An Instructor's Assistant for Team-Training in Dynamic Multi-agent Virtual Worlds

Stacy C. Marsella & W. Lewis Johnson

Information Sciences Institute / Univ. of Southern California
4676 Admiralty Way, Marina del Rey, CA 90292
{marsella, johnson}@isi.edu

Abstract. The training of teams in highly dynamic, multi-agent virtual worlds places a heavy demand on an instructor. We address the instructor's problem with the PuppetMaster. The PuppetMaster manages a network of monitors that report on the activities in the simulation in order to provide the instructor with an interpretation and situation-specific analysis of student behavior. The approach used to model student teams is to structure the state space into an abstract situation-based model of behavior that supports interpretation in the face of missing information about agent's actions and goals.

1 Introduction

Teams of people operating in highly dynamic, multi-agent environments must learn to deal with rapid and unpredictable turns of events. Simulation-based training environments inhabited by synthetic agents can be effective in providing realistic but safe settings in which to develop skills these environments require (e.g., [14]). To faithfully capture the unpredictable multi-agent setting, virtual world simulations can be inhabited by multitudes of synthetic agents which ideally exhibit the same kind of complex behaviors that human participants would exhibit. For example, Distributed Interactive Simulation (DIS) training sessions involve student teams interacting with potentially thousands of synthetic and human agents within a very dynamic battlefield simulation [14].

However such virtual environments present a problem for the instructor who must evaluate and control rapidly evolving training sessions. Information from any one student's perception of events may be unavailable and agents (human or synthetic) may not be able to explain their own motives. Furthermore, abstracting from the behavior of individuals to the behavior and goals of teams requires additional effort. Conversely, there may be too much low level information for the instructor to absorb and interpret. And as the number of interacting agents grows, the difficulty increases. Accordingly, it may be difficult to determine what teams are doing and why they are doing it.

We address the instructor's problem with a synthetic assistant, the PuppetMaster, which provides a high-level interpretation and assessment of the teams. As the teacher's automated assistant, the PuppetMaster dynamically assigns probes that monitor student teams and synthetic agents in the simulation. The situation in the virtual world is then assessed from the perspective of high level training objectives. The assessment includes events, trends and aggregate reporting over multiple entities (e.g., for revealing teamwork). The resulting evaluation is used to compose presentations that reduce the instructor's effort during the exercise and assists the instructor's review with students after the exercise.

The central concern of this paper is how the PuppetMaster forms its assessment of student teams. The dynamics and pedagogical goals for the domains we have been considering present a considerable challenge for student assessment. In dynamic environments, such as DIS battlefields, there are multiple tasks and multiple agents/teams performing those tasks. Moreover, there is no overall plan of action either at the individual or team level that can be guaranteed to achieve the goal; events can enfold in such a way that obviates any plan. Student teams must learn to operate in both goal-directed and reactive fashions. To assess how well a team is doing, it is critical to appropriately model the students' loosely scripted interactions with the world, in the face of changing plans, partial information and irrelevant information. And instructional support needs to foster development of initiative and the ability to react or replan appropriately.

This can present a problem for modeling students and/or their task performance [15, 2, 10, 4, 1] as well as more generally plan recognition[11, 5]. For instance, student monitoring is often tied to detailed modeling of the task and matching of actions or events. Such approaches tend to have most ready application when the skill to be acquired can be fully modeled in terms of what specific actions need to be taken and the order in which they should be taken. Approaches to modeling in more dynamic settings (e.g., [4]) do weaken somewhat how actions and goals are fitted into a plan. But the pedagogical intent to do a detailed fitting or causal analysis remains along with the presumption of access to all the information and action modeling that supports that analysis. Related work in agent tracking (e.g., [13]) infers plans that are a mix of plan-based and reactive procedures but also presumes detailed modeling of an agent's actions. Finally, there also needs to be a way to aggregate the modeling of an individual agent's behavior into the modeling of team behavior.

To model student teams, we have adapted an approach from reactive planning research, the Situation Space [12, 7]. A situation space structures states of the world into classes of problem solving histories whereby an agent's recognition of its *current situation* guides its goal-directed behavior. The PuppetMaster models both reactive and goal driven behaviors within the framework of a situation space model of the entire student team. The *current situation* provides a top-down focus for monitoring, inferring missing information and assessing behavioral trends.

