Designing a Mathematical Manipulatives Tutoring System Using POMDPs

Georgios Theocharous
Intel Corporation
Intel Labs
Santa Clara

Nicholas Butko
Cognitive Science Department
University of California
San Diego

Matthai Philipose
Intel Corporation
Intel Labs
Seattle

Abstract

Many elementary mathematics teachers believe that learning improves significantly when students are instructed with physical objects such as coins, called manipulatives. Unfortunately, teaching with manipulatives is a time consuming process that is best with personalized 1-to-1 tutoring. In this paper, we explore the feasibility of an automated physical and personal tutoring solution. We collect, annotate and analyze a rich video data set of teaching sessions. We demonstrate that there is significant structure in the data that can be captured in formal computational decision making models such as Partially Observable Markov Decision Processes. Specifically, we identify the actions the teachers use and how those are conditioned and manipulate different states, such as mood states, mathematical concept states, and coin configuration states. We identify good metrics of teaching performance, as well as define an observation space. We finally present early prototype systems.

Introduction

The use of physical objects, such as coins, rods, cubes, patterns and other concrete objects called manipulatives, is a widely accepted approach for teaching abstract and symbolic mathematical concepts in kindergarten and early grades (Piaget and Szeminska 1941; Bruner 1966; Montessori 1964; Sowell 1989). These researchers showed that interaction with concrete objects provides the basis for abstract thoughts. For example, a child might construct an understanding of the meaning of a 5 cent coin by counting a set of 5 pennies and then associating the value of 5 cents with the physical characteristics of a nickel. The child will also be able understand the meaning of the number 5 through the process of grouping the pennies one by one.

Unfortunately, teaching early mathematics with coin manipulatives is a time consuming process and ideally occurs as a personal 1-on-1 tutoring with a teacher. Each session may last up to 30 minutes and may have to be repeated many times through the school season before the student finally develops cognitive structures for the different concepts, which include naming the coins, sorting them by size and value, counting them, and adding their values.

ICAPS 2010 POMDP Practitioners Workshop, May 12, 2010, Toronto, Canada.

In general, personalized tutoring systems have the potential to radically alter the quality of education by offering a unique learning experience to each student in the classroom (Anderson and Reiser 1985). Bloom (Bloom 1984) reports that students tutored 1-to-1 outperform their peers taught in classroom settings by as much as two standard deviations. In the 25 years since Bloom's report, 2σ has become the gold-standard goal against which tutoring technologies are measured. So far, personalized tutoring systems have found success in areas such as high school math and computer programming (Koedinger and Anderson 1997; Corbett 2001). Unfortunately, for earlier education the development of personalized tutoring systems has not been as successful yet, due to the fact that during those years learning is not just about mentally manipulating abstract and symbolic structures, but also involves physical object manipulation as well as intense student and teacher social interaction.

Dealing with physical objects and social interaction can be challenging, because in addition to the tutoring difficulties, it requires rich data-driven technologies of machine perception and planning under uncertainty. There are video archives of teachers instructing their students that could be analyzed by these technologies (Pea, Lindgren, and Rosen 2006). However, to our knowledge, there are no data sets of teaching interactions specifically using mathematics manipulatives in a 1-to-1 setting. Our goal in this paper is to explore wether is feasible to collect rich data of teaching interactions, annotate them sufficiently and investigate wether there is structure in the data that can be modeled in some formal computational decision making models. Specifically, using our data we identify the actions the teachers use and how those are conditioned and manipulate different states, such as mood states, mathematical concept states and coin configuration states. We identify good metrics of teaching performance, as well as demonstrate machine perception algorithms for perceiving the different states. We show how these dimensions of actions, states, metrics and perceptions formulate a sequential decision making process under uncertainty, which can be captured formally with the mathematical framework of Partially Observable Markov Decision Processes (POMDPs).

In the rest of the paper, we first describe our data set. We then illustrate the different structured parts of the data that formulate the POMDP parameters, such as the actions, states, reward function, and observations. Finally, we present early prototype systems.

The 2σ Dataset

To understand the design dimensions of creating a real-world tutoring system, we collected data of tutoring sessions between a primary education teacher and students. Because this tutoring in this dataset was 1-to-1, we named it the 2σ *Dataset*, reflecting the challenge and promise of powerful personalized tutoring technologies. In all, there were 10 sessions with different students spanning K–2. In each session, the teacher taught skill-appropriate math problems involving manipulating coins. For Kindergartners, such problems might be "Can you show me all the Pennies?" while for second graders they might be complicated subtraction problems. Each tutoring session lasted approximately 20 minutes, and was recorded using 4 simultaneous camera views and audio (Figure 1). During the study, we also interviewed the teacher regarding appropriate curricula for the K–2 age groups, about her own teaching habits, and about her assessment of each child's learning after a tutoring session.

