

# Nonverbal Action Selection for Explanations Using an Enhanced Behavior Net

Javier Snaider<sup>1</sup>, Andrew M. Olney<sup>2</sup>, and Natalie Person<sup>3</sup>

<sup>1</sup> Computer Science Department & Institute for Intelligent Systems, University of Memphis, Memphis, TN, USA

<sup>2</sup> Institute for Intelligent Systems, University of Memphis, Memphis, TN, USA

<sup>3</sup> Department of Psychology, Rhodes College, Memphis, TN, USA

{jsnaider, aolney}@memphis.edu, person@rhodes.edu

**Abstract.** In this paper we present a novel approach to the nonverbal action selection problem for an agent in an intelligent tutoring system. We use a variation of the original Maes' Behavior Net that has several improvements that allow modeling action selection using the content of the utterance, communicative goals, and the discourse history. This Enhanced Behavior Net can perform action selection dynamically, reprioritize actions based on all these elements, and resolve conflict situations without the use of sophisticated predefined rules.

**Keywords:** Nonverbal action selection - intelligent tutoring system – behavior net – gestures.

## 1 Introduction

Pedagogical agents, like human tutors, frequently must give explanations, provide feedback, and refer to external resources when teaching students [1][2][3]. The synchronization of nonverbal behaviors (e.g., gestures) with dialogue can direct more effectively student attention and intensify engagement [4][5]. However, when there are a large number of behaviors to choose from, the problem of selecting the optimal behavior becomes nontrivial [6]. In this study, we combine previous research on action selection to the problem of animation selection for a pedagogical agent that gives explanations, provides feedback to students' contributions, and refers to a workspace that displays images.

The application context for the present study is the Guru intelligent tutoring system (ITS) for biology [7]. Our approach is similar to previous approaches for generating nonverbal behaviors [8][9][10]. However, we make three distinct contributions. First, our approach involves coordinating all the events that are afforded in a multimedia display. Images, diagrams, dialogue, and text are presented to the user on the multimedia display. Second, the agent's behaviors must be pedagogically appropriate and tailored to the student's current understanding of the material. Hence, rather than using only surface features of a text, as in the previous work generating nonverbal behaviors mentioned above, we are also including information about what words and

concepts are pedagogically relevant to a student at a particular point in time. Thirdly, the agent must perform action selection dynamically by keeping track of the discourse history and reprioritizing actions based on that history.

An enhanced version of the Maes' Behavior Net [11] was used to address the challenge of synchronizing verbal, nonverbal, and multimedia outputs in pedagogical explanations. This approach has several advantages: for example, the BN can automate the process at a tremendous time savings in authoring effort. Also, the BN dynamically reacts to the conditions of the current student resulting in tailored instruction and synchronized deployment of verbal, nonverbal, and multimedia elements to that student.

The following section describes gestures in expert human tutoring in some detail; these gestures represent the standard behaviors that the BN should implement. The subsequent section describes the actions and goals for nonverbal behaviors. Next, the Enhanced Behavior Net (EBN) algorithm is described along with the changes we made in the data structure and in the calculation of the behaviors' activation. The last two sections discuss the conclusions and future directions of this work.

## 2 Gestures in Expert Human Tutoring

Previous work has investigated the kinds of gestures that occur in expert human tutoring [12]. Williams et al. analyzed the gestures of different tutors in ten tutoring sessions. All sessions consisted of naturalistic one-on-one tutoring on diverse subjects such as algebra, chemistry, and biology. From each session, 200 turns were selected for gesture coding, totaling 2,000 turns. The seven gesture categories and are presented in Table 1. Thirty-five action categories are nested within the major gesture categories. Not all are listed due to space constraints.

**Table 1.** Gesture Categories and Actions

<b>Category</b>	<b>Action</b>	<b>Description</b>
Deictic	Point at workspace	Pointing gestures
Iconic	Animate subject matter	Illustrate what is said with concrete semantic meaning
Beat	Count on fingers Point upwards	Emphasize aspects of dialogue; rhythmic in nature
Personal	Cross arms across chest Scratch itch	Do not involve other participant or shared workspace
Gaze	Look at student	Indicates where tutor is looking
Paralinguistic	Gesture for student to take notes Shrug shoulders	Metacommunicative nonverbal speech acts
Action	Write on workspace Thumb through pages	Specific tutor actions on workspace

Williams et al. found that tutors used gestures differentially based on the pedagogical/communicative intent of their utterance, with the exception of beat gestures which occurred throughout. Moreover, their analyses of gesture frequency in

tutoring indicated that three of the four most common gestures involved the workspace (looking at/pointing to/writing on). The high frequency of workspace related gestures in expert human tutoring underscores the relevance of workspace related gestures for pedagogical agents.