2 The Need

We explored the design of our instructor's assistant in virtual world simulation for training military tank platoons. The entities in this simulation environment include teams of trainees in tank simulators as well as synthetic forces generated by software such as the ModSAF (Modular Semi-Autonomous Forces) program [3]. There are four tank simulators in a typical platoon exercise, each manned by four students, and approximately fifty synthetic forces. Army instructors manage the exercise, performing the roles of superior officer (e.g., issuing orders) and adversary (e.g, dynamically creating and tasking opposing forces). After the exercise, instructors provide feedback to students.

An exercise for a tank platoon might involve traveling in *wedge* formation to a location in order to occupy a position which blocks the opposing force. While in a wedge formation, there are guidelines as to how the individual tanks follow each other, keep each other in sight, etc. As the platoon travels to the position along the virtual terrain, it may encounter friendly or opposing forces to which it has to respond appropriately. These forces can be either synthetic or human agents and encounters with them can rapidly unfold in unpredictable ways. Likewise, the platoon may encounter terrain features that can slow down and/or damage its vehicles. Thus, the multi-agent dynamic qualities of the simulation can rapidly impact what needs to be done to satisfy the broadly scripted exercise, as well as whether it can be satisfied.

Based on our observations of these exercises, instructors do not micromanage or even micro-evaluate the students. For instance, they don't analyze student behavior on an action-by-action basis unless the effect of a student's actions is entirely anomalous in the current situation. Additionally, there are also individual differences in how instructors evaluate specific skills in the training context.

These characteristics are consistent with the nature of the domain and the skills that need to be acquired. Because of the dynamic, multi-agent domain, there is no guaranteed plan for achieving the goals set out for the students. Without such a plan, or other agreed upon objective basis, micro-evaluation at the level of each and every individual action is problematic. Furthermore, understanding the rationale for every action may require modeling the domain from the perspective of every student, e.g., down to the level of every rock which must be circumnavigated to avoid tread damage. Also, because of the dynamic domain, it is necessary to foster the development of initiative in the teams consistent with situation-appropriate goals and behaviors.

Although instructors do an excellent job of managing training exercises, the tools that they have at their disposal are severely lacking. Instructors are often forced to write notes to themselves on paper as the exercise proceeds, and must rely on subjective impressions. Also, as an exercise unfolds over time, relevant information concerning transient events or trends may be missed.

Instructors need to know how well the platoon is maintaining their formation over time, which requires analysis of the formation over time, when it is appropriate. A formation that is appropriate when traveling may be inappropriate when engaging the enemy. If the formation is poor, an analysis from the team's perspective is necessary. Is it a problem in maintaining visual contact, difficulty navigating the terrain, or problems with the vehicles? If the enemy is nearby, is the team in a position to spot them or be spotted by them?

The information necessary to address these concerns is present in the simulation, if one looks in the right place at the right time. The goal of the PuppetMaster is to extract the information from the simulation automatically and to analyze it both from the team's and instructor's perspectives, allowing the instructor to focus on where to provide instructional feedback.

3 Situation Spaces in the PuppetMaster

The PuppetMaster works within the Probes system which includes virtual world monitors and a display manager. Based on a description of the training objectives of each exercise, the PuppetMaster dynamically assigns the monitors (typically embedded in instrumented ModSAF agents) depending on what information is required to recognize and analyze the current phase of the exercise. These monitors collect the requested data and report back, on a regular basis, or when interesting events occur, or in response to queries from the PuppetMaster. The PuppetMaster uses the data it receives to interpret and assess the training exercise. Output to the instructor is controlled by a display manager. PuppetMaster's understanding of the training exercise is organized around a Situation Space. It uses the Situation Space to control monitoring, form an assessment of the team and selectively report the assessment.

As developed in reactive planning research, Situation Spaces structure states of the world into situations and links between situations. The behavior of the agent is organized around its current situation which determines the (sub)goal(s) it should try to achieve (or maintain) as well as how to monitor the world. In turn, monitoring determines whether traversal to a new situation has occurred. Traversal could be caused by successful achievement of a (sub)goal or unexpected turn of events which could be advantageous or dis-advantageous. Paths through the situation space represent alternative abstract partial plans which incorporate both goal-directed behavior as well as goal-directed responses to unforeseen events.

To make these ideas more concrete, consider the situation space depicted in Fig. 1 for the simplified example problem presented earlier. Recall the problem is to travel in a wedge formation to a blocking position, responding to encounters with the opposing force en route. The current situation at the start of the exercise is assumed to be "Traveling" (Wedge Formation). When a platoon is in this situation, their goals include traveling along some route towards the blocking position, and maintenance goals of maintaining a wedge formation and scanning for the enemy. As they pursue these goals, expected or un-expected transitions can occur that must be monitored. For instance, the students must monitor for reaching their objective, based on the transition to the Goal situation. Also, they must monitor for contact with other forces. This could cause a transition to an "Action on Contact." The goals in this new situation would include assessing the threat, targeting opposing forces and insuring one is not an easy target. When in the Action on Contact situation, transitions to "Disabled" or back to "Traveling" must be monitored.

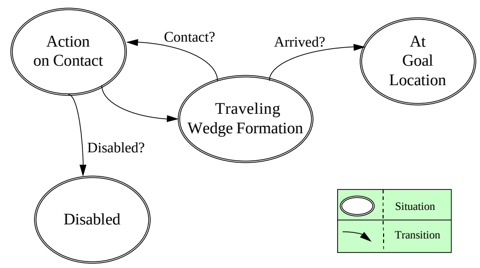

Fig. 1. Example Situation Space.

Several characteristics of Situation Spaces can be gleaned from this example. There can be an indefinite number of paths through the space. This allows it to compactly express the dynamic characteristics in these training environments whereby students can undergo repeated and unexpected encounters. The necessary reactivity is modeled within an overall goal-directed declarative plan. And the set of possible plans is modeled at a high level, with each situation modeling many possible problem solving histories.

In addition to their value in planning, situation spaces are a good declarative model around which to organize analysis of team behavior in a dynamic world. To serve that end, we transformed them from an aid for planning to an aid for analysis of team behavior.

3.1 Situation Spaces from an Instructor's Perspective

We made two moves to use Situation Spaces to analyze team behavior. First, we transformed the goals indicated by a situation from goals-to-be-achieved to analyses-to-be-performed. For instance, the goal associated with traveling in a wedge, along a route to the blocking position, becomes an analysis goal to determine how well the team is traveling. Similarly, the Action On Contact goals of "targeting" and "avoid being an easy target" become how well they are targeting and avoid being targeted.

The other move we made was to allow for multiple perspectives. This move is necessary because the student team may have a different perspective from the PuppetMaster as to what the current situation is, and such discrepancies provide key pedagogical assessments. For instance, the students may transition to an Action On Contact situation and fire at what they think is the opposition, but actually are distant rocks. Meanwhile, the PuppetMaster, due to its larger perspective, will know that the transition to action on contact was a mistake. Conversely, a viable enemy threat may exist which the students have not seen.

A consequence of maintaining multiple perspectives is that, for those transitions that could result in bifurcation of perspective, there needs to be a transition arc for each perspective. For the transition from the Traveling to the Action On Contact situations there need to be 2 arcs, one for pedagogical perspective and one for team perspective. (Figure 1 shows only one of the arcs.)

Another consequence of multiple perspectives is the need to determine which situation is the appropriate basis for monitoring and reporting analyses. We have found it sufficient to report analyses based on the team perspective, for several reasons. Reporting an analysis of the team from a perspective different from that guiding their behavior tends to generate marginally useful information. In contrast, analyzing team behavior from their perspective generates the information needed to infer the point at which they finally transition into the "correct" situation. For instance, when the team is traveling in a wedge, there are certain behaviors they are supposed to exhibit such as constantly scanning their turrets and following each other at a certain distance. These behaviors are not typically appropriate during active engagements so their cessation is a key indicator for inferring transition of the team perspective. On the other hand, the pedagogical perspective is key in analyzing failures of the team in assessing their situation.

Our experience is that the bifurcation into two perspectives persists only for very short periods of time. There are several reasons for this. The situation space is at a very high level of abstraction so a bifurcation tends to indicate a dramatic misread of the environment. In addition, given the pressure that the domain exerts on the team's behavior, such a misread is likely not to persist for long. For example, incoming rounds are a solid indicator to the team that those "rocks" off in the distance are actually the opposition firing at them.

Allowing for these two perspectives has proved sufficient to date in characterizing the state of a platoon exercise. However, we could go further. Recall there are four tanks in a platoon. One might break the team perspective into individual tank perspectives, thus having PuppetMaster track what it considers to be each tank's perspective on the current situation. Now discrepancies between tank perspectives reflect factors such as breakdown in coordination.

This move towards multiple perspectives assumes that the perspectives share the same situation space. However, if there were multiple platoons being assessed, performing different missions, then different situation spaces would be more appropriate. Conversely, we may model a single agents overall behavior as being composed of distinct situations in multiple situation spaces. These points beg the question of composition over situations which would need further study.

In Probes, a situation space is defined as a set of situations, with each situation being a 3-tuple consisting of:

- List of predicate functions for transitions of situation perspectives
- Set of evaluation functions for forming an assessment
- Initialization function which invokes monitors pertinent to the situation.

The Army's instructional material, reinforced by our observations of actual training sessions, provided the high-level structure of the situation space (types of situations, transitions and evaluations). From there, the actual transition, evaluation and initialization functions were defined. Across all the exercises for platoon training, the instructional material was laid out in a modular fashion that shared exercise goals, situations and assessment criteria. This suggests that the application of situation spaces across the full suite of platoon exercises could be semi-automated with considerable reuse of transition and evaluation functions.

3.2 Situation-based vs. Event-based Modeling

Structuring the state space into a situation space achieves several ends. Monitoring can be organized top-down around the situation, which in turn usefully constrains interpretation and assessment. Behavioral trends can be monitored and assessed according to their appropriateness within a situation. Changes in behavioral trends within a situation can be used to infer missing information, such as whether the platoon has spotted the enemy and is going into action on contact. Unnecessary details about the state of the world are not monitored. And the transitions between situations provide a principled basis of coupling analysis of planned and reactive behaviors.

A key feature of the approach is that the high level analysis appropriate to the situation is based mainly on partial state descriptions and trends in those descriptions. This is quite different from work on student modeling and plan recognition which use a complete action model as the basis of checking/inferring a plan with an agent's actions. For instance, recognizing that a platoon is trying to travel in a wedge formation would be at best difficult if the recognition was based on the low level actions that the 16 crew members in the 4 tanks were executing. And evaluating this team behavior is best done at the level of the abstract, partial state descriptions, especially trends in those descriptions, and not at the level of individual discrete actions these various team members are performing. The relation between the levels is not strongly fixed. Moreover, recognition and evaluation is best done at the level at which this joint behavior has consequences in this dynamic environment. Still, causal analysis at the level of individual actions would be of diagnostic utility but only in the context of the higher level analysis and not as a way of deriving the higher level analysis. Some form of constrained plan recognition, for example, may be useful for determining why a wedge formation is falling apart.

4 Example

Let's now consider an example run of the Probes system. In the following example, the team is training on the exercise discussed earlier, traveling in a wedge to a blocking location. Blue is the platoon being assessed whereas Red is the opposition. Here, both sides are comprised of synthetic agents being run by ModSAF, but in an actual exercise Blue would be human crews in simulators.

Probes provides instructors with a coordinated set of presentations. Figure 2 shows two presentations: the situation space for the exercise on the left, with the current situation (team perspective) highlighted. On the right is a log of high-level events and assessments. The PuppetMaster automatically determines which situation is currently in force, by monitoring activities in the simulation. Probes also uses synthesized speech to announce situation transitions.

The exercise analysis log records events and analyses that are likely to be relevant to the current situation. Note that during "Situation Travel (Wedge)" (traveling in wedge), the PuppetMaster reports that vehicles in the platoon are

Fig. 2. Probes' Situation Presentation.

in a poor wedge formation. In contrast, when Situation Act occurs (i.e., Blue is in an Action on Contact) PuppetMaster starts reporting that the platoon is still in a wedge when it should probably be going "on line", in effect modifying their formation in a fashion consistent with the exigencies of an active engagement. At this point it stops assessment of the wedge alignment, since the wedge formation is no longer appropriate for the current situation.

These analyses rely on trends in partial state descriptions, in particular, persistence in the relations between tanks over time. The analyses are not trying to evaluate (or infer) travel in a wedge by reasoning about the actions which the four tanks are executing (or the 16 crew members).

The log also reports when Red spots Blue. PuppetMaster's monitoring can access state information internal to the synthetic agents which reveals tactics and situation assessments. When Red spots Blue, their internal assessment recognizes a threat and, based on that assessment, the PuppetMaster's pedagogical perspective transitions to Action on Contact ("ActualSit Act"). The pedagogical perspective also would have transitioned if Blue's team perspective had transitioned and the transition was valid (e.g., the opposing force did not turn out to be rocks). In the case of the Blue forces (especially when they are human crews), PuppetMaster infers their intent based on the current situation, the objectives of the exercise and trends in the partial state descriptions that are being monitored. For instance, if Red can (potentially) be spotted and is in range, plus Blue has turrets aiming at Red, breaks out of wedge formation, or flattens the wedge, then a situation transition for Blue's perspective can be inferred.

As the situation changes, Probes displays statistics about the unit's performance as appropriate to the current situation. Figure 3 depicts two displays specific to the poor wedge formation. On the left is an "Analysis Details" window, which the instructor accesses by clicking on the poor wedge notification in the event log. This causes Probes to do some shallow reasoning and reveal relevant factors such as local terrain conditions, damage to the vehicles, etc. On the right is a wedge evaluation gauge which pops up automatically when the platoon is in a wedge and depicts the distances between tanks and the overall depths of the wedge. Ideally, Probes analysis can be coupled to a 3D (stealth) view of the virtual world and we have experimented with ways of annotating the 3D display with relevant analysis data.

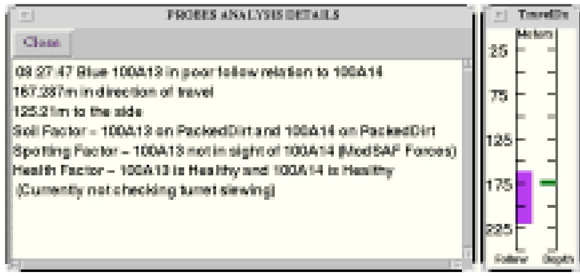

Fig. 3. Probes' presentation of specific details concerning flawed wedge formation.

5 Status

The Probes system has been fully implemented, with all the capabilities illustrated in the "Example" section. We are currently pursuing collaborations that will allow us to evaluate Probes in various training contexts.

Our research is also focusing on techniques for automating the construction of situation spaces, including the use of a scenario generation as a front-end to the situation space construction [9] and decision tree induction for deriving evaluation functions [8]. We are studying the implications of applying situation spaces to more complex teams of students where individual students would be best modeled as distinct situation spaces. In particular, we are concerned with how to formalize the relation between these situation spaces. Work in formal models of teamwork (e.g., [6]) may potentially be useful here.

We expect the design will have applicability to domains that share the unpredictable nature of training of tank platoons. Currently, the applicability of the method is being evaluated in other domains.

6 Concluding Remarks

Simulation-based training of teamwork skills for dynamic, unpredictable multi-agent settings present special challenges from a pedagogical standpoint. It is our view that the instructor can benefit immensely from an assistant that is an intelligent interpreter of events in the simulation. Further, only an automated assistant can play this role: human agents have great difficulty even assimilating the information coming from all the agents' viewpoints and as a consequence have difficulty assessing that information. To be useful, such an assistant also needs to be able to interpret events from a viewpoint of instructional objectives. Finally, it must be able to handle missing information as well as different kinds of agent architectures.

We have developed the PuppetMaster that can assess a training session from the standpoint of instructional objectives. This assistant uses a situation-based approach to modeling behavior that is consistent with the information it has available and the analyses it needs to perform.