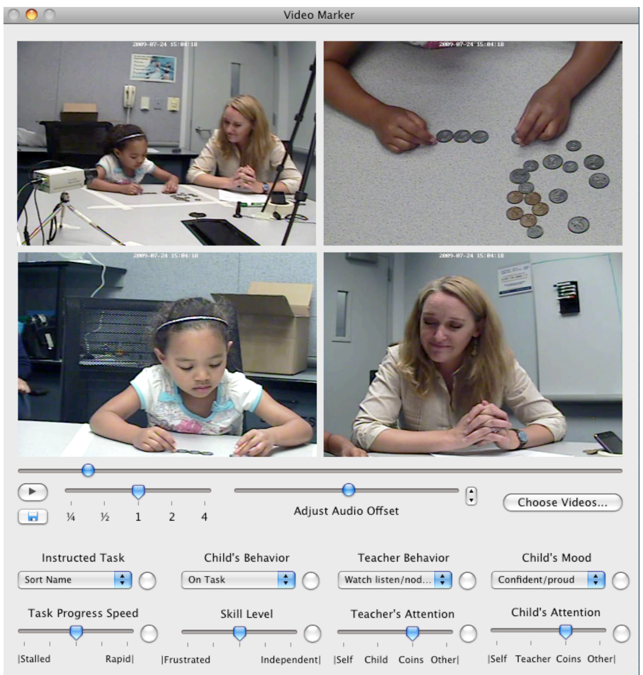

Figure 1: The 2σ Dataset. Recordings were made of 1-to-1 tutoring sessions. Four video cameras simultaneously captured the interaction from four angles: the whole scene, the table workspace, the student's face, and he teacher's face. A video annotating tool shown in the picture above, allowed to annotate the video frames-by-frame with labels including child's mood and teacher's behavior.

Our research team, which consisted of the teacher, a developmental psychologist, and a couple of machine learning and perception researchers agreed on a common set of labels for the teacher actions and child moods. These labels were chosen in such a fashion that would be general enough and common sense that would be hard to dispute. For example, some of the teacher actions included setting up coins, reinforcing the child, by saying "good job" and watching and listening. Some of the child behaviors included thinking, tired, and confused. Furthermore, the labels were chosen for the most part to be mutually exclusive in order to make the domain easier to annotate as well as being sufficient enough to explore the possibility of the existence of structure in the data. Detailed definitions and results on the labels are presented in later sections of the paper.

Our experiments include data from all 10 of the teaching sessions, with very detailed annotations for 2 of the children that we refer to in this paper as child "A" and child "B". Both of them were at the same level and knew little about coins. In both cases, the teacher taught the students using a similar sequence of lessons that included coin names, coin values, sorting, and counting by 1s, 5s and 10s. In later sections we give detailed diagrams of the lesson sequences and concepts.

Actions

From careful observation of the video data we extracted a set of representative teacher actions. Below we give names to these teacher actions and give examples to define their meanings.

- *Set up coins*. The teacher collects the coins in a single group and may separate some out.
- *Diagnose*. The teacher asks question to diagnose whether the child masters some of the concepts that need to be taught. For example, the teacher asks, "do you know what these are?" and points to pennies.
- *Illustrate/teach task*. For example the teacher says, "this is a penny".
- *Ask to do a task*. For example the teacher says, "can you separate all the pennies out?"
- *Hint: point/question/inform/explain*. During a task the teacher sometimes intervenes when the child is not making progress towards the correct solution. For example, if the child moved a dime in a group of pennies the teacher might just point to it and explain that this is not a penny and should be moved away.
- *Watch nod/shake head*. During a task the teacher watches and either gives positive or negative feedback through facial expressions.
- *Reinforce*. After a task is finished the teacher says "good job!" or "Excellent" or sometimes gives a high-five.
- *Explain final result*. For example, once the child separates all the pennies out, the teacher points to them and says "these are all the pennies".
- *Motivate to continue*. For example, the teacher says "you are a Rock star, would you like to play some more games?"

Figures 2 show graphs of the sequence and frequency of teacher actions and how they relate to each other. The graphs are for the two children A and B. It is striking to see that both exhibit very similar transition dynamics and structure, which is evidence that the teacher uses similar teaching strategies

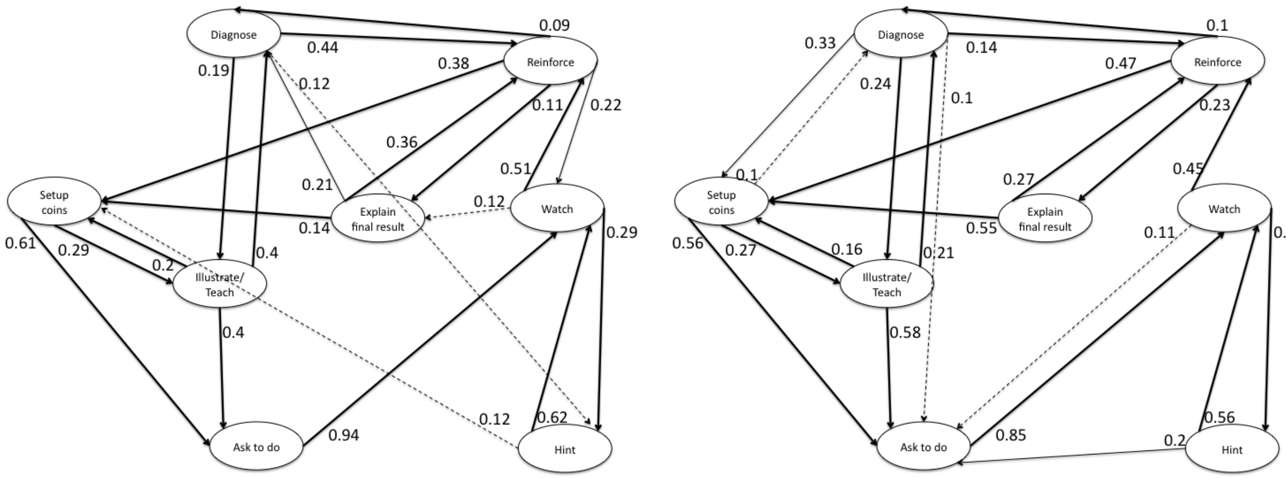

Figure 2: These are graphs of the teacher actions for child A (on the left) and child B (on the right). The darker arrows indicate dominant transitions and are the same in both graphs. The numbers are the transition probabilities observed in the data. We only included edges that had more than 0.09 probability. The graphs shows that the teacher initially diagnoses and sets up coins, which eventually lead to either some teaching or directly asking the child to perform a task. When the task begins the teacher interleaves watching and hinting. When the task finishes the teacher reinforces, sometimes explains the final result and eventually goes back to setting up a new task. The dotted edges are specific to each graph, but have very weak probabilities. The thin continuous edges have stronger probabilities and are specific to each graph too. It is obvious that the teacher strategy has structure that can be seen across different children, which we were able to elicit despite the fact that the data was annotated by two separate people.

across all the children. It also validates the objectiveness of our teacher action definitions.

States

Student mood states

From careful observation of the video data we extracted a set of key situations that represent the mood state of the child. Below we describe these states and give example scenarios to define their meaning.

- *Interested.* In this state the child is either looking at the coin or teacher and listening, but not necessarily thinking about any math concept, but rather being ready to play math games with the coins.
- *Thinking.* In this state the child is manipulating coins to solve one of the assigned tasks.
- *Tired/bored.* In this state the child has already had enough lessons in the session and is beginning to lose her interest and visibly not paying attention to either the coins or the teacher.
- *Confused.* In this state the child, seems to be making the wrong moves with the coins on the table. Other signs are hovering over the coins and being undecided as to what to do and also looking up at the teacher trying to get hints as to what is the right thing to do.
- *Confident/proud.* In this state the child is smiling, moves coins correctly and fast, and when she finishes a task sits back with a smile and lets the teacher look at her accomplishment. She smiles even more when the teacher gives her reinforcement after a successful task completion.
- *Distracted.* In this state the child momentarily switches her attention to different objects and situations other than coins and the teacher. These would be the cameras for example or the chair.
- *Frustrated.* In this state the child is unsure how to make progress, and has stopped trying.

Examples of the above mood states are shown in Figure 3.

Concept states

According to the California department of education, by the end of kindergarten students should be able to count, recognize, represent, name and order up to 30 objects. By the end of grade one, students should be able to identify and know the value of coins and show how different combinations of coins equal the same value. They should be able to count by 2s, 5s and 10s to 100 (Cal 1999). In this paper we target kindergarten and first grade concepts. To elicit the relationships between the different concepts that a child needs to learn and how each concept allows for new concepts to be built upon, we examined our data. We looked at the sequence of lesson being taught and how the teacher was able to progress to harder lesson while interleaving those with diagnosing concepts from previous lessons. In Figure 4 we show the order of lessons and how they relate in terms of concepts and in Figure 5 we derive a cleaner concept space graph.

