

Scalable POMDPs for Diagnosis and Planning in Intelligent Tutoring Systems

Jeremiah T. Folsom-Kovarik¹, Gita Sukthankar², Sae Schatz¹, and Denise Nicholson¹

¹Institute for Simulation and Training
University of Central Florida
3100 Technology Parkway, Orlando, FL 32826
{jfolsomk, sschatz, dnichols}@ist.ucf.edu

²School of Electrical Engineering and Computer Science
University of Central Florida
4000 Central Florida Blvd, Orlando, FL 32816
gitars@cecs.ucf.edu

Abstract

A promising application area for proactive assistant agents is automated tutoring and training. Intelligent tutoring systems (ITs) assist tutors and tutees by automating diagnosis and adaptive tutoring. These tasks are well modeled by a partially observable Markov decision process (POMDP) since it accounts for the uncertainty inherent in diagnosis. However, an important aspect of making POMDP solvers feasible for real-world problems is selecting appropriate representations for states, actions, and observations. This paper studies two scalable POMDP state and observation representations. *State queues* allow POMDPs to temporarily ignore less-relevant states. *Observation chains* represent information in independent dimensions using sequences of observations to reduce the size of the observation set. Preliminary experiments with simulated tutees suggest the experimental representations perform as well as lossless POMDPs, and can model much larger problems.

Introduction

Intelligent tutoring systems (ITs) are computer programs that teach or train adaptively. ITs can choose material, plan tests, and give hints that are tailored for each tutee. Just as working one-on-one with a personal tutor increases learning compared to classroom study, an ITS can help students learn more than they would with a linear teaching program. ITs have created substantial learning gains (e.g., Koedinger, Anderson, Hadley, & Mark, 1997; VanLehn et al., 2007). ITs proactively assist both tutors and tutees by diagnosing underlying causes for tutees' behavior and adapting feedback or hints to tutor them efficiently.

Diagnosis in ITs requires recognizing tutees' intent from behaviors such as answering questions or interacting with a training simulator. Each behavior observed could

stem from several different knowledge states or mental conditions. Deciding which conditions are present is key to choosing effective feedback or other interventions.

Several sources of uncertainty complicate diagnosis in ITs. Learners' plans are often vague and incoherent (Chi, Feltovich, & Glaser, 1981). Learners who need help can still make lucky guesses, while learners who know material well can make mistakes (Norman, 1981). Understanding is fluid and changes often. Finally, it is difficult to exhaustively check all possible mistakes (VanLehn & Niu, 2001).

Proactive *interventions* such as material or hint selection require planning and timing. Interventions can be more or less helpful depending on cognitive states and traits (Craig, Graesser, Sullins, & Gholson, 2004). In turn, unhelpful interventions can distract or create tutee states that make tutoring inefficient (Robison, McQuiggan, & Lester, 2009).

Like many other user assistance systems, tutoring is well modeled as a problem of sequential decision-making under uncertainty. Solving a partially observable Markov decision process (POMDP) can yield an intelligent tutoring policy. This potential was recognized in early work by Cassandra, who proposed teaching with POMDPs (1998).

Learner Model

A model of the tutee is central to proactive tutoring. Many learner models are possible to represent with a POMDP, and this paper describes one simple model structure M . In general, labeled data can inform model parameters, but all parameters herein are based on empirical results drawn from the literature and subjected to a sensitivity analysis.

A particular POMDP is defined by a tuple $\langle S, A, T, R, \Omega, O \rangle$ comprising a set of hidden states S , a set of actions A , a set of transition probabilities T that specify state changes, a set of rewards R that assign values for reaching states or accomplishing actions, a set of observations Ω , and a set of emission probabilities O that define which observations appear (Kaelbling, Littman, & Cassandra, 1998).

POMDP States

Hidden states in an ITS POMDP with the M structure describe the current mental status of a particular tutee. This status is made up of *knowledge states* and *cognitive states*.

Knowledge states describe tutee knowledge. K is the domain-specific set of concepts to be tutored. Each element of K represents a specific misconception or a fact or competency the tutee has not grasped, together termed *gaps*. The ITS should act to discover and remove each gap.

The set C of cognitive states represent transient or permanent properties of a specific tutee that change action efficacy. Cognitive states can also explain observations.

Cognitive states are not domain-specific, although the domain might determine which states are important to track. In M , C tracks boredom, confusion, frustration, and flow. Previous research supplies estimates for how these affective states might change learning effectiveness (Craig et al., 2004). Other ITSs may include additional factors in the state representation such as workload, interruptibility, personality, or demographics.