Additional Details: http://www.isi.edu/~marsella/pp/

Acknowledgements: The authors wish to acknowledge the assistance of Robert Balzer's group at ISI. This work was funded by DARPA Computer Aided Education and Training Initiative. Opinions expressed are those of the authors and not of the funding agencies or the US Government.

References

1. Alexe, C. & Gecsei, J. A Learning Environment for the Surgical Intensive Care Unit. In Frasson, C., Gauthier, G. & Lesgold, A. (Eds), *Intelligent Tutoring Systems, ITS '96 Proceedings.* Berlin:Springer, 1996. 439-447.
2. Anderson, J.R., Boyle, C.F., Corbett, A.T. & Lewis, M.W. R. Cognitive Modelling and Intelligent Tutoring. In *Artificial Intelligence*, 42(1), 1990. Pp. 7–49.
3. Calder, R.B., Smith, J.E, Courtemanche, A.J., Mar, J.M.F. & Ceranowicz, A.Z. ModSAF Behavior Simulation and Control. In *Proc of the 3rd Conference on Computer-Gen. Forces and Behavioral Representation.* Orlando, Fla: Univ. of Central Fla. Inst. for Simulation and Training, 347-356.
4. Hill, R. & Johnson, W.L. Situated plan attribution for intelligent tutoring. In *Proc. of the Nat. Conf. on AI.* Menlo Park, CA: AAAI Press, 1994.
5. Kautz, H. & Allen, J.F. Generalized Plan Recognition. In *Proceedings of the National Conference on Artificial Intelligence.* Menlo Park, CA: AAAI Press, 1986.
6. Levesque, H.J,, Cohen, P.R. & Nunes, J.H.T. On Acting Together. In *Proc. of the Eighth Natl. Conf. on AI.* Los Altos: CA, Morgan Kaufman, 1990. Pp 94–99.
7. Marsella, S. C. & Schmidt, C.F. Reactive planning using a "Situation Space". In Hendler, J. (Ed.) *Planning in Uncertain, Unpredictable, or Changing Environments (Working Notes 1990 AAAI Spring Symp.).* Available as SRC TR 90-45 of the Systems Research Ctr, Univ of Maryland, College Park, MD. 1990.
8. Murthy, S. K., Kasif, S. & Salzberg, S. A System for Induction of Oblique Decision Trees. In *Journal of AI Research*, 2, 1994, Pp. 1-32.
9. Pautler, D., Woods, S. & Quilici, A. Exploiting Domain-Specific Knowledge To Refine Simulation Specifications. Proc. of the Twelfth IEEE International Conf. on Automated Software Engineering, Nevada, November 1997. Pp 117-124.
10. Rickel, J. & Johnson, J.L. Integrating Pedagogical Capabilities in a Virtual Environment Agent. In Johnson, W.L. (Ed.) *Proc. of The First International Conference on Autonomous Agents.* New York: The ACM Press, 1997. Pp 30–38.
11. Schmidt, C.F., Sridharan, N.S. & Goodson, J.L. The plan recognition problem: an intersect. between psychology and artificial intelligence. AI, 11(1,2), 1978. 45–83.
12. Schmidt, C.F., Goodson, J.L., Marsella, S.C., Bresina, J.L., Reactive Planning Using a Situation Space. Proc of 1989 AI Systems in Govern. Conf, IEEE, 1989.
13. Tambe, M. & Rosenbloom, P. Architectures for Agents that Track Other Agents in Multi-Agent Worlds. In Wooldridge, M., Muller, J.P. & Tambe, M.(Eds) *Intell. Agents II, Agent Theories, Arch., and Lang.* Berlin:Springer-Verlag, 1996. 156–170.
14. Tambe, M., Johnson, W.L., Jones, R., Koss, F., Laird, J.E., Rosenbloom, P.S. & Schwamb, K. Intelligent Agents for interactive simulation environments. *AI Magazine*, 16(1), 1995.
15. VanLehn, K. Student Modelling. In Polson, M.C. & Richardson, J.J. (Eds.), *Intelligent Tutoring Systems.* Hillsdale, NJ: Lawrence Erlbaum, 1988. Pp. 55–78.