Leveraging Student Goal Setting for Real-Time Plan Recognition in Game-Based Learning

Alex Goslen[1], Dan Carpenter[1], Jonathan P. Rowe[1], Nathan Henderson[1], Roger Azevedo[2], James Lester[1]

[1] North Carolina State University, Raleigh NC, USA
{amgoslen, dcarpen2, jprowe, nlhender, lester}@ncsu.edu
[2] University of Central Florida, Orlando, FL, USA
roger.azevedo@ucf.edu

Abstract. Goal setting and planning are integral components of self-regulated learning. Many students struggle to set meaningful goals and build relevant plans. Adaptive learning environments show significant potential for scaffolding students' goal setting and planning processes. An important requirement for such scaffolding is the ability to perform student plan recognition, which involves recognizing students' goals and plans based upon the observations of their problem-solving actions. We introduce a novel plan recognition framework that leverages trace log data from student interactions within a game-based learning environment called CRYSTAL ISLAND, in which students use a drag-and-drop planning support tool that enables them to externalize their science problem-solving goals and plans prior to enacting them in the learning environment. We formalize student plan recognition in terms of two complementary tasks: (1) classifying students' selected problem-solving goals, and (2) classifying the sequences of actions that students indicate will achieve their goals. Utilizing trace log data from 144 middle school students' interactions with CRYSTAL ISLAND, we evaluate a range of machine learning models for student goal and plan recognition. All machine learning-based techniques outperform the majority baseline, with LSTMs outperforming other models for goal recognition and naive Bayes performing best for plan recognition. Results show the potential for automatically recognizing students' problem-solving goals and plans in game-based learning environments, which has implications for providing adaptive support for student self-regulated learning.

Keywords: Plan Recognition, Game-Based Learning, Self-Regulated Learning.

1 Introduction

Self-regulated learning (SRL) describes learning that is guided by metacognition, strategic action, and motivated behavior [17, 20]. A key attribute of SRL is its focus on goal-driven learning. Self-regulated learners formulate goals and develop plans for achieving them, which are monitored and adapted based upon learners' self-evaluated progress [22]. Goal setting and planning is particularly important in scientific inquiry where learning is guided by students' curiosity and motivation for acquiring knowledge, and where students need well-defined plans to carry out productive investigations [9]. Self-regulated learners set goals and sub-goals to complete a learning task [23]. To achieve their goals, students build plans that outline approaches, such as strategies or sequences of actions they intend to enact [22].

Learning environments that support goal setting and planning foster positive emotions and can create opportunities for student success [4]. Adaptive learning environments provide a way to scaffold student goal setting and planning in a manner that is individualized to each student. An important component of adaptive scaffolding is recognizing student goals and plans while the learner solves problems within the learning environment [1]. The task of plan recognition is focused upon predicting an individual's high-level goal, and the plan for achieving it, based on lower-level observations of the individual's strategies and actions. Goal recognition is considered a special case of plan recognition where the prediction task is focused only on recognizing high-level goals [3]. While there has been considerable work on modeling student knowledge in adaptive learning environments, limited research has been done on student plan recognition.

This paper presents a novel student plan recognition framework that uses machine learning to build goal and plan recognition models to predict students' problem-solving goals and the series of actions students intend to achieve them. The framework is evaluated with CRYSTAL ISLAND, a game-based learning environment for middle school microbiology, in which students utilize a novel planning support tool that encourages them to externalize their goal setting and planning processes during science problem solving. We utilize trace log data from students' interactions with the planning support tool, as well as their other problem-solving actions in the game, to train multi-label classification models to predict students' goals and plans. Specifically, we predict labels derived from student goals and a cluster-based representation of planned actions for the goal recognition and plan recognition tasks, respectively. We present results from a comparison of six machine learning-based classification techniques (support vector machines, random forest, naive Bayes, logistic regression, multilayer perceptron, long short-term memory networks) for modeling student goals and plans in CRYSTAL ISLAND. Our findings indicate that long short-term memory (LSTM) networks show promise in both goal and plan recognition tasks, which have potential to inform real-time scaffolding to support student goal setting and planning.

