

Enhancing Multimodal Goal Recognition in Open-World Games with Natural Language Player Reflections

Anisha Gupta¹, Dan Carpenter¹, Wookhee Min¹, Jonathan Rowe¹,
Roger Azevedo², James Lester¹

¹North Carolina State University, Raleigh, NC 27695, USA

²University of Central Florida, Orlando, FL 32816, USA

¹{agupta44, dcarpen2, wmin, jprowe, lester}@ncsu.edu, ²roger.azevedo@ucf.edu

Abstract

Open-world games promote engagement by offering players a high degree of autonomy to explore expansive game worlds. Player goal recognition has been widely explored for modeling player behavior in open-world games by dynamically recognizing players' goals using observations of in-game actions and locations. In educational open-world games, in-game reflection tools can help students reflect on their learning and plan their strategies for future gameplay. Data generated from students' written reflections can serve as a source of evidence for modeling player goals. We present a multimodal goal recognition approach that leverages players' natural language, written reflections along with game trace log features to predict player goals during gameplay. Results show that both the highest predictive performance and best early prediction performance are achieved by deep learning-based, multimodal goal recognition models that utilize both written reflection and gameplay features as input. These models outperform unimodal deep learning models as well as a random forest baseline. Multimodal goal recognition using natural language reflection data has significant potential to enhance goal recognition model performance and player modeling to support the creation of engaging and adaptive open-world digital games.

Introduction

Open-world game environments provide players with the opportunity to freely explore the game world and accomplish milestones in any order (Alexander, 2017; Aung et al., 2019). While this may facilitate a more personalized and immersive gameplay experience, the numerous possibilities of action sequences in such game environments introduces significant challenges in game design. For example, game designers need to ensure coherence in the flow of the game's narrative and support players' progress in the game, regardless of the player's idiosyncratic sequence of in-game actions (Min et al., 2014; Hooshyar et al., 2019). In educational settings, digital games contextualize learning and problem solving within engaging virtual environments that leverage the motivational characteristics of narrative such as

believable characters, compelling plots, and storyworld events (Riedl & Bulitko, 2013). Effective open-world educational games should support students to achieve their current plans or nudge them towards recommended milestones conducive to learning. However, a key challenge posed by open-world game-based learning environments is how to effectively provide such in-game support and maintain a coherent storyline, while students' plans and intentions, which are latent from the game's perspective, are dynamically changing during gameplay (Min et al., 2013).

To address this challenge, player goal recognition has been explored in previous work (Min et al., 2016a; Ha et al., 2012), where the system tries to predict the immediate next milestone that a player is likely to accomplish given their prior in-game actions. Goal recognition can inform effective intervention strategies in educational games (Ha et al., 2014). Accurate goal recognition models can help guide players towards the goals they are trying to achieve or inform narrative adaptations for players. Goal recognition models can also aid researchers and game designers in understanding which actions are correlated with different goals in the game, as well as challenges that players commonly face while working toward in-game goals.

The non-linear nature of open-world game environments usually translates to a set of goals which the players manually discover by exploring the game. This is particularly challenging since player actions can sometimes appear to be haphazard and undirected, which may also lead players to unintentionally reach certain milestones (i.e., predefined goals) without having a particular plan in mind. Prior work has explored the use of traditional machine learning algorithms with handcrafted features to address this (Ha et al., 2011; Mott et al., 2006). Recent work using deep learning methods has eliminated the need for feature engineering and allowed for more generalizable methods (Min et al., 2014; Min et al., 2016b; Min et al., 2017a).

In addition to low-level gameplay actions, other modalities such as eye gaze data (Min et al., 2017b) have been investigated to improve the performance of goal recognition models by leveraging an expanded set of predictive features that may also indicate player goals. Another source of evidence that may support improved goal recognition in open-world games is students’ natural language, written reflections during and after gameplay. Recent years have seen growing interest in the role of written reflection in game-based learning environments (Cloude et al., 2021). While students’ written reflections can offer insights about students’ strategies and future plans, it may also be a noisy data modality since the written reflections might not be clear about students’ plans or might not align with students’ goals achieved later in the game. This makes students’ written reflections challenging to interpret in open-world educational games with ill-defined goals. Moreover, to leverage these natural language responses as a modality for real-time goal recognition, they have to be automatically evaluated, introducing challenges of dealing with misspelled words and ungrammatical language.

In this paper, we present a long short-term memory (LSTM)-based multimodal goal recognition framework with decision-level fusion of multimodal data. We examine the effectiveness of incorporating students’ written reflections as evidence of planning and intent along with low-level gameplay action logs for the task of goal recognition. In the current context, we define goal recognition as the task of identifying students’ subgoals in an open-world educational game in which students seek to achieve a single overarching objective of solving an interactive science mystery. We measure the performance of multimodal goal recognition models using both accuracy and early prediction metrics (Min et al., 2016b). The performance of our deep learning models are evaluated against a random forest baseline using gameplay data collected from 156 middle school students interacting with an open-world educational game called CRYSTAL ISLAND. Our evaluation reports what combination of predictive features based on written reflections and gameplay log modalities achieve the highest predictive performance on multimodal player goal recognition.

Related Work

Player modeling is a computational task for modeling player cognition, behavior, and affective states to enable player-adaptive game experiences. Player modeling has a wide range of applications including procedural content generation, adaptive feedback, and game balancing (Yannakakis et al., 2013; Hooshyar et al., 2019). Dynamically recognizing players’ goals and intentions, a key player modeling task, in open-world games holds great promise for designing interactive narratives and providing adaptive support. To effectively model uncertainty, classic plan and goal recognition frameworks have leveraged probabilistic models (Charmiak

and Goldman, 1993; Pynadath and Wellman, 2000; Ha et al., 2011; Geib and Goldman, 2009). More recently, Pereira et al. (2020) explored plan recognition-as-planning techniques for designing heuristics that can quickly and effectively estimate player goals. Min et al. (2016a) presented action embedding-based LSTM networks for player goal recognition in open-world game environments and reported improved performance compared to the previous state-of-the-art techniques. While multimodal machine learning has been widely explored for different aspects of player modeling such as affect (Henderson et al., 2020) and interest (Emerson et al., 2020) modeling, there is limited work for multimodal goal and plan recognition. Min et al. investigated multimodal goal recognition, which achieves improved predