Physics

As an example, assume that an intelligent tutor is provided with the following instructional script:

ASK: *Suppose a football player is running in a straight line at a constant speed, while carrying a football. Suddenly he throws the football straight up, over his head, and keeps running at the same speed. Explain where it will land.*

EXPECTED ANSWER: *In accordance with Newton’s first law of motion, the ball will continue to travel in the same direction, unless there is an opposing force. In this case, the main opposing force is gravity, but gravity only affects the vertical component of the motion vector. It is true that the particles in the air, such as oxygen molecules, exert an opposing force, which pushes against the ball in the opposite direction, but this force is very small compared to the accelerating force of gravity.*

Now, compare the following two transcripts from an imaginary tutoring session:

Transcript A

Tutor: Suppose a football player is running in a straight line at a constant speed, while carrying a football. Suddenly he throws the football straight up, over his head, and keeps running at the same speed. Explain where it will land.

Learner A: It will continue along in his direction and come back down where he can catch it.

Tutor: I agree. But why?

Learner A: Well, the ball and the runner both have the same forward speed.

Tutor: Okay...and....

Learner A: So when gravity pulls the ball down, it just affects the vertical speed.

Tutor: What about the horizontal speed?

Learner A: It remains the same. It’s Newton’s second law of motion.

Tutor: Do you mean Newton’s first law?

Learner A: Oh right. First law.

Tutor: Good. Let’s go on to the next problem. What if, instead of a ball, it’s a balloon? On the moon?

Transcript B

Tutor: Suppose a football player is running in a straight line at a constant speed, while carrying a football. Suddenly he throws the football straight up, over his head, and keeps running at the same speed. Explain where it will land.

Learner B: i don’t know. somewhere behind him.

Tutor: Why?

Learner B: cause it will fall back behind him so he can’t catch it.

⁶ The syntax used here is JavaScript Object Notation (json), a message format used for communication between agents in multiagent systems. See <http://www.json.org/>.

Tutor: Why do you think that?
Learner B: i just think it will fall behind him
Tutor: Okay...but why?
Learner B: cause
Tutor: Okay, let's look at that video again....

Clearly, there is quite a lot going on here; however, a key difference between the two transcripts is that whereas Learner A's contributions are all *relevant* (consistent with the expected contribution) and *novel* (providing additional information); in the same sense, only Learner B's first contribution is novel, and none are relevant.⁷ More generally, the relevance and novelty of given discourse move may be mapped onto a Venn diagram such as that shown in

Figure 9-2.

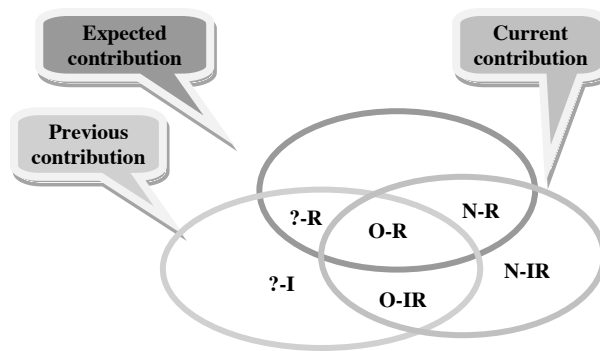


Figure 9-2. Mapping relevance and novelty

This gives a matrix (Table 9-1) with four quadrants:

Table 9-1. Novelty and Relevance Matrix

	Old	New
Relevant	O-R (Not new but relevant)	N-R (New and relevant)
Irrelevant	O-IR (Not new and irrelevant)	N-IR (New but irrelevant)

Assuming that this kind of information would be useful to other agents (notably the agent handling the conversation with the learner), we can now ask how an agent might go about determining the relevance and novelty of a given utterance.

⁷ Note that we are using the term “relevant” here in a special, non-intuitive sense. Whereas a contribution may be “relevant” in the sense that it relates in some way to the topic, it is considered irrelevant if it is inconsistent with a model answer.

Quantitative Measures for Novelty and Relevance

A rough-and-ready relevance measure for a given learner contribution to a dialog with an intelligent tutor can be defined as the extent to which the learner’s answer to a tutor’s question is “semantically similar” to the answer the tutor expects. In the same way, a measure of the novelty of the learner’s most recent contribution can be defined as the extent to it resembles the learner’s previous contributions to the same conversation, i.e., attempts to answer the same question. In other words, if the R-N agent is passed two strings – one representing the expected answer and the other the learner’s most recent contribution – then it can come up with a relevance score using any one of several methods used to compute semantic similarity, including Latent Semantic Analysis (LSA; Landauer & Dumais, 1997), Hyperspace Analogue to Language (HAL; Burgess, Livesay & Lund, 1996), Latent Dirichlet Allocation (LDA; Blei, Ng & Jordan, 2003), Non-Latent Similarity (NLS; Cai et al., 2004); Word Association Space (WAS; Steyvers, Shiffrin & Nelson, 2002), and Pointwise Mutual Information (PMI; Recchia & Jones, 2009). For a discussion of the use of LSA in ITSs see Hu et al. (2007).

Also, so long as it knows that the topic has not changed (e.g., the tutor is still prompting for an answer to the same question), then, using the same method, it can calculate the semantic similarity of the learner’s most recent contribution to her previous contributions. This produces values seen in Table 9-2.

Table 9-2. Sample relevance and novelty measures

	Old	New
Relevant	0.4 (O-R)	0.2 (N-R)
Irrelevant	0.1 (O-IR)	0.3 (N-IR)

In addition, two other measures are obtained by combining current and previous contributions. A Current Relevant Contribution (CRC) score is obtained by adding O-R and N-R (in this case, 0.6), while the CRC score combined with all previous CRC scores gives a Total Coverage (TC) score.

Together, for any given dialog move, these six measures may be viewed as constituting a micromodel of the learner’s “current state.” An agent that is capable of evaluating a given dialog move on these measures can pass along the micromodel it has built for the use of other agents in the community. For example, in Transcript A, an agent’s analysis of Learner A’s contribution “Well, the ball and the runner both have the same forward speed” might be communicated as follows:

```
json={"target":"all","learner":"1259","input string":"Well, the ball and the runner both have the same forward speed","time":"2013-02-05 11:25:27 UTC","R/N":"0.3","R/O":"0.4", "I/N":"0.28","I/O":"0.01", "CRC":"0.17","TC":"0.24"}
```

...whereas Learner B’s contribution “cause” could yield the following:

```
json={"target":"all","learner":"1147","input string":"cause","time":"2013-02-05 11:25:27 UTC","R/N":"0.0", "R/O":"0.0","I/N":"0.0","I/O":"0.10","CRC":"0.04","TC":"0.033"}
```

Over a series of moves, the cumulative scores constitute what Hu & Martindale (2008) refer to as a Learner Characteristic Curve (LCC), which may be displayed in the form of a set of graphs, as illustrated in Figure 9-3.

Design Recommendations for Intelligent Tutoring Systems - Volume 1: Learner Modeling

This information, the relevance and novelty of a given utterance, combined with the cumulative relevance measures, can be viewed as constituting a small, localized micromodel of the learner's current cognitive state and as such has practical utility. In fact, it is on the same level as micromodels developed by other agents in a multiagent ITS community that provides both real-time and historical information about a user's apparent affect.




Figure 9-3. Sample LCC output

For example, an agent that monitors a learner's facial expressions might make an assertion in the form:

```
json={"target":"all","learner":"1147","facial  
expression":"puzzled","value":".3","time":"2013-02-01 10:45:22 UTC"}
```

which constitutes a component micromodel of the learner's affective state. Combined with the preceding message from the R-N agent, the conversation manager (more specifically, a response generation agent) now has two pieces of evidence to consider, i.e., that the user appears confused and that the novelty and relevance measures for the user's most recent utterance were both low. Given this information, the response agent could decide to send a message such as this to an avatar:

```
json={"target":"avatar","learner":"1147","output string":"You seem confused.  
Are you?","time":"2013-02-01 10:45:24"}
```

Assuming they have ways of understanding messages such as these, other agents in the system can use them for their own purposes. For example, a response generation agent might use the information in the first message to generate the turn:

Learner A: Well, the ball and the runner both have the same forward speed.