2 Related Work

Goals and plans are critical in SRL. Winne and Hadwin's Information Processing Theory of SRL posits that, throughout goal setting, planning, and enactment, students are continually monitoring and controlling how their learning is unfolding so that they are in control of their learning processes, and they are monitoring how effective these processes are in contributing to learning, information processing, and task completion [21, 22]. This implies that students know to set subgoals, use the appropriate and effective cognitive and metacognitive SRL strategies, and adapt the use of these strategies. However, how middle school students set goals and plans during science problem solving is not well understood, leaving key questions regarding how to effectively support student goal setting and planning in science learning environments [20].

Despite the importance of student goal setting and planning in SRL, there has been relatively little work on devising computational models of student plan recognition in adaptive learning environments. An important exception is work on Andes, an intelligent tutoring system for physics, which utilized Bayesian networks to model student plans and make predictions about student actions during problem solving [6]. This work exemplifies a successful application of plan recognition that informs adaptive support to provide students with specialized help through hints.

Prior work has also investigated a restricted form of student plan recognition, i.e., student goal recognition, using trace log data from student interactions with a game-based learning environment. A set of eleven goals were inferred from player activity. Authors explored a variety of event representations, models, and different evaluation metrics for accuracy and efficiency [10, 13, 15]. The most recent work found using one-hot encoding vectors to represent in-game events as input for LSTMs achieved the best performance predicting these game activity-derived goals [14]. Additionally, prior work has highlighted similarities between natural language processing and plan recognition, demonstrating the effectiveness of applying various natural language processing techniques (NLP) to plan recognition tasks [2, 7].

In this work, we extend these findings by devising a novel student plan recognition framework that uses students' in-game actions and planning support tool usage as observed input and leverages neural embedding-based representations of student action sequences from students' externalized plans to produce target labels. This framework utilizes two multi-label classifiers to compare six machine learning-based classification techniques for modeling student goals and plans in CRYSTAL ISLAND. Our aim is to demonstrate that a machine learning-based framework for student plan recognition can accurately model student goals and plans during science problem solving in a game-based learning environment.

3 Goal Setting and Planning in CRYSTAL ISLAND

3.1 Planning Support Tool in CRYSTAL ISLAND

To investigate predictive models of student goal setting and planning during science problem-solving, we utilize a game-based learning environment for middle school microbiology. CRYSTAL ISLAND features an interactive science mystery that engages students in a process of scientific inquiry as they investigate the source of a mysterious disease outbreak on a remote island research station. Students assume the role of an infectious disease investigator who is tasked with diagnosing the outbreak and recommending a treatment and prevention plan.

In order to support student goal setting and planning in CRYSTAL ISLAND, we have developed a planning support tool that incorporates design concepts from visual programming languages [19] and AI planning [8]. Specifically, students utilize a block-based visual interface to assemble hierarchical (i.e., two-layer) plans consisting of high-level goals and low-level sequences of actions that can be enacted in CRYSTAL

ISLAND (Fig. 1). Students choose from a palette of pre-defined goal and action blocks in the tool. The goal blocks represent possible subgoals that students may wish to achieve on their way to solving the mystery, which are the overarching goal of the problem-solving scenario. Example goals include “Learn about outbreak” and “Report evidence-based diagnosis”. Each action block lists specific steps that students can take to achieve a goal. Example actions include “Read about how diseases spread” and “Use scanner to test objects”. Goal and action blocks are connected to form plans. For example, if a student sets a goal to “Explore Island”, they can place movement actions such as “Go to Infirmary” under the goal block to indicate a necessary step needed to complete the specified goal.