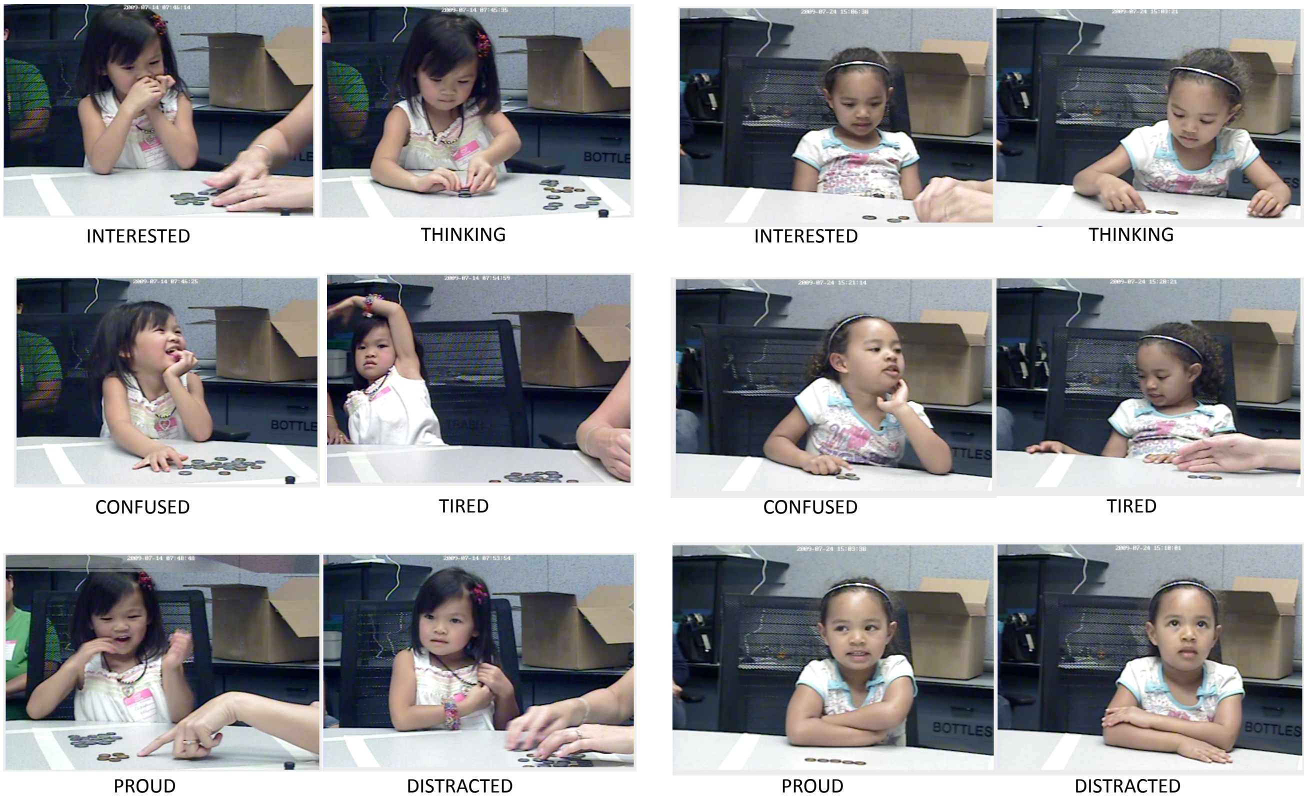

Figure 3: These images describe visually the different mood states that can be observed in the data for child A on the left and child B on the right. Images are surprisingly similar. This result is again a visual evidence that there is significant structure in the data that can be easily annotated and used in our system.

Coin states

From video observations (as can be seen in Figure 3) it was clear that the coin operations involved were mainly groupings on a table of the various types of coins. Therefore, we decided that the coin states are the number of different clusters of coins and the types of coins within the clusters.

Transition dynamics

From our data. we also explored the transition probabilities among the different types of state variables conditioned on the teacher actions.

Mood state dynamics: For the mood state dynamics we counted the frequency of change among the variables for different actions. We did this for child A and B as shown in Table 1 . These tables show sample results for the action *Reinforce*. It is interesting that this action in both students elevates pride, especially when the child is thinking, or is interested. It can also lead to distraction when the student is interested, because the student becomes so confident that she begins to lose her focus. When reinforcement is given while the child is still proud it elevates the interest in both cases, but with some danger of creating confusion. When the child is frustrated the reinforce action motivates her to start thinking again. Similar tables and conclusions as these can be drawn for the rest of the actions.

Concept state dynamics: In figures 4 and 5 is obvious that there is not a single, chain-structured learning trajectory. By choosing the "Ask a question" action, the teacher initiates a learning activity that may ultimately help the child switch forward to a more advanced understanding within a single topic. By choosing the "Diagnose" action, the teacher tests whether the child has mastered a concept.

Coin state dynamics: For the coin state dynamics, coins start from some initial setup and end up in some goal configuration. Each action such as setting up the coins and hinting along the way leads to goal configurations, as can be seen from Figures 8 and 7.

The reward function

In the teaching sessions, we observed the teacher often waits until the child is clearly on the wrong track and won't recover before interfering. From this, we concluded that she tries to minimize the number of hints that she gives. She also tries to get the student to learn as many lessons as possible. These are two quantities that should be part of our reward function. Additionally, there are intermediate states that seem good to be always on such as interest and confidence. A simple goal-based reward function that penalizes hints and tries to cover as many concepts as possible should automatically provide for a policy that boosts confidence to