Tutor: Okay...and....

and the information in the second message to produce:

Learner B: cause

Tutor: Okay, let's look at that video again....

As another example, a pedagogical agent might take an LCC representing repeated non-novel, "irrelevant" learner contributions as evidence of a possible misconception.

Opening the Model to the Learner

In a system where the learner model is “open” to the learner (Bull, 2004; Bull & Pain, 1995; Dimitrova et al, 2001; Kay, 1997; Mitrovic & Martin, 2002), a user interface agent might use the micromodel to create a set of graphs, as in Figure 9-4.

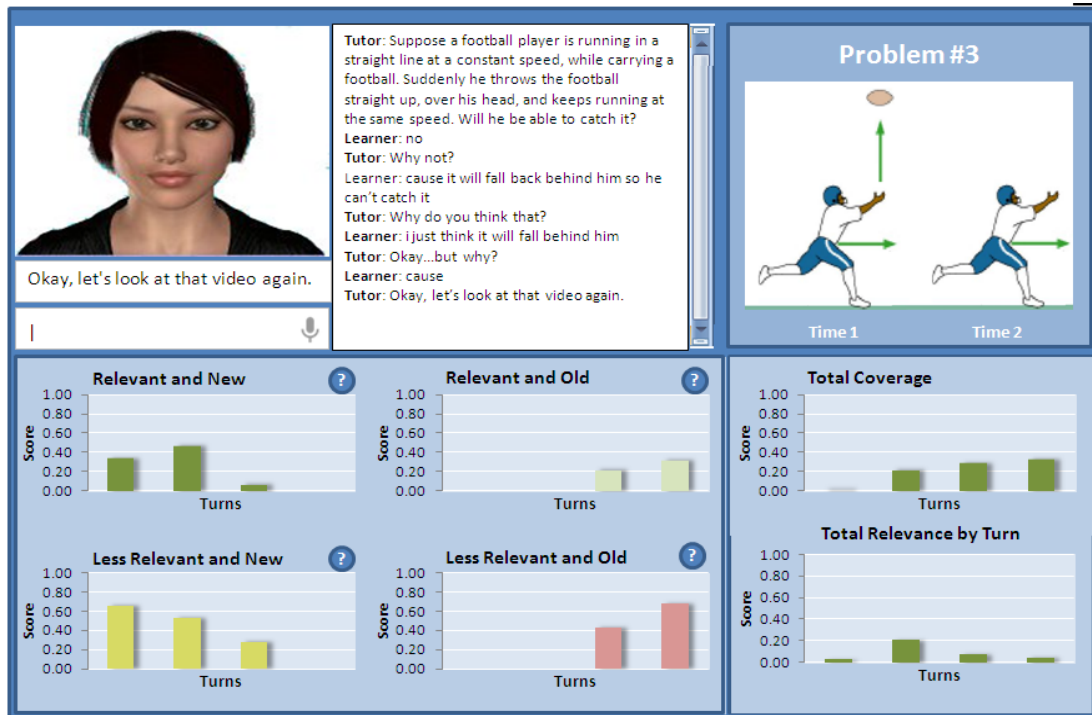


Figure 9-4. Opening the micromodel to the learner

Giving the learner feedback on the relevance and novelty of her discourse moves in this way could conceivably encourage her to focus her attention on maintaining higher levels of relevance and novelty than she might otherwise, thereby increasing the likelihood that the conversation will lead to real learning.

Discussion

In this chapter, we have explained how a specialized agent within a conversation-based ITS can monitor a learner’s discourse moves on two dimensions: novelty and relevance. In this way, over a series of moves, it builds up a “micromodel” of the learner’s cognitive state, called a Learner Characteristic Curve (LCC), which it can then pass along in the form of messages (assertions of belief) to other agents, assuming a common ACL. As a result, the agent can contribute to the solutions of larger problems without needing to know what the solution is. Importantly, such an agent is both reusable and replaceable. It is reusable in the sense that it can be used in any number of different ITSs where measures of novelty and relevance are considered useful in some way. It is replaceable in the sense that another agent that performs the same analysis, but more effectively, could be brought in to take over. This form of loose coupling (Orton & Weick, 1990; Weick, 1976) allows for the rise of mutations at both the local (agent) and system

(multiagent) levels, thus allowing gradual evolution toward increasingly sophisticated systems that may eventually approach the effectiveness of a highly skilled human teacher.