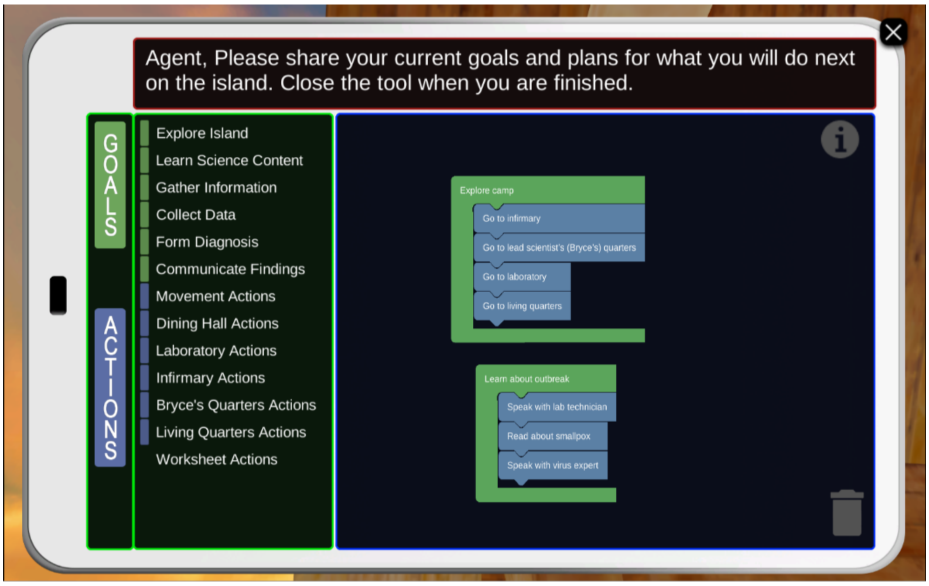

Fig. 1. Planning support tool in the CRYSTAL ISLAND learning environment.

Prior to engaging with CRYSTAL ISLAND, students watch a short, narrated video that introduces the planning support tool and demonstrates how to use the tool to build a plan. Once students begin using the game, they are prompted early on to set their own goal(s) and build plans using the tool. Students use the tool by dragging and dropping goal and action blocks onto a virtual canvas that serves as the planning area. After they have formulated a plan, they can close the tool and choose to enact their plan (or not) within the CRYSTAL ISLAND virtual environment. If students complete a goal or want to remove a goal that they previously chose, they can drag the block to a trash icon in the planning support tool. Upon deleting a goal block, students are prompted to indicate whether they reached the discarded goal or not. Students are presented with mandatory prompts to use the tool at major milestones in the science mystery, as well as every thirty minutes during gameplay, and may also voluntarily access the planning support tool at any time.

3.2 Goal Setting and Planning Dataset

A study was conducted with 144 middle school students in the United States. Of these students, 60% were female and the average age was 13.2 years. Students played CRYSTAL ISLAND remotely during asynchronous science class time due to a transition to remote learning during the COVID-19 pandemic. Students were instructed to access the game over a two-day span and were not given a time limit to complete the game. Students also completed pre- and post-tests to assess science content knowledge, along with a brief demographic survey. The pre- and post-tests consisted of 17 multiple choice questions about microbiology that could be answered based on the curricular content in CRYSTAL ISLAND. Interaction logs of students' actions within the game and usage of the planning support tool were logged automatically. Students on average played the game for 94.7 minutes (SD = 47.7).

4 Student Plan Recognition in Game-Based Learning

We present a student plan recognition framework that utilizes trace log data from students' planning support tool usage and gameplay to induce multi-label classification models to predict student goals and plans during science problem solving in the CRYSTAL ISLAND game-based learning environment. The input to the student plan recognition models is a feature vector representation of student actions distilled from students' trace log data from the game. Students' goals and plans from the planning support tool are used to devise labels for training the plan recognition models using a supervised learning approach. Specifically, each student action is annotated with a goal label and plan label that signify the goal students are attempting next and the set of actions they plan to take to achieve that goal, respectively. Below we describe the event sequence representation, labeling approach, and evaluation methods utilized in the student plan recognition framework.

4.1 Event Sequence Representation

Student interactions with CRYSTAL ISLAND generate trace log data that consists of timestamped sequences of actions taken by students while playing the game. We refer to these as event sequences. Based on prior work, each student action in an event sequence is represented by three types of features: action types, action arguments, and locations [14].

- **Action type.** Action type refers to categories of in-game activities undertaken by the student within the learning environment. These actions ranged from viewing posters and reading articles about viruses and bacteria to scanning items and talking to characters. For example, "Movement" signifies moving to a particular location or "Conversation" means a student had a conversation with a non-playable character in the game. There were 9 total action types.
- **Action argument.** Action arguments provide more details about the action type. For example, if the action type is "BooksAndArticles", the title of the book or arti-

cle the student read is included as the action argument. There were 108 unique action arguments.
- **Location.** Location represents the region of the virtual island where the action took place. If the action type is "Movement", the location is the place where the student moved to. There were 24 unique locations in the game.

To prepare the dataset for student plan recognition, event sequences were segmented according to student usage of the planning support tool. The intuition for this approach is that students externalize their goals and plans using the planning support tool. Afterward, they enact their plans by performing actions in the game. An event sequence concludes when the student next reopens the planning support tool and changes their goals or plans, thereby initiating a new event sequence. In other words, an event sequence begins with the first student action after the planning support tool is closed. The event sequence concludes with the last student action before next opening the planning support tool. In total, there were 400 event sequences across all students. The length of event sequences ranged from 1 to 454, with a median of 30. The event sequences were constructed cumulatively to allow for action-level prediction, with the maximum length of a sequence being 30. For example, events one through 30 between planning support tool uses would translate to 30 rows of data, the first row only containing the first event, the second containing the first and second event, and so on up to 30. Because LSTMs require fixed-length input sizes, sequences of less than length 30 were zero-padded. Once the event sequences were created, we used one-hot encoding to convert student actions into a vector representation. One-hot encoding vectors have been shown to work effectively in prior work on student goal recognition in game-based learning environments [14].

Each plan that students constructed in the planning support tool consisted of a goal and a set of actions. We utilized student goals from the planning support tool to devise labels for the goal recognition task, and we used sets of actions from the planning support tool to devise labels for the plan recognition task. Event sequences were assigned labels based upon students' plans from their prior use of the planning sup-port tool. To illustrate, consider the following example. A student opens the planning support tool and creates a plan consisting of a goal and a set of actions (i.e., Plan 1). The event sequence that follows this planning support tool interaction is assigned a goal and plan label based upon the goals and set of actions that are included in Plan 1.

4.2 Goal Recognition Labels

The planning support tool allows students to select from 20 possible goals and was designed so that each goal falls into one of 5 categories: (1) Collect Data, (2) Communicate Findings, (3) Form Diagnosis, (4) Learn Science Content, and (5) Gather Information. For our analysis, these five categories serve as goal labels, rather than using all 20 lower-level goals. Since students can create multiple plans at a time, we formalized goal recognition as a multi-label classification task, assigning each event sequence a binary label vector in which each element of the vector corresponds to a possible goal category. The dataset had the following distribution of goal categories:

(1) Collect Data: 22%, (2) Communicate Findings: 4%, (3) Form Diagnosis: 13%, (4) Learn Science Content: 22%, and (5) Gather Information: 40%.

4.3 Plan Recognition Labels

The planning support tool allows students to select from 55 possible actions to build plans for achieving their intended goals. Similar to goals, the palette of actions in the planning support tool was divided across six action categories. We utilized these higher-level categories to represent the actions in students' plans. Students' plans typically contained more than one action associated with a goal, with an average of 2.58 (SD = 1.96) actions per goal. To convert the action sets into labels for student plan recognition, the following procedure was applied. First, all actions in a plan were concatenated using the same order that students specified in the planning support tool. Next, SpaCy word embeddings were applied to each categorical action set [18]. The resulting embeddings were averaged for each set of actions in a plan. Next, k-means clustering was applied to the word embeddings to separate the plans into clusters. The number of clusters was determined visually using the Elbow method, resulting in 4 distinct groups of action sets [5]. The resulting clusters were used to derive 4 possible class labels for plan recognition.

When reviewing patterns of action categories within the clustering, it seemed that the most used action category in each plan aligned within the clusters. Cluster 0 (9%) represents plans that mostly contain "Read Science Content". Cluster 1 (30%) represents primarily "Explore" action category usage. Cluster 2 (33%) represents plans that contain mostly "Gather and Scan Items", and Cluster 3 (28%) represents plans that contain mostly "Speak with Characters".

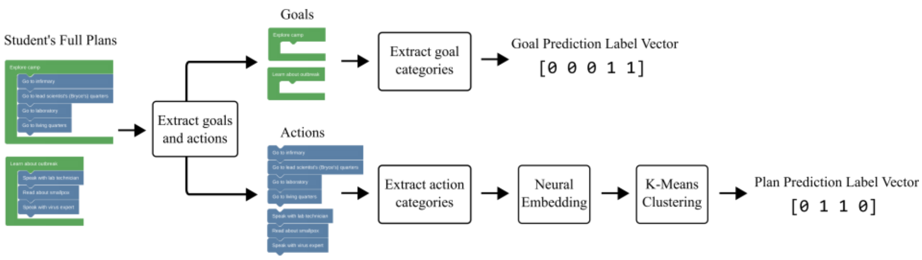

Fig. 2. Procedure for translating student plans into multi-label vectors for student goal recognition (top) and student plan recognition (bottom).

These labels were assigned to event sequences in a multi-label fashion, similar to the goal recognition task. Figure 2 illustrates the process for translating students' plans into label vectors for goal recognition and plan recognition, respectively.