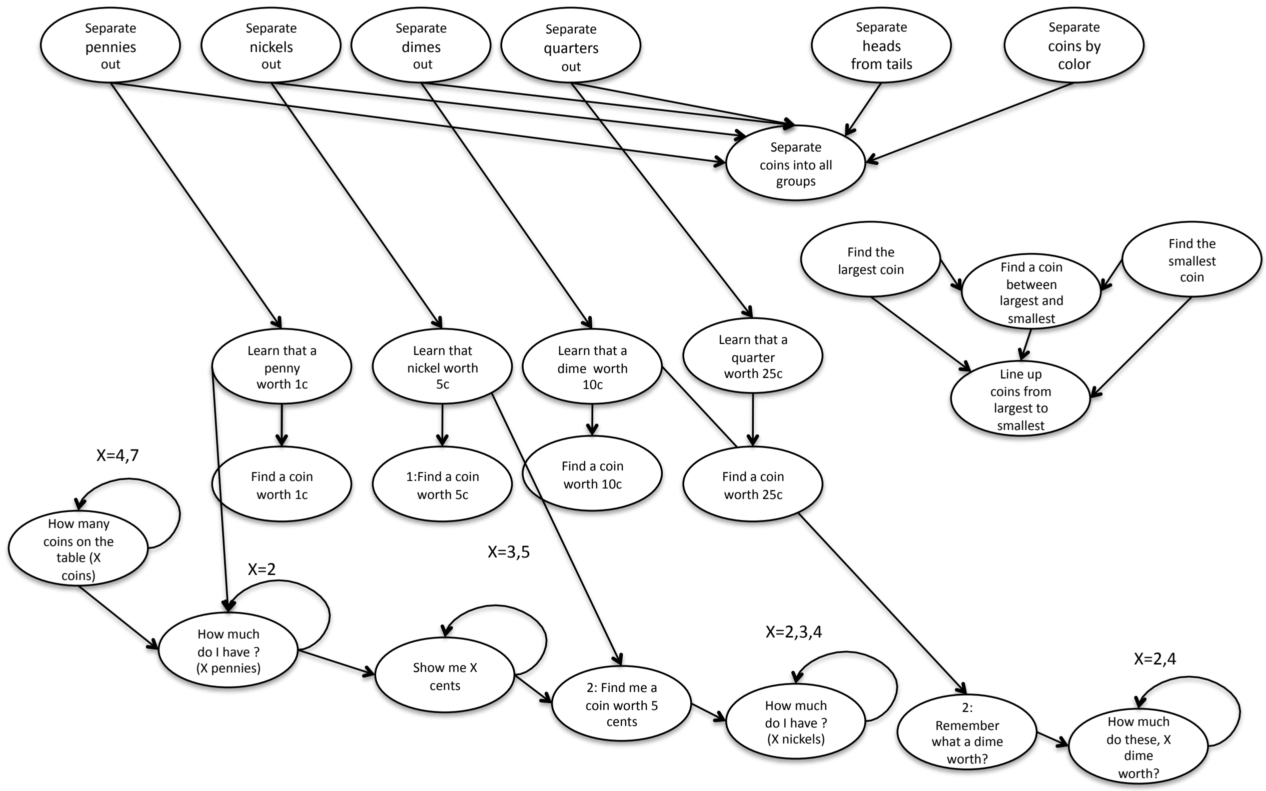

Figure 4: The circles describe the lessons that were observed in one of the video sessions. The order of the lessons was given from left to right and top to bottom. The arcs indicate the fact that in order to be able to do a lesson one needs to know concepts leant from a previous lesson. The teacher often traversed those arcs backward every time a new lesson was about to start, to diagnose and make sure that the building concepts were already in place. This graph was extracted by observing the lesson sequence for child B and clearly demonstrates that there is a structured lesson order.

promote interest for example, while diagnosing as needed the true concept level of the student.

Observations

Mood state perception

Perceiving the mental states that we described earlier can be achieved through facial expression recognition. Facial expression recognition can be understood by characterizing the individual muscles in the face and their appearance when they move. A taxonomy of these movements is described by the Facial Action Coding System (FACS) (Ekman, Friesen, and Hager 2002). FACS based expression recognition systems have shown dramatic progress recently. In some cases, are as accurate at predicting mental states as human experts trained in facial analysis and much better than untrained humans (Littlewort, Bartlett, and Lee 2009).

We used the Computer Expression Recognition Toolbox (CERT) (Bartlett et al. 2006) to measure 106 frame-by-frame indices of facial information, including information about facial expression, head position, gaze direction, and so on (Figure 6B). We used a linear-regression-based model to predict the focus of the student's attention from these facial indices. Using this method, the computer was able to correctly predict where the student was attending over 90% of the time. More importantly, there was a systematic 80% correlation between actual attentional shifts and predicted attentional shifts (Figure 6D). As the student becomes fatigued, the system perceives a dramatic rise in attention shifts. Such a signal could readily be used by to judge when to infer the Tired/boredom mood state.

Coin state perception

Coin detection Coins are essentially circular objects. A traditional Computer Vision approach would leverage this geometrical observation by using the simple and mathematically elegant Hough Transform approach. This approach first searches an image for edges, and then considers the mathematical space of all possible circles to which each edge component might belong (Yuen et al. 1990). Where a large number of edges could belong to the same circle, we decide that this circle is a coin.

A more recent class of approaches to finding any object in a scene is called the Cascaded Classifier approach, and it was originally used to find faces very quickly in images (Viola and Jones 2001). While the Hough Transform approach is rooted in Geometry, the Cascaded Classifier is based on

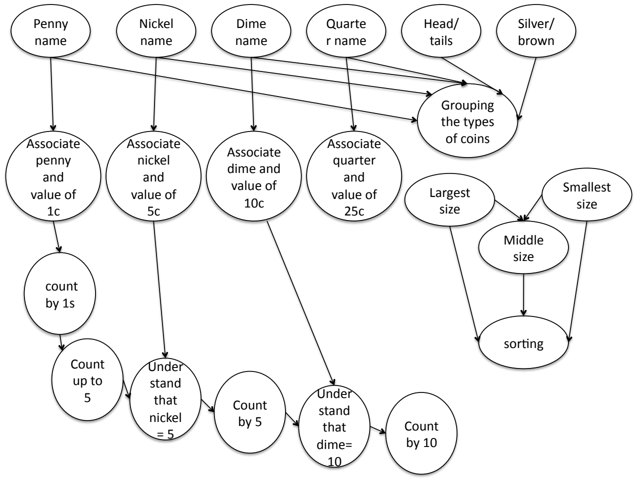

Figure 5: From Figure 4, one can derive a concept space and the order in which the concepts need to be taught. This graph shows concepts learnt during kindergarten and first grade. They involve concepts such as coin names, largest versus smallest, sorting, association of coin names with values, understanding coin values trough counting and learning to count by 1s, 5s and 10s.