Our pedagogical agent, Guru, is designed around the same pedagogical scenario described by Williams et al. In the Guru environment, students interact with a full bodied animated pedagogical agent and a multimedia panel as shown below. The Guru tutor interacts with the workspace and text, images, and diagrams on the workspace appear and disappear at relevant points in the tutoring session.

The utterances generated by the Guru agent are controlled by a dialogue manager which selects dialogue based on the student's current understanding and the current pedagogical goals (e.g. introduce a new topic, provide scaffolding, assess student understanding). However, the nonverbal behavior of Guru is not controlled by the dialogue manager. Instead, nonverbal behaviors are managed by the EBN; the pedagogical goals and agent actions that are handled by the EBN are described below.

### 3 Pedagogical Goals and Actions

Based on the gesture analyses conducted by Williams et al. [12] and careful observations of expert tutoring videos, we have identified a subset of pedagogical goals and agent actions that correspond to the nonverbal behaviors that are prevalent in one-to-one tutoring. The seven pedagogical goals are *Emphasize Concept*, *Ground Concept*, *Ask Question*, *Wait for Answer*, *Provide Positive Feedback*, *Provide Negative Feedback*, and *Provide Neutral Feedback*. These goals, along with information in the dialogue history, drive the Guru agent's nonverbal actions. The agent actions are Sway Back and Forth, Interlock Fingers, Cross Arms, Gesture Left then Right, Hands Out, Head Tilt Nod, Look at Whiteboard, Left Hand Out, Right Hand Out, Shrug Shoulders, Head Tilt, Point to Whiteboard, Head Tilt Left Then Right, Head Shake Yes, Head Shake No, Animate Subject Matter, Smile, and Grimace.

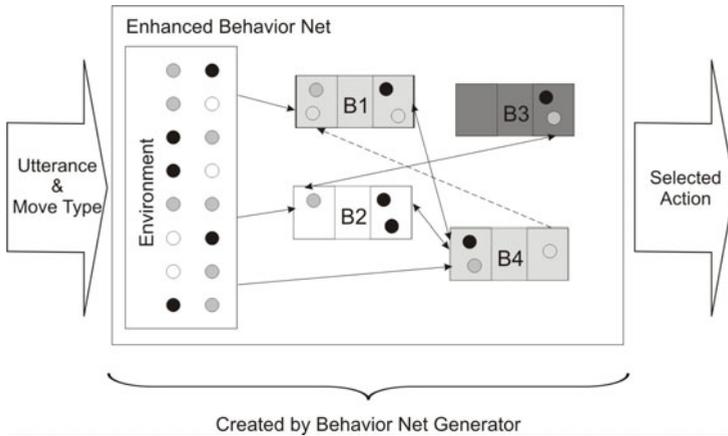
Some agent actions map onto to more than one goal when they are used in different contexts. For example, the agent may cross its arms when waiting for a student to answer or when delivering negative feedback. Although multiple actions can map onto a goal, the effectiveness of each action for a particular goal may differ. For example, the action Gesture Left may map to the goal Emphasize Concept, but Point to Whiteboard has a stronger effect because the action is more precise.

Pedagogical and discourse goals change over the course of an explanation. When a new concept or topic is first introduced, e.g. Cholesterol, the agent should ground the discourse referent by highlighting the image -- also known as *grounding* the referent [13]. However, once the referent has been grounded, it is no longer necessary to keep grounding it. Thus grounding is an example of a pedagogical (and conversational) goal that is sensitive to the discourse history.

### 4 Enhanced Behavior Net

Maes extensively describes the Behavior Net, an action selection mechanism, in her original papers [11][14]. We will refer to this original implementation as MASM

(Maes' Action Selection Mechanism) in the rest of this paper. The basic components of MASM are *behaviors*. A behavior is composed of an *action*, a list of logic literals called *preconditions* and a *result list*. The result list represents the consequences of the action and it is also composed of literals, but in this case, it is divided in two sub lists: the *add list* and the *delete list*. The add list is the list of literals that become true after the action execution and the literals in the delete list become false. Each behavior has a real value attribute called *activation*. The behaviors receive activation from the *environment*, the list of literals that are true in the system. If a literal is in the *environment*, this means that this literal is true, all behaviors that have this same literal in their precondition receive some activation.



**Fig. 1.** The Enhanced Behavior Net. B1 to B4 are behaviors. The shades of gray represent their activations. The circles in the environment represent conditions and goals. The circles in the behaviors represent preconditions (left) and the result list (right).

Behaviors also receive activation from the *goals*, a list of literals whose desired end-state is true. Behaviors that have a goal as part of their add list gain some activation. Finally, the behaviors comprise a network of *successors* and *predecessors*. When the results of one behavior are the preconditions of another, then the former is the predecessor and the latter is the successor and they are linked. Behaviors spread activation among their successors and predecessors. There are also *conflictor links* that are outside the scope of this paper.