4.4 Model Selection and Evaluation

We examined six different supervised learning techniques to induce multi-label classifiers for student goal recognition and plan recognition: support vector machines (SVM), random forest (RF), naive Bayes (NB), logistic regression (LR), multi-layer perceptron (MLP), and long short-term memory (LSTM) networks. These models were selected to establish a general baseline of results. Since this task has not been completed previously, we chose mostly non-sequential models to analyze patterns of overall performance. We performed nested 5-fold cross validation using an iterative grid search for hyperparameter tuning of all six models. Due to the limited representation of some of the labels, we could not choose a k any greater than 5 without having one of the classes no longer represented in the training or test set. We used a stratified student-level split within the nested cross validation to maintain a similar class distribution across the training and test sets and to prevent data leakage between folds. For the non-LSTM models, we took the sum of the one-hot encoding vector across events to handle different lengths of sequences and created a single vector representing the number of times each type of action occurs in a sequence. The LSTM received the entire one-hot encoding vector as input.

We utilized the macro-average F-measure to evaluate the models. F-measure has been shown to be a good indicator of model performance in multi-label classification tasks because it highlights incorrectly classified labels by basing the calculations on false positives and false negatives [11, 12]. Since false positives and false negatives are instances that can create user frustration, they are important indicators of performance in an adaptive learning environment. In addition, we have an uneven distribution of classes for both the goal and plan recognition tasks. Macro-average F-measure works well on imbalanced datasets because it computes the average for each class label separately and then aggregates them together [16]. Therefore, this metric is well suited for evaluating models intended for use in adaptive learning environments.

5 Results

To investigate the effectiveness of the machine learning-based goal recognition and plan recognition models, we compared all models against a baseline model that always predicts the majority class.

5.1 Goal Recognition Results

Goal recognition results for all six models are shown in Table 1. All models except random forest improved on the baseline in four out of five goal categories. Random forest appeared to overfit to the majority class, and it performed similarly to the baseline model. In some cases, an imbalance of the class labels causes classifiers to ignore the less-represented classes, which could cause a model to overfit to the majority class. Because random forest makes decisions based on information gain, it makes sense that it would often favor choosing the majority class. The LSTM was among the top two highest-performing models for four out of five classes, including one of the least represented goal categories (i.e., Form Diagnosis). SVM, NB, LR and MLP all improved on the baseline with respect to the macro-average F-measure. The LSTM showed the greatest improvement on the baseline with a 42% relative improvement in the F-measure.

Table 1. Average F-measure for each classification model and goal category in the student goal recognition task. Distributions in results represent the test set and are averaged across 5 folds of cross validation.

	Collect data	Comm. Findings	Form diagnosis	Learn science content	Gather info.	Overall
N dist.	21%	3%	3%	24%	49%	
	F	F	F	F	F	Macro F
Maj.	0.00	0.00	0.00	0.00	0.74	0.15
SVM	0.20	0.07	0.22	0.20	0.71	0.28
RF	0.00	0.00	0.00	0.00	**0.74**	0.15
NB	**0.42**	0.12	0.23	**0.43**	0.58	0.35
LR	0.24	0.16	0.40	0.27	0.67	0.35
MLP	0.29	0.19	0.31	0.16	0.64	0.32
LSTM	0.32	**0.35**	**0.47**	0.35	0.62	**0.42**

5.2 Plan Recognition Results

Table 2 shows the plan recognition results for all six machine learning models, as well as the baseline. For the plan classes 0, 1 and 2, all machine learning-based models improved on the baseline. Naive Bayes showed the highest macro-average F-measure for plan classes 0 and 1. This could be due to the model attributing most input actions to all four plan classes, causing the results to be improved. The multi-layer perceptron outperformed the baseline model on the majority plan class, which indicates it more precisely predicted the majority plan class than any other approach. The LSTM performed best again for the least represented plan class. All models improved on the macro-average F-measure compared to the majority baseline.

Table 2. Average F-measure for each classification model and plan class in the student plan recognition task. Distributions in results represent the test set and are averaged across 5 folds of cross validation.