Probability Theory and is expected to be more robust to uncertainty and ambiguity in images. Our results showed that the Cascaded Classifier approach has about 10-100x fewer false alarms for any level of misses compared to the Hough Transform approach, while analyzing the image over six times faster.

Coin grouping We implemented a fast spectral clustering method similar to (Shi and Malik 2000) that combines stability criteria from both eigenvalues and eigenvectors to estimate the number of clusters and identify the clusters. Compared to older clustering techniques like k-means clustering and mixture-of-Gaussians, spectral clustering makes groups that match human agreement better (Ng, Jordan, and Weiss 2001). Examples of the groups tracked by our system are shown in Figure 7.

Still, even two humans may look at the coins and disagree on how many groups the student has made. New techniques allow us to quantify this ambiguity in the number of groups (Lewis 2009). Here, the 2σ Dataset gave us an important insight. We observed an interesting group refinement teaching behavior: the teacher would refine the groups made by the student so that they would be less ambiguous (Figure 7, bottom right). A future research avenue will be to identify ambiguous groups and encourage the student to make them clearer, helping to improve her illustrative communication skills.

Prototypes

As an ongoing part of our explorations we have implemented two major prototypes. The first one, in Figure 8 shows the physical coin detection and grouping of our system..

Table 1: This a sample of transition probabilities that can be estimated from the data. This was done for the action Reinforce for children A and B. Each row of the table shows the probability distribution of the mood changing to the moods in the different columns. The probability support show remarkable similarity across the children despite the fact that labels came from two different people. The names of the moods are abbreviated versions of the moods defined in the student mood states section.

ChldA	Intr	Thnk	Tird	Conf	Prd	Distr	Frst
Intr		0.1	0.1	0.1	**0.5**	**0.2**	
Thnk	**0.43**		0.14	**0.29**	0.14		
Tird	**0.2**	**0.2**			**0.6**		
Conf		**0.4**	**0.40**		**0.2**		
Prd	**0.33**	0.13	**0.20**	0.07		**0.27**	
Distr	**0.33**	0.17			**0.5**		
Frst		**1.0**					
ChlldB							
Intr		**0.5**	0.17		0.17	0.17	
Thnk	0.13				**0.87**		
Tird							
Conf.							
Prd	**0.27**	**0.73**					
Distr	**0.5**				**0.5**		
Frst							

The second, in Figure 9 shows a first POMDP-based version. This version assumed that coin perception was completely observable and that the concept space was linear from easier to hardest concept. In addition, it only considered a single mood state of *distracted*. This model was implemented in the factored framework of algebraic decision diagrams, where the policy was computed using the Symbolic Perseus package (Poupart 2005).

We tested this model in a virtual coin manipulative platform combined with the facial expression recognizer. We implemented five lessons where each one was about separating a particular type of coin, starting from pennies then nickels, dimes, quarters and dollars. Each lesson included the additional type of coin that needed to be separated out. Each lesson was split into two, where in one lesson there was detailed hinting (a teaching lesson), and in the next lesson the child was asked to do the task without any hints (a testing lesson). The number of levels was equal to the number of lessons. The policy that emerged when the POMDP was solved would first try to diagnose the students level and then try to guide her through the completion of all lessons, by first teaching and then testing, and falling back to teaching when the student could not finish a testing lesson. When the student attention faded off, the tutor tried to bring it back by asking the student to pay attention.

Summary and Conclusions

In this paper we proposed to build a mathematical math coin tutoring system. Unlike previous work on intelligent tutor-

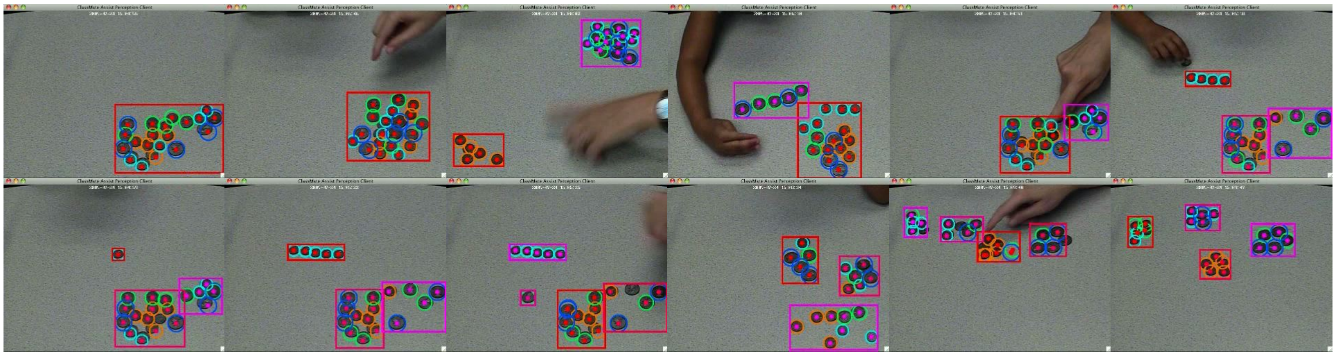

Figure 7: Groups of coins identified by our system. In the bottom right two examples, the student has made four groups of coins that could be seen as three groups (left). The teacher then helps her separate the groups to become more clear (right).