A behavior becomes *executable* when all its preconditions are true and its activation is over a threshold  $T$ . All executable behaviors compete and the action of the winner is executed. If no behavior fulfills both conditions to be executable,  $T$  is diminished and another cycle is performed until one behavior is selected.

The original MASM was studied in detail by Tyrell [6] and he reported on three strengths and limitations that are particularly relevant to the current discussion of modeling discourse and pedagogical rules. First, MASM uses Boolean conditions: either a feature is present in the environment or not. However, several of the discourse/pedagogical goals require graded values of features to represent recency and effectiveness. Second, MASM does not include negated preconditions. This

makes it difficult to represent behaviors that apply to things that have not happened yet, like grounding a referent. Third, when conflicting goals compete, MASM does not handle the passing of activation correctly in all situations. Several variations and improvements that correct some of these problems and extend its applicability to more complex domains have been proposed. See for example: [15][16]. The following paragraphs address these issues in detail.

In regard to the Boolean limitation of MASM, a degree of fuzziness is needed to represent if a specific action was executed recently. The literals of MASM were replaced by *conditions* that have real-valued activation and *importance*. A condition is considered true if it exceeds a defined threshold. The activation of conditions decays over time. *Importance* is used to distinguish the importance of each precondition. Preconditions with more importance contribute more to the activation of the behavior. With regard to the MASM limitation on negation, MASM presumes a closed world assumption, meaning that literals not present in the environment are considered false. This assumption has some drawbacks. Some behaviors may need a specific condition to be false in order to be executable. For example, a word must be not *emphasized\_recently* to be emphasized. Our solution is to allow preconditions to be *negated*. Conditions that are not present in the precondition list of a behavior are considered irrelevant for that particular behavior. The environment has only *positive* conditions. When activation is passed from the environment, behaviors with negated preconditions receive activation as a function of the complement of the activation of positive condition in the environment, i.e. one minus the activation of the condition.

It is important to notice that the environment now contains all the conditions of the system including goals. The activation of each of them indicates the degree of certainty of this condition. Each goal has an attribute called *desirability*. Behaviors can receive activation from the goals, if they can produce one of them. The amount of activation that a goal contributes with behaviors is relative to the difference between its desirability and its activation. Then, a goal with full activation is already fulfilled and it is not necessary to select behaviors that produce this goal.

An important module of the EBN is the Behavior Network Generator (BNG). This module creates the Behavior Net required for the material the tutor is going to teach using only basic components of the dialogue, key concepts, multimedia assets, and a predefined set of actions and goals. The BNG populates the graph of the behavior net by defining the environment, the network of behaviors, and their connectivity. A complete description of the BNG is beyond the scope of this paper.

The Guru Tutor uses the EBN to dynamically synchronize verbal, non-verbal, and multimedia presentations. The EBN takes a raw tutor utterance and a pedagogical move type from the dialogue manager, e.g. "QUESTION", "DIRECT\_INSTR", "FEEDBACK\_OK", etc. The EBN first adjusts the activation of the environment's conditions based on the received utterance and the move type. The EBN inspects the text of the utterance for words that relate to key concepts and multimedia assets using string matching. Conditions in the EBN environment are updated based on these matches, e.g. if "mitosis" is present, then the multimedia condition for "mitosis" will receive activation, as will the key concept node for "mitosis". Then the behaviors' activation is recalculated. If any behavior is executable and over the selection threshold then this behavior is selected. Otherwise, the threshold is decreased and more activation is spread among behaviors. This process is repeated until one

behavior is selected. Finally, the activation of all behaviors and conditions are decayed. There are defined some background behaviors with a constant, never decaying, low activation. This guarantees that one action is eventually selected.

The output of the EBN is an action that specifies some nonverbal behavior or change in the multimedia panel, synchronized with the agent's speech using SAPI 5 bookmarks that allows synchronization of the agent with TTS events, e.g. word boundaries.

## 5 Conclusions

In this paper we present a novel approach to the nonverbal action selection problem for an intelligent tutoring agent, Guru. We introduce a variation of the original Behavior Net that has several improvements: negated preconditions for behaviors, continuous values of activation instead of Boolean conditions, desirability for goals, and a new set of activation passing equations for better performance. For example, the activation of conditions tracks the time since the last execution of one action. The desirability for goals is used to model their importance and the negated preconditions are used for example to allow the execution of an action only if has not been executed recently. Summing up, this Enhanced Behavior Net allows a better modeling of the action selection problem.