What, then, are the implications of the preceding argument for a “generalized intelligent framework” like GIFT? More precisely, how can we, as a community of ITS researchers and developers, possibly move forward from our current world of “domain-specific, unique, one-of-a-kind, largely domain-dependent solutions focused on a single pedagogical strategy” (Sottolare, Brawner, Goldberg & Holden, 2012:1) toward a future of open, domain-independent systems with shareable, reusable tools and components, efficient authoring, transportable learner models, cross-platform functionality, and so forth. Given our rapidly evolving technological environment (e.g., the sudden ubiquity of smartphones and tablets, the explosion of massive, text-based knowledge representations in the Semantic Web, the magnetic attraction of social media such as Facebook and YouTube, etc.), it seems unlikely that a single, intelligent tutoring solution will ever be more than temporarily useful. Rather, it seems what we can look forward to, and should build toward as a community of practice, is some evolving collection of workable, partial, and loosely coupled solutions, which, while provisional, are built in such a way that they can evolve both with and apart from each other. Specifically, the recommendation is that we begin to think seriously about the adoption of a common ACL such as KQML (Finin, Fritzson, McKay & McEntire, 1994), FIPA-ACL (O’Brien & Nicol, 1998), or JADE (Bellifemine, Caire, Poggi & Rimassa, 2008) as the basis for a new generation of agent-based intelligent learning systems that are capable of *autonomous cooperation* (Hülsmann, Scholz-Reiter, Freitag, Wucisk & De Beer, 2006; Windt, Böse & Philipp, 2005). Importantly, because the contributing agents can have their own internal databases and algorithms, we can move in this direction without wholesale reengineering of our existing systems. Rather, once we agree on a common ACL and begin to build a shared ontology (which can occur incrementally), then we can “simply” encapsulate our existing (and evolving) systems in wrappers that allow these systems, whatever their function, to take part in whatever communities may arise, both adding and receiving value, in ways that are now only dimly imaginable.

References

- Allen, J. E., Guinn, C. I. & Horvitz, E. (1999). Mixed-initiative interaction. *Intelligent Systems and their Applications, IEEE, 14*(5), 14-23.
- Bakhtin, M. M. [1930s] 1981. *The dialogic imagination: Four essays*. Ed. Michael Holquist. Trans. Caryl Emerson and Michael Holquist. Austin and London: University of Texas Press.
- Bellifemine, F., Caire, G., Poggi, A. & Rimassa, G. (2008). JADE: A software framework for developing multiagent applications. *Lessons learned. Information and Software Technology, 50*(1), 10-21.
- Bittencourt, I. I., de Barros Costa, E., Almeida, H. D., Fonseca, B., Maia, G., Calado, I. & Silva, A. D. (2007). Towards an ontology-based framework for building multiagent intelligent tutoring systems. In *Workshop on software engineering for agent-oriented systems, III, João Pessoa, 2007*. Porto Alegre, SBC, 53-64.
- Blei, D. M., Ng, A. Y. & Jordan, M. I. (2003). Latent dirichlet allocation. *The journal of machine learning research, 3*, 993-1022.
- Bull, S. & Pain, H. (1995). Did I say what I think I said, and do you agree with me? Inspecting and questioning the student model. In J. Greer (Ed.), *Proceedings of World Conference on Artificial Intelligence in Education, Association for the Advancement of Computing in Education*, 501-508
- Bull, S. (2004). Supporting learning with open learner models. *Planning, 29*(14), 1.
- Burgess, C. & Lund, K. (1997). Modeling parsing constraints with high-dimensional context space. *Language and Cognitive Processes, 12*, 177-210.