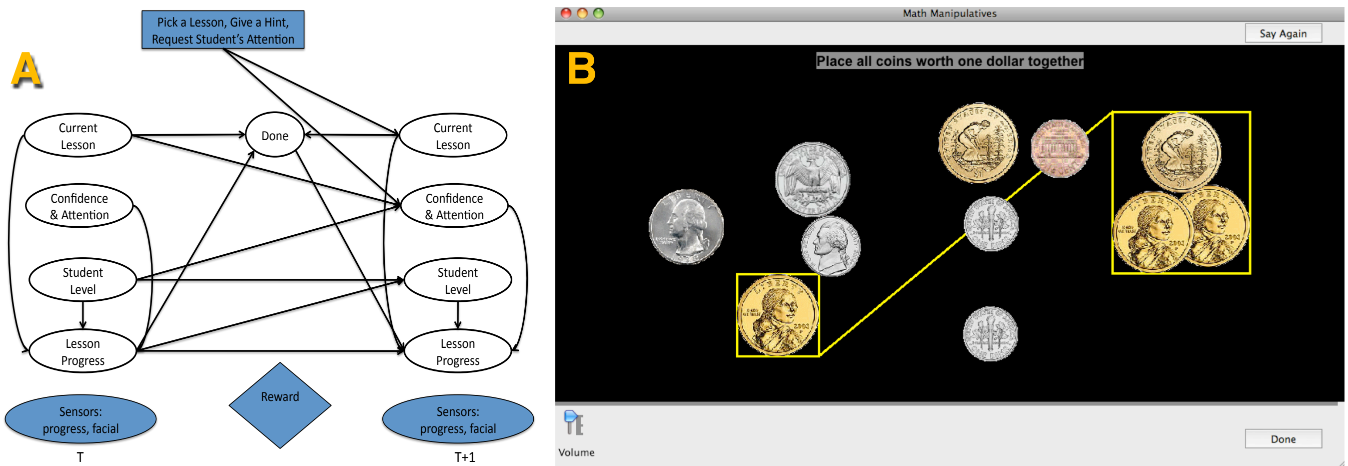

Figure 9: The Figure above on the left shows the 2 time-slice POMDP model, which takes as inputs observations about the progress of tutoring and student facial expressions. The actions decide what lesson to give next and whether it should be accompanied with detailed visual and text to speech hints. The actions can also ask the student to pay more attention. The model reasons about the student level, and whether she is paying attention. The reward function is positive only when the current lesson matches the student level and negative otherwise. The intuition is that a student would not learn if the lesson was too easy or too difficult. The figure on the right shows a virtual coin platform where this POMDP was tested at. Due to the virtual platform the progress variable was completely observable.

ing solutions we faced the challenge of solving the problem of physical object interaction as well as the problem of student mood recognition through non-invasive sensors such as cameras. By collecting video teaching sessions and analyzing them, we showed throughout the paper that there is remarkable structure in the data, that can be captured in the form of a Partially Observable Markov Decision Process. Finally, we showed a preliminary POMDP model that has begun to exhibit intelligent teaching strategies and was able to reason about simple child moods, diagnose the child's level and advance it.

Still much work remains in completing this system with a thorough elicitation of all the parameter of the domain and computing and evaluating a POMDP tutoring policy that resembles complete teacher strategies. In the coming months we plan on collecting and analyzing data with more students, learn the POMDP transition dynamics for the different states and actions, using machine learning to improve and augment the prediction of POMDP observations, and finally validating the whole system in actual classrooms.

References

Anderson, J. R., and Reiser, B. J. 1985. The lisp tutor: it approaches the effectiveness of a human tutor. *BYTE* 10(4):159–175.

Bartlett, M. S.; Littlewort, G. C.; Frank, M. G.; Lainscsek, C.; Fasel, I. R.; and Movellan, J. R. 2006. Automatic recognition of facial actions in spontaneous expressions. *Journal of Multimedia* 1(6):22.

Bloom, B. S. 1984. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher* 13(6):4–16.