Plan class	0	1	2	3	Overall
N dist.	8%	27%	28%	36%	
	F	F	F	F	Macro F
Maj.	0.00	0.00	0.00	0.55	0.14
SVM	0.36	0.35	0.20	0.29	0.30
RF	0.31	0.41	0.00	0.18	0.22
NB	**0.53**	**0.54**	0.17	0.48	**0.43**
LR	0.46	0.50	0.21	0.43	0.40
MLP	0.29	0.19	**0.31**	**0.64**	0.32
LSTM	0.48	0.47	**0.31**	0.38	0.40

6 Discussion

Overall, the machine learning-based models show clear improvement with respect to macro-averaged F-measure over a naive baseline on the student goal and plan recognition tasks. Prior work on student goal recognition found LSTMs to be the best performing model on a multiclass goal recognition task [14]. Our work extends these findings by showing that LSTMs also perform effectively for goal recognition in a multi-label context. Student plan recognition proved to be a more difficult task than student goal recognition. Unlike goal recognition, there was not a single model that performed best across all plan classes. For example, naive Bayes showed the highest macro-average F-measure, but its predictions were consistently every plan class for a given set of input actions. This type of prediction is not ideal to inform run-time scaffolding because it does not provide a precise indication of what students are planning.

The imbalanced labels in the dataset presented challenges in training and evaluating the models for student goal recognition and plan recognition. However, it is representative of the types of plans generated by students through their use of the planning support tool in CRYSTAL ISLAND. Notably, we saw planning support tool usage decrease over time, with students trending toward using the tool frequently in the first half of the game, but less so as time went on. There were also different levels of granularity associated with the different goal categories and plan classes. For example, goals related to gathering information typically occurred early in the game, and they encompassed a relatively broad set of possible actions. In comparison, goals in the Communicate Findings category ideally occurred after a student formed a hypothesized diagnosis, which typically occurs later in the game. The steps involved to communicate findings are directly outlined in the game, and as a result, one would expect plans related to this goal to occur less frequently. Encouragingly, the results show the promise of using machine learning-based multi-label classification techniques for student goal and plan recognition despite the inherent challenges of imbalanced data.

The wide variety of student plans also presented distinctive challenges for plan recognition. Some students frequently used the planning support tool and updated plans without being prompted, while other students opened and closed the planning support tool only when required. This limits our framework because if students do not update their plans, our framework interprets all input actions as being towards the same goal and plan. Similarly, if students use the planning support tool sparingly, then the goal and action labels might not be fully representative of the event sequences enacted in between planning support tool uses Further enhancements to the framework could be added by identifying when a plan has been completed through gameplay or a goal, so it is not singularly relying on students to update their goals and plans. Additionally, more work could be done to predict goal abandonment based on how long a goal or plan persists in the planning support tool interactions. Such improvements could alter the distribution in goal and plan labels and potentially help with recognition performance. Additionally, more work could be done to predict goal abandonment based on how long a goal or plan persists in the planning support tool interactions. Such improvements could alter the distribution in goal and plan labels and potentially help with recognition performance.

7 Conclusion

Goal setting and planning are key components of self-regulated learning. Adaptive learning environments show significant promise for adaptively scaffolding students' goal setting and planning processes, but they require computational models of student plan recognition to do so. This work presents a student plan recognition framework that leverages student goals and plans captured during interactions with a novel planning support tool in a game-based learning environment for middle school microbiology. Students' goals and plans were used to derive labels to formalize goal and plan recognition as multi-label classification tasks. Several machine learning techniques were evaluated to predict students' goal and plan labels based upon observations of their problem-solving actions in the game. In both tasks, we saw significant improvement on the majority baseline with most machine learning models. LSTMs showed particular promise in both the goal recognition and plan recognition tasks with respect to their ability to perform well across all classes.

The results indicate the potential of integrating student plan recognition models into real-time adaptive learning environments. Plan recognition models could be used to drive adaptive scaffolding in the form of open learner models of student goal setting and planning processes, or they could drive adaptive hints and prompts related to student SRL. Additionally, future work could investigate additional nuances of student goal setting and planning, which will contribute to more robust models because students can work towards multiple goals and plans at a time or abandon goals and plans without updating their planning support tool. Lastly, exploring additional sequential models and a multi-task learning approach to student goal recognition and plan recognition is a promising direction for future work.