Design Recommendations for Intelligent Tutoring Systems - Volume 1: Learner Modeling

- Cai, Z., McNamara, D. S., Louwerse, M., Hu, X., Rowe, M. & Graesser, A. C. (2004). Nls: Non-latent similarity algorithm. In K. Forbus, D. Gentner & T. Regier (Eds.), *Proceedings of the 26th annual meeting of the cognitive science society* (p. 180-185). Mahwah, NJ: Erlbaum.
- Carbonell, J. R. (1970). AI in CAI: Artificial intelligence approach to computer assisted instruction. *IEEE Transactions on Man-Machine Systems* 11(4): 190-202.
- Chaib-draa, B. & Dignum, F. (2002). Trends in agent communication language. *Computational Intelligence*, 18(2), 89-101.
- Chen, W. & Mizoguchi, R. (2004). Learner model ontology and learner model agent. *Cognitive Support for Learning-Imagining the Unknown*, 189-200.
- Corbett, A. T., Koedinger, K. R. & Anderson, J. R. (1997). Intelligent tutoring systems. *Handbook of human-computer interaction*, 849-874.
- El Mokhtar En-Naimi, A. Z., Amami, B., Boukachour, H., Person, P. & Bertelle, C. (2012). Intelligent Tutoring Systems Based on the Multiagent Systems (ITS-MAS): The Dynamic and Incremental Case-Based Reasoning (DICBR) Paradigm. *IJCSI International Journal of Computer Science Issues* 9(6), 112-121.
- Elson-Cook, M. (1993). Student modelling in intelligent systems. *Artificial intelligent review*, 7(3-4), 227-240.
- Finin, T., Fritzson, R., McKay, D. & McEntire, R. (1994, November). KQML as an agent communication language. In *Proceedings of the third international conference on Information and knowledge management* (pp. 456-463). ACM.
- Garrod, S. & Pickering, M. J. (2004). Why is conversation so easy? *Trends in cognitive sciences*, 8(1), 8-11.
- Graesser, A. C., Conley, M. W. & Olney, A. (2012). Intelligent tutoring systems. *APA handbook of educational psychology*. Washington, DC: American Psychological Association.
- Graesser, A. C., Jeon, M. & Dufty, D. (2008). Agent technologies designed to facilitate interactive knowledge construction. *Discourse processes*, 45, 298-322.
- Graesser, A. C., Olney, A. M., Haynes, B. C. & Chipman, P. (2005). AutoTutor: A cognitive system that simulates a tutor that facilitates learning through mixed-initiative dialogue. In C. Forsythe, M. L. Bernard & T. E. Goldsmith (Eds.), *Cognitive systems: Human cognitive models in systems design*. Mahwah, NJ: Lawrence Erlbaum.
- Graesser, A.C., Lu, S., Jackson, G.T., Mitchell, H., Ventura, M., Olney, A. & Louwerse, M.M. (2004). AutoTutor: A tutor with dialogue in natural language. *Behavioral research methods, instruments, and computers*, 36, 180-193.
- Grice, H. Paul. "Logic and conversation. *Speech acts*, ed. by Peter Cole and Jerry Morgan, 41-58." (1975).
- Hu, X., Cai, Z., Wiemer-Hastings, P., Graesser, A. C. & McNamara, D. S. (2007). Strengths, limitations, and extensions of LSA. *The handbook of latent semantic analysis*, 401-426.
- Hu, X. & Martindale, T. (2008, January). Enhance learning with ITS style interactions between learner and content. In *The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC)* (Vol. 2008, No. -1). National Training Systems Association.
- Hülsmann, M., Scholz-Reiter, B., Freitag, M., Wucisk, C. & De Beer, C. (2006). Autonomous cooperation as a method to cope with complexity and dynamics?—A simulation based analyses and measurement concept approach. In Y. Bar-Yam (Ed.), *Proceedings of the International Conference on Complex Systems (ICCS 2006)*. Boston, MA, USA (Vol. 2006).
- Kay, J. (2001). Learner control. *User Modeling and User-Adapted Interaction*, 11(1), 111-127.
- Kone, M. T., Shimazu, A. & Nakajima, T. (2000). The state of the art in agent communication languages. *Knowledge and Information Systems*, 2(3), 259-284.
- Landauer, T. K. & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review; Psychological Review*, 104(2), 211.