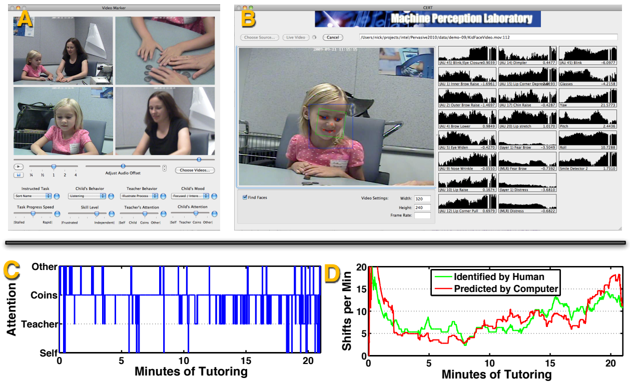

Figure 6: Analyzing Facial Expressions. **A:** We created a tool to quantify aspects of teacher-student interaction, including "What is the student paying attention to?" **B:** The Computer Expression Recognition Toolbox (Bartlett et al. 2006) automatically extracted 106 facial indices from each frame in the teaching session. **C:** *Y- Axis:* Human judgments of the student's attention targets. *X-Axis:* Time-point during a 22 minute turing session. **D:** The rate of the student's attention shifts rise as the student becomes fatigued. Note the high correlation between human observations and ClassmateAssist's predictions (the predictions have been scaled down by a factor of 2.6 to compensate for noise).

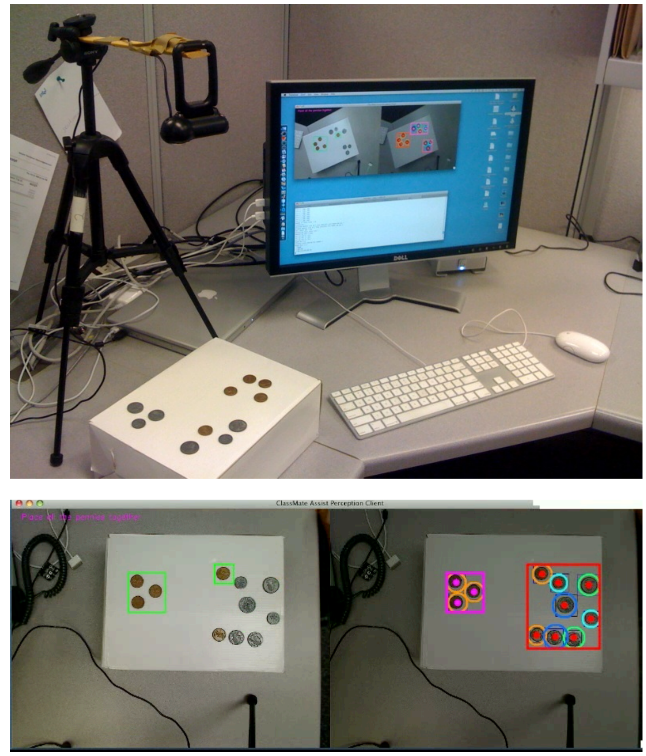

Figure 8: This is our current prototype system in action. The top image shows the setup. The bottom right image shows the different clusters and types of coins within the clusters. The bottom left image is a Hint action that is accompanied with a text to speech message asking to move that penny with the group of pennies.

Bruner, J. S. 1966. *Toward a Theory of Instruction*. Cambridge, MA: Harvard University Press.

California Department of Education. 1999. *Mathematics Content Standards for California Public Schools Grade Twelve Kindergarten Through*.

Corbett, A. T. 2001. Cognitive computer tutors: Solving the two-sigma problem. In *Proceedings of the 8th International Conference on User Modeling 2001*, volume 2109 of *Lecture Notes In Computer Science*, 137–147. Springer-Verlag.

Ekman, P.; Friesen, W.; and Hager, J. 2002. *Facial Action Coding System (FACS): Manual and Investigator's Guide*. Salt Lake City, UT: A Human Face.

Koedinger, K. R., and Anderson, J. R. 1997. Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education* 8:30–43.

Lewis, J. M. 2009. Finding a better k: A psychophysical investigation of clustering. In *CogSci 2009 Proceedings*, 315–320.

Littlewort, G. C.; Bartlett, M. S.; and Lee, K. 2009. Automatic coding of facial expressions displayed during posed and genuine pain. *Image and Vision Computing* 27(12):1797 – 1803. Visual and multimodal analysis of human spontaneous behaviour:.

Montessori, M. 1964. *The Montessori Method*. New York: Schoken Books.

Ng, A. Y.; Jordan, M. I.; and Weiss, Y. 2001. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems 14*, 849–856. MIT Press.

Pea, R.; Lindgren, R.; and Rosen, J. 2006. Computer-supported collaborative video analysis. In *7th International Conference of the Learning Sciences*.

Piaget, J., and Szeminska, A. 1941. *The Childs Conception of Number*. New York: W. W. Norton Co.

Poupart, P. 2005. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. Ph.D. Dissertation, University of Toronto, Toronto.

Shi, J., and Malik, J. 2000. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(8):888–905.

Sowell, E. J. 1989. Effects of manipulative materials in mathematics instruction. *Journal for Research in Mathematics Education* 20(5):498–505.

Viola, P., and Jones, M. 2001. Rapid object detection using a boosted cascade of simple features. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yuen, H. K.; Princen, J.; Illingworth, J.; and Kittler, J. 1990. Comparative study of Hough Transform methods for circle finding. *Image and Vision Computing* 8(1):71–77.