Acknowledgements. This research was supported by funding from the National Science Foundation under grant DUE-1761178. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

1. Azevedo, R., & Wiedbusch, M. (in press). Theories of metacognition and pedagogy applied in AIED systems. In du Boulay (Ed.), Handbook of Artificial Intelligence in Education. The Netherlands: Springer.
2. Baikadi, A., Goth, J., Mitchell, C. M., Ha, E. Y., Mott, B. W., & Lester, J. C. (2011, October). Towards a computational model of narrative visualization. In Workshops at the Seventh Artificial Intelligence and Interactive Digital Entertainment Conference, (pp 2–9).
3. Blaylock, N., & Allen, J. (2003, August). Corpus-based, statistical goal recognition. In International Joint Conference on Artificial Intelligence, (Vol. 3, pp. 1303-1308).
4. Boekaerts, M., & Pekrun, R. (2015). Emotions and emotion regulation in academic settings. Handbook of educational psychology, 3, 76-90.
5. Bholowalia, P., & Kumar, A. (2014). EBK-means: A clustering technique based on elbow method and k-means in WSN. International Journal of Computer Applications, 105(9).

6. Conati, C., Gertner, A., & Vanlehn, K. (2002). Using Bayesian networks to manage uncertainty in student modeling. User modeling and user-adapted interaction, 12(4), 371-417.
7. Geib, C. W., & Steedman, M. (2007, January). On Natural Language Processing and Plan Recognition. In International Joint Conference on Artificial Intelligence, (Vol. 2007, pp. 1612-1617).
8. Ghallab, M., Nau, D., & Traverso, P. (2016). Automated planning and acting. Cambridge University Press.
9. Graesser, A., & McNamara, D. (2010). Self-regulated learning in learning environments with pedagogical agents that interact in natural language. Educational Psychologist, 45(4), 234-244.
10. Ha, E. Y., Rowe, J. P., Mott, B. W., & Lester, J. C. (2011, October). Goal recognition with Markov logic networks for player-adaptive games. In Seventh Artificial Intelligence and Interactive Digital Entertainment Conference, 32-39.
11. Liu, S. M., & Chen, J. H. (2015). A multi-label classification based approach for sentiment classification. Expert Systems with Applications, 42(3), 1083-1093.
12. Madjarov, G., Kocev, D., Gjorgjevikj, D., & Džeroski, S. (2012). An extensive experimental comparison of methods for multi-label learning. Pattern recognition, 45(9), 3084-3104.
13. Min, W., Ha, E. Y., Rowe, J., Mott, B., & Lester, J. (2014, September). Deep learning-based goal recognition in open-ended digital games. In Tenth Artificial Intelligence and Interactive Digital Entertainment Conference, 37-43.
14. Min, W., Mott, B., Rowe, J., Taylor, R., Wiebe, E., Boyer, K. E., & Lester, J. (2017, September). Multimodal goal recognition in open-world digital games. In Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference, (pp. 80–86).
15. Min, W., Mott, B. W., Rowe, J. P., Liu, B., & Lester, J. C. (2016, July). Player Goal Recognition in Open-World Digital Games with Long Short-Term Memory Networks. In Twenty-Fifth International Joint Conference on Artificial Intelligence, (pp. 2590-2596).
16. Pereira, R. B., Plastino, A., Zadrozny, B., & Merschmann, L. H. (2018). Correlation analysis of performance measures for multi-label classification. Information Processing & Management, 54(3), 359-369.
17. Schunk, D. H., & Greene, J. A. (2017). Historical, contemporary, and future perspectives on self-regulated learning and performance (pp. 1-15). Routledge.
18. Srinivasa-Desikan, B. (2018). Natural Language Processing and Computational Linguistics: A practical guide to text analysis with Python, Gensim, spaCy, and Keras. Packt Publishing Ltd.
19. Weintrop, D., & Wilensky, U. (2019). Transitioning from introductory block-based and text-based environments to professional programming languages in high school computer science classrooms. Computers & Education, 142, 103646.
20. Winne, P.H., & Azevedo, R. (2022). Metacognition. In K. Sawyer (Ed.), Cambridge handbook of the learning sciences (3rd, pp. 93-113). Cambridge, MA: Cambridge University Press.
21. Winne, P. H., & Hadwin, A. E. (1998). Studying as self-regulated learning (pp. 291-318). Routledge.
22. Winne, P. H., & Hadwin, A. F. (2008). The Weave of Motivation and Self-Regulated Learning. Motivation and Self-regulated Learning: Theory, Research, and Applications, 297.
23. Winne, P. H. (2018). Theorizing and researching levels of processing in self-regulated learning. British Journal of Educational Psychology, 88(1), 9-20.