One of the main advantages of this architecture is the use of goals and conditions to model the requirements instead of static rules. The net of behaviors, inherent in the structure of the architecture, can handle dynamic changes and resolve conflicts that otherwise require complex rules to resolve. The activation and decaying mechanisms of the goals and conditions allow the previous history of selected actions and other relevant conditions in the future action selections to be taken into account. A first implementation of the EBN was integrated into the Guru ITS, and the preliminary testing experiments were performed using material from different lectures. For more comprehensive testing, we will expand our current set of animations to better reflect the varied output of the EBN.

We recognize that our current implementation has limitations, and we are entertaining additional ways to improve in the system. First, the actual implementation of the BNG is fixed for our specific domain: the Guru ITS. A more generic approach is to implement it with a mechanism where the user or developer could specify the structure of behaviors, conditions, goals and other elements of the system. This will extend the use of this architecture beyond the scope of Guru. Also, a standard output format, like BML, could facilitate this same goal. Another possible improvement is to perform a more sophisticated preprocessing with the input utterance.

**Acknowledgments.** The research reported here was supported by the Institute of Education Sciences, U.S. Department of Education, through Grant R305A080594 to the University of Memphis. The opinions expressed are those of the authors and do not represent views of the Institute or the U.S. Department of Education.

## References

1. Graesser, A.C., McNamara, D.S., VanLehn, K.: Scaffolding deep comprehension strategies through Point&Query, AutoTutor, and iSTART. *Educational Psychologist* 40(4), 225–234 (2005)
2. Lester, J.C., Towns, S.G., Fitzgerald, P.J.: Achieving affective impact: Visual emotive communication in lifelike pedagogical agents. *International Journal of Artificial Intelligence in Education* 10(3-4), 278–291 (1999)
3. Rickel, J., Johnson, W.L.: Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence* 13(4), 343–382 (1999)
4. Atkinson, R.K.: Optimizing learning from examples using animated pedagogical agents. *Journal of Educational Psychology* 94(2), 416 (2002)
5. Hershey, K., Mishra, P., Altermatt, E.: All or nothing: Levels of sociability of a pedagogical software agent and its impact on student perceptions and learning. *Journal Educational Multimedia and Hypermedia* 14(2), 113–127 (2005)
6. Tyrell, T.: *Computational Mechanisms for Action Selection*. PhD Thesis, University of Edinburgh, UK (1993)
7. Olney, A.M., Graesser, A.C., Person, N.K.: Tutorial Dialog in Natural Language. In: Nkambou, R., Bourdeau, J., Mizoguchi, R. (eds.) *Advances in Intelligent Tutoring Systems*. Studies in Computational Intelligence, vol. 308, pp. 181–206. Springer, Heidelberg (2010)
8. Bergmann, K., Kopp, S.: GNetIc – using bayesian decision networks for iconic gesture generation. In: Ruttkey, Z., Kipp, M., Nijholt, A., Vilhjálmsson, H.H. (eds.) *IVA 2009*. LNCS, vol. 5773, pp. 76–89. Springer, Heidelberg (2009)
9. Lee, J., Marsella, S.C.: Nonverbal behavior generator for embodied conversational agents. In: Gratch, J., Young, M., Aylett, R.S., Ballin, D., Olivier, P. (eds.) *IVA 2006*. LNCS (LNAI), vol. 4133, pp. 243–255. Springer, Heidelberg (2006)
10. Neff, M., Kipp, M., Albrecht, I., Seidel, H.P.: Gesture modeling and animation based on a probabilistic re-creation of speaker style. *ACM Transactions on Graphics (TOG)* 27(1), 5 (2008)
11. Maes, P.: How to do the right thing. *Connection Science* 1, 291–323 (1989)
12. Williams, B., Williams, C., Volgas, N., Yuan, B., Person, N.: Examining the Role of Gestures in Expert Tutoring. *Intelligent Tutoring Systems*, 235–244 (2010)
13. Clark, H.H., Brennan, S.E.: Grounding in communication. In: Resnick, L.B., Levine, J.M., Teasley, S.D. (eds.) *Perspectives on Socially Shared Cognition*, pp. 127–149. American Psychological Association, Washington, DC (1991)
14. Maes, P.: Modeling Adaptive Autonomous Agents. *Artificial Life* 1, 135–162 (1993)
15. Decugis, V., Ferber, J.: Action selection in an autonomous agent with a hierarchical distributed reactive planning architecture. Paper Presented at the Second International Conference on Autonomous Agents, Minneapolis, MN USA (1998)
16. Negatu, A., Franklin, S.: An action selection mechanism for ‘conscious’ software agents. *Cognitive Science Quarterly* 2, 363–386 (2002); special issue on Desires, goals, intentions, and values: Computational architectures. Guest editors Maria Miceli and Cristiano Castelfranchi