Design Recommendations for Intelligent Tutoring Systems - Volume 1: Learner Modeling

- Lavendelis, E. & Grundspenkis, J. (2009). Design of multiagent based intelligent tutoring systems. *Scientific Journal of Riga Technical University. Computer Sciences*, 38(38), 48-59.
- Mitrovic, A., Martin, B. & Suraweera, P. (2007) Intelligent tutors for all: Constraint-based modeling methodology, systems and authoring. *IEEE Intelligent Systems*, 22(4), 38-45.
- Nkambou, R., Mizoguchi, R. & Bourdeau, J. (2010). *Advances in intelligent tutoring systems*. Heidelberg: Springer.
- O'Brien, P. D. & Nicol, R. C. (1998). FIPA—towards a standard for software agents. *BT Technology Journal*, 16(3), 51-59.
- Orton, J. D. & Weick, K. E. (1990). Loosely coupled systems: A reconceptualization. *Academy of management review*, 203-223.
- Person, N. K., Bautista, L., Kreuz, R. J., Graesser, A. C. & Tutoring Research Group. (2000). The dialog advancer network: A conversation manager for AutoTutor. In *ITS 2000 Proceedings of the Workshop on Modeling Human Teaching Tactics and Strategies* (pp. 86-92).
- Picard, R. (2006). Building an affective learning companion. Keynote address at the *8th International Conference on Intelligent Tutoring Systems*, Jhongli, Taiwan.
- Psofka, J., Massey, L. D. & Mutter, S. A. (Eds.). (1988). *Intelligent tutoring systems: Lessons learned*. Lawrence Erlbaum.
- Recchia, G. & Jones, M. N. (2009). More data trumps smarter algorithms: Comparing pointwise mutual information with latent semantic analysis. *Behavior research methods*, 41(3), 647-656.
- Ritter, S., Anderson, J. R., Koedinger, K. R. & Corbett, A. (2007). Cognitive Tutor: Applied research in mathematics education. *Psychonomic bulletin & review*, 14(2), 249-255.
- Schatz, S. & Folsom-Kovarik, J. T. (Winter, 2011). Return on Investment: A practical review of learner modeling techniques. *M&S Journal*. www.msco.mil
- Sleeman D. & J. S. Brown. (1982). (Eds.). *Intelligent tutoring systems*. Orlando, Florida: Academic Press.
- Sottolare, R.A., Brawner, K.W., Goldberg, B.S. & Holden, H.K. (2012). *The Generalized Intelligent Framework for Tutoring (GIFT)*. Orlando, FL: U.S. Army Research Laboratory – Human Research & Engineering Directorate (ARL-HRED).
- Steyvers, M., Shiffrin, R. M. & Nelson, D. L. (2002). Semantic spaces based on free association that predict memory performance. *Triple Festschrift honoring Lyle Bourne, Walter Kintsch, and Tom Landauer*.
- VanLehn, K. (2006). The behavior of tutoring systems. *International journal of artificial intelligence in education*, 16(3), 227-265.
- Weick, K. E. (1976) Educational organizations as loosely coupled systems. *Administrative Science Quarterly*, 21, 1-19.
- Wells, G. (2007). Semiotic mediation, dialogue and the construction of knowledge. *Human Development*. 50(5). pp 244-274.
- Wenger, E. (1987). *Artificial intelligence and tutoring systems—Computational and Cognitive Approaches to the Communication of Knowledge*. Los Altos, CA: Morgan Kaufmann.
- Windt, K., Böse, F. & Philipp, T. (2005). Criteria and application of autonomous cooperating logistic processes. In *Proceedings of the 3rd International Conference on Manufacturing Research. Advances in Manufacturing Technology and Management*. Gao, JX and Baxter, DI and Sackett, PJ (Eds.).
- Woolf, B. P. (2008). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann.
- Wolfe, C. R., Fisher, C. R., Reyna, V. F. & Hu, X. (2012). Improving internal consistency in conditional probability estimation with an Intelligent Tutoring System and web-based tutorials. *International Journal of Internet Science*, 7, 38-54.

Design Recommendations for Intelligent Tutoring Systems - Volume 1: Learner Modeling

Zouhair, A., En-Naimi, E. M., Amami, B., Boukachour, H., Person, P. & Bertelle, C. (2012, October). Intelligent tutoring systems founded on the multiagent incremental dynamic case based reasoning. In *Information Science and Technology (CIST), 2012 Colloquium* (pp. 74-79). IEEE.