POMDP Actions

Many hints or other actions could tutor any particular gap. In M , the ITS's controller decides which gap to address (or none). A pedagogical module, separate from the POMDP, is posited to then choose an intervention that tutors the target gap effectively in the current context. The pedagogical module is a simplification, but not a requirement. Alternatively, the POMDP could also assume fine control over hints and actions. However, efficacy estimating functions for each possible intervention would then be needed.

Whether the POMDP controls interventions directly or through a pedagogical module, each POMDP action i has a chance to correct each gap j with base probability a_{ij} . The interventions available to the ITS determine the values of a , and they must be learned or set by a subject-matter expert. For simplicity, no actions in M introduce new gaps.

The tutee's cognitive state may further increase or decrease the probability of correcting gaps according to a modifying function $f_c: [0,1] \rightarrow [0,1]$. In M , values for f_c approximate trends empirically observed during tutoring (Craig et al., 2004). Transitions between the cognitive states are also based on empirical data (Baker, Rodrigo, & Xolocotzin, 2007; D'Mello, Taylor, & Graesser, 2007). Actions which do not succeed in improving the knowledge state have higher probability to trigger a negative affective state (Robison et al., 2009). Therefore, an ITS might improve overall performance by not intervening when no intervention is likely to help, for example because there is too much uncertainty about the tutee's state.

Action types not included in M 's example, such as knowledge probes or actions targeting cognitive states alone, are also possible to represent with a POMDP.

POMDP Observations

Like actions, the experimental representations in the present paper are also agnostic to ITS observation content.

In M , an external assessment module is posited to preprocess observations for transmission to the POMDP, a design decision motivated by the typically complex and non-probabilistic rules that are used to map tutee behaviors into performance assessments aligned with the state space.

The assessment module preprocesses ITS observations of tutee behavior into POMDP observations of assessment lists. POMDP observations in M are indicators that certain gaps or states are present or absent based on tutee behavior at that moment. The POMDP's task is to find patterns in those assessments and transform them into a coherent diagnosis in its belief state. In M , each POMDP observation contains information indicating the presence of one gap and the absence of one gap in the current knowledge state. A non-informative blank can also appear in the present or the absent dimension. The ITS can incorrectly assess absent gaps as present when a tutee *slips*, and present gaps as absent when the tutee makes a *guess* (Norman, 1981).

Key Problem Characteristics

Often, tutoring problems have four characteristics that can be exploited to compress their POMDP representations.

First, in many cases certain knowledge gaps or misconceptions are difficult to address in the presence of other gaps. For example, when a person learns algebra it is difficult to learn exponentiation before multiplication or multiplication before addition. This relationship property allows instructors in general to assign a partial ordering over the misconceptions or gaps they wish to address. A skilled instructor ensures that learners grasp the fundamentals before introducing topics that require prior knowledge.

To reflect the way the presence of one gap, k , can change the possibility of removing a gap j , a final term $d_{jk}: [0,1]$ is added to M which describes the difficulty of tutoring one concept before another. The probability that an action i will clear gap j when other any gaps k are present, then, becomes $f_c(a_{ij} \prod_{k \neq j} (1-d_{jk}))$.

Second, in tutoring problems each action often affects a small proportion of K . To the extent that an ITS recognizes knowledge gaps as different, customized interventions often target individual gaps. Therefore, $a_{ij}=0$ for relatively many combinations of i and j , yielding a sparse matrix.

Third, the presence or absence of knowledge gaps in the initial knowledge state can be close to independent, that is, with approximately uniform co-occurrence probabilities. This characteristic is probably less common in the set of all tutoring problems than the previous two, but such problems do exist. For example, in cases when a tutor is teaching new material, all knowledge gaps will be likely to exist initially, satisfying the uniform co-occurrence property.

Finally, observations in ITS problems can be constructed to be independent by having them convey information about approximately orthogonal dimensions. One such construction is the definition of an observation in M , which contains information about one gap that is present and one gap that is absent. These two components are orthogonal – it is even possible to observe the presence and the absence of the same gap, though one assessment would be untrue.

Experimental Representations

State Queuing

In tutoring problems, a partial ordering over K exists and can be discovered, for example by interviewing subject-matter experts. By reordering all gap pairs j, k in K to minimize the values in d where $j < k$ and breaking ties arbitrarily, it is further possible to choose a strict total ordering over the knowledge states. This ordering is not necessarily the optimal action sequence because of other considerations such as current cognitive states, initial gap likelihoods, or action efficacies. However, it gives a heuristic for reducing the number of states a POMDP tracks at once.

A *state queue* is an alternative knowledge state representation. Rather than maintain beliefs about the presence or absence of all knowledge gaps at once, a state queue only maintains a belief about the presence or absence of one gap, the one with the highest priority. Gaps with lower priority are assumed to be present with the initial probability until the queue reaches their turn, and POMDP observations of their presence or absence are ignored except insofar as they provide evidence about the priority gap. POMDP actions attempt to move down the queue to lower-priority states until reaching a terminal state with no gaps.

The state space in a POMDP using a state queue is $S = C \times (KU\{done\})$, with *done* a state representing the absence of all gaps. Whereas an enumerated state representation grows exponentially with the number of knowledge states to tutor, a state queue grows only linearly.

State queuing places tight constraints on POMDPs. However, these constraints may lead to approximately equivalent policies and outcomes on ITS problems, with which they align. If ITSs can only tutor one gap at a time and actions only change the belief in one dimension at a time, it may be possible to ignore information about dimensions lower in the precedence order. And the highly informative observations that are possible in the ITS domain may ameliorate the possibility of discarding some.

Observation Chaining

An opportunity to compress observations lies in the fact that ITS observations can contain information about multiple orthogonal dimensions. Such observations can be serialized into multiple observations that each contain non-orthogonal information. A process for accomplishing this is *observation chaining*, which compresses observations when emission probabilities of different dimensions are independent, conditioned on the underlying state.

Under observation chaining, observations decompose into families of components. A family contains a fixed set of components that together are equivalent to the original observation. Decompositions are chosen such that emission probabilities within components are preserved, but emission probabilities between components are independent.

As an example, imagine a subset of Ω that informs a POMDP about the values of two independent binary variables, X and Y . To report these values in one observation,

Ω might contain four elements $\{XY, X'Y, XY', X'Y'\}$. An observation chain would replace these elements with one family of two components, $\{X, X', Y, Y'\}$. Whenever one of the original elements would have been observed, a chain of two equivalent components is observed instead.

An observation chain can describe many simultaneous gap assessments and can have any length, so a token *end* is added to signal the end of a chain. In the simplified model M , $\Omega = (KU\{blank\}) \cup (K'U\{blank\}') \cup \{end\}$, with *blank* an observation that contains no information. Real-world ITS observations could also include evidence about C .

In POMDPs with observation chaining, S contains an unlocked and a locked version of each state. Observing the *end* token moves the system into an unlocked state, while any other observation sets the equivalent locked state. Transitions from unlocked states are unchanged, but locked states always transition back to themselves. Therefore, observation chains make Ω scale better with the number of orthogonal observation features, at the cost of doubling $|S|$.

Experiments

Three experiments were conducted on simulated students to empirically evaluate ITSs based on different POMDP representations. Examining the performance of different representations helped estimate how the experimental POMDPs might impact real ITS learning outcomes.

Policies generated using a POMDP solver were used to tutor simulated students that acted according to the ground truth of a model with structure M . The simulated students responded to ITS actions by clearing gaps, changing cognitive states, and generating observations with the same probabilities the learner model specified (specific values used are in *Experimental Setup*).

The first experiment was designed to explore the effects of problem size on the experimental representation loss during tutoring. Lossless POMDPs were created that used traditional, enumerated state and observation encodings and preserved all information from the ground truth model, perfectly reflecting the simulated students. Experimental POMDPs used either a state queue or observation chains, or both. For problems with $|K|=|\Omega|=1$, the experimental representations do not lose any information. As K or Ω grow, the experimental POMDPs compress the information in their states and/or observations, incurring an information loss and possibly a performance degradation.

Since the lossless representation had a perfect model of ground truth, the initial hypothesis in this experiment was that students working with an ITS based on the experimental representations would finish with scores no worse than 0.25 standard deviations above students working with a lossless representation ITS (Glass' delta). This threshold was chosen to align with the Institute of Education Science's standards for indentifying substantive effects in educational studies with human subjects (U. S. Department of Education, 2008). A difference of less than 0.25 standard deviations, though it may be statistically significant, is not considered a substantive difference.

The second experiment measured the sensitivity of the experimental representations to the tutoring problem parameters a and d . As introduced above, d is a parameter describing the difficulty of tutoring one concept before another. When d is low, the concepts may be tutored in any order. Since a state-queue POMDP is constrained in the order it tries to tutor concepts, it might underperform a lossless POMDP on problems where d is low. Furthermore, the efficacy of different tutoring actions can vary widely in real-world problems. Efficacy of all actions depends on a , including those not subject to ordering effects.

Simulated tutoring problems with a in $\{.1, .3, .5, .7, .9\}$ and d in $\{0, .05, .1, .15, .2, .25, .5, .75, 1\}$ were evaluated. To maximize performance differences, the largest possible problem size was used, $|K|=8$. In problems with more than eight knowledge states, lossless representations were not able to learn useful policies under some parameter values before exhausting all eight gigabytes of available memory.

The experimental representations were hypothesized to perform no more than 0.25 standard deviations worse than the lossless representations. Parameter settings that caused performance to degrade by more than this limit would indicate problem classes for which the experimental representations would be less appropriate in the real world.

The third experiment explored the performance impact of POMDP observations that contain information about more than one gap. Tests on large problems ($|K|\geq 16$) showed the experimental ITSs successfully tutored fewer gaps per turn as size increased. One cause of performance degradation was conjectured to be the information loss associated with state queuing. As state queues grow longer, the probability increases that a given observation’s limited information describes only states that are not the priority state. If the information loss with large K impacts performance, state queues may be unsuitable for representing some real-world tutoring problems that contain up to a hundred separate concepts (e.g., Payne & Squibb, 1990).

To assess the feasibility of improving large problem performance, the third experiment varied the number of gaps that could be diagnosed based on a single observation. Observations were generated, by the same method used in the other experiments, to describe n equal partitions of K . When $n=1$, observations were identical to the previous experiments. With increasing n , each observation contained evidence of more gaps’ presence or absence. Observations with high-dimensional information were possible to encode with observation chaining. This experiment would not conclusively prove that low-information observations degrade state queue performance, but would explore whether performance could be improved.

Experimental Setup

ITS performance in all experiments was evaluated by the number of gaps remaining after t ITS-tutee interactions, including no-op actions. In these experiments, the value $t=|K|$ was chosen to increase differentiation between conditions, giving enough time for a competent ITS to accomplish some tutoring but not to finish in every case.

All gaps had a 50% initial probability of being present. Any combination of gaps was equally likely, except that starting with no gaps was disallowed. Simulated students had a 25% probability of starting in each cognitive state.

The set of actions included one action to tutor each gap, and a no-op action. Values for a and d were $a_{ij}=0.5$ where $i=j$, and 0 otherwise; and $d_{ij}=0.5$ where $i<j$, and 0 otherwise. Nonzero values varied in the second experiment.

The reward structure was different for lossless and state-queue POMDPs. In lossless POMDPs, any transition from a state with at least one gap to a state with j fewer gaps earned a reward of $100j/|K|$. For state-queue POMDPs, any change of the priority gap from a higher priority i to a lower priority $i-j$ earned a reward of $100j/|K|$. No changes in cognitive state earned any reward. The reward discount was 0.90. Goal states were not fully observable.

Policy learning was conducted with the algorithm Successive Approximations of the Reachable Space under Optimal Policies (SARSOP). SARSOP is a point-based learning algorithm that quickly searches large belief spaces by focusing on the points that are more likely to be reached under an approximately optimal policy (Kurniawati, Hsu, & Lee, 2008). SARSOP’s ability to find approximate solutions to POMDPs with many hidden states makes it suitable for solving ITS problems, especially with the large baseline lossless POMDPs. Version 0.91 of the SARSOP package was used in the present experiment (Wei, 2010).

Policy training and evaluation used a standard controller and no lookahead. Informal experiments suggested that lookahead increased training time and memory requirements without significantly improving outcomes.

Training in all conditions consisted of value iteration for 100 trials. Informal tests suggested that learning for more iterations did not produce significant performance gains. Evaluations consisted of 10,000 runs for each condition.

Results

The first experiment did not find substantive performance differences between experimental representations and the lossless baseline. Figure 1 shows absolute performance degradation was small in all problems. Degradation did tend to increase with $|K|$. Observation chaining alone did not cause any degradation except when $|K|=8$. State queues, alone or with observation chains, degraded performance by less than 10% of a standard deviation for all values of $|K|$.

The second experiment tested for performance differences under extreme values of a and d . Observation chaining did not cause a substantive performance degradation under any combination of parameters. Differences did appear for POMDPs with a state queue alone or a state queue combined with observation chaining. With a state queue alone, POMDPs performed substantively worse than the lossless baseline on tutoring problems with $d<0.25$. With a state queue and observation chaining combined, POMDPs performed substantively worse on problems with $d<0.20$. The improvement over queuing alone may be attributable to more efficient policy learning possible with smaller Ω .

Performance Comparisons

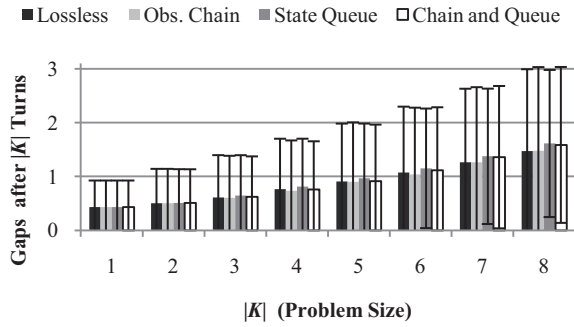


Figure 1: Direct comparison to a lossless POMDP on problems with $|K| \leq 8$, $a = .5$, $d = .5$ shows the experimental representations did not substantively degrade tutoring performance.

Figure 2 shows the performance degradation in the second experiment POMDP using both state queues and observation chains, as a percentage of the lossless performance. Performance changes that were not substantive, but were close to random noise, ranged from 0% to 50% worse than the lossless representation. The first substantive difference came at $d = .15$, with a 53% degradation. The most substantive difference was 179% worse, but represented an absolute difference of only .22 (an average of .34 gaps remaining, compared to .12 with a lossless POMDP).

The third experiment (Figure 3) showed the effect on tutoring performance of observations with more information. For $|K| = 32$, tutoring left an average of 9.4 gaps when each observation had information about one gap, and 3.7 when each observation described 16 gaps. For $|K| = 64$, the number of gaps remaining went from 21.5 to 9.2. Finally, when $|K| = 128$ performance improved from leaving 47.9 gaps to leaving just 22.3 with information about 32 gaps in each observation. However, doubling the available information again did not result in further improvement.

Parameter Sensitivity (Chain and Queue)

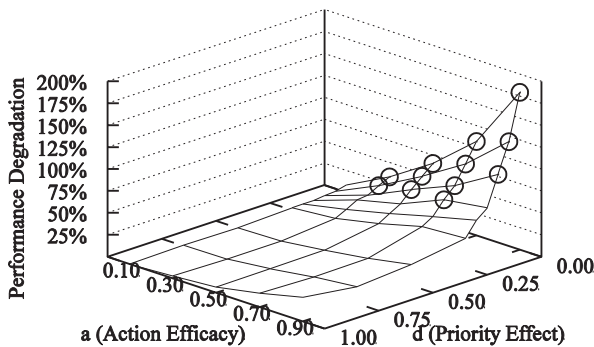


Figure 2: An experimental (chaining and queuing) POMDP can underperform a lossless representation on unsuitable problems. When $|K| = 8$, as here, the absolute difference is small but still substantive at some points (circled) where $d < 0.2$.

Increasing Information (Chain + Queue)




Figure 3: Performance on large problems can be improved up to approximately double, if every observation contains information about multiple features. Observation chaining greatly increases the number of independent dimensions a single observation can completely describe (the number above each column).

Discussion

Together, the simulated experiments in this paper show encouraging results for the practicality of POMDP ITSs.

The first experiment demonstrated that information compression in the experimental representations did not lead to substantively worse performance for problems small enough to compare. The size limit on lossless representations ($|K| \leq 8$) is too restrictive for many ITS applications. The limit was caused by memory requirements of policy iteration. Parameters chosen to shorten training runs could increase maximum problem size by only one or two more. Furthermore, one performance advantage of using enumerated representations stems from their ability to encode complex relationships between states or observations. But the improvement may not be useful if it is impractical to learn or specify each of the thousands of relationships.

The second experiment suggested that although certain types of problems are unsuitable for the experimental representations, problems with certain characteristics are probably suitable. Whenever gap priority had an effect on the probability of clearing gaps $d \geq 0.2$, state queues could discard large amounts of information (as they do) and retain substantively the same final outcomes. Informal discussions with subject-matter experts suggest many real-world tutoring problems have priorities near $d \approx 0.5$.

The third experiment demonstrated one method to address performance degradation in large problems. Adding more assessment information to each observation empirically improved performance. Observing many dimensions at once is practical for a subset of real-world ITS problems. For example, it may be realistic to assess any tutee performance, good or bad, as demonstrating a grasp of all fundamental skills that lead up to that point. In such a case one observation could realistically rule out many gaps.

In the third experiment, problems with $|K| = 128$ and more than 32 gaps described in each observation did not finish policy learning before reaching memory limits. This sug-

gests an approximate upper limit on problem size for the experimental representations. POMDPs that model 128 gaps (1,024 states) could be sufficient to control ITSs that tutor many misconceptions, such as (Payne & Squibb, 1990). Even larger material could be factored to fit within the limit by dividing it into chapters or tutoring sessions.

Conclusions

POMDPs' abilities to learn, plan, and handle uncertainty could help intelligent tutoring systems proactively assist tutors and tutees. Simulation experiments suggested that with the modifications introduced herein, POMDP solvers can generate policies for problems as large as real ITSs confront.

Further modifications might improve ITS performance with the compressed representations. First, a hybrid state representation might only queue a subset of knowledge states, tracking others without queuing if greater precision improves proactive behavior. Second, since state queues discard some state information, their beliefs could be supplemented by external structures such as Bayesian nets or particle filters interposed between the assessment source and the POMDP. Observation chains could transmit beliefs from these structures to the POMDP with high fidelity.

Next steps will include applying a POMDP ITS to a real-world tutoring problem with human tutees.

Acknowledgements

This work is supported in part by the Office of Naval Research Grant N0001408C0186, the Next-generation Expeditionary Warfare Intelligent Training (NEW-IT) program and AFOSR award FA9550-09-1-525. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ONR, AFOSR, or the US Government. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

Baker, R., Rodrigo, M., & Xolocotzin, U. (2007). *The Dynamics of Affective Transitions in Simulation Problem-Solving Environments*. Paper presented at the 2nd Int'l. Conf. on Affective Computing and Intelligent Interaction.

Cassandra, A. R. (1998). A survey of POMDP applications. In *Working Notes of the AAAI Fall Symposium on Planning with Partially Observable Markov Decision Processes* (pp. 17–24). Orlando, FL: AAAI.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5(2), 121–152.

Craig, S. D., Graesser, A. C., Sullins, J., & Gholson, B. (2004). Affect and learning: an exploratory look into the role of affect in learning with AutoTutor. *Journal of Educational Media*, 29, 241–250.

D'Mello, S., Taylor, R. S., & Graesser, A. C. (2007). *Monitoring affective trajectories during complex learning*. Paper presented at the 29th Annual Meeting of the Cognitive Science Society.

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2), 99–134.

Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8(1), 30–47.

Kurniawati, H., Hsu, D., & Lee, W. S. (2008). *SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces*. Paper presented at the Robotics: Science and Systems Conference.

Norman, D. A. (1981). Categorization of action slips. *Psychological Review*, 88(1), 1–15.

Payne, S. J., & Squibb, H. R. (1990). Algebra mal-rules and cognitive accounts of error. *Cognitive Science: A Multidisciplinary Journal*, 14(3), 445–481.

Robison, J., McQuiggan, S., & Lester, J. (2009). *Evaluating the Consequences of Affective Feedback in Intelligent Tutoring Systems*. Paper presented at the Int'l. Conf. on Affective Computing and Intelligent Interaction.

U. S. Department of Education. (2008). What Works Clearinghouse: Procedures and Standards Handbook (Version 2.0). Retrieved Aug. 21, 2009, from <http://ies.ed.gov/ncee/wwc/references/idocviewer/doc.aspx?docid=19&tocid=1>

VanLehn, K., Graesser, A. C., Jackson, G. T., Jordan, P., Olney, A., & Rosé, C. (2007). When Are Tutorial Dialogues More Effective Than Reading? *Cognitive Science*, 31, 3–62.

VanLehn, K., & Niu, Z. (2001). Bayesian student modeling, user interfaces and feedback: A sensitivity analysis. *International Journal of Artificial Intelligence in Education*, 12(2), 154–184.

Wei, L. Z. (2010). Download - APPL. Retrieved April 29, 2010, from <http://bigbird.comp.nus.edu.sg/pmwiki/farm/appl/index.php?n=Main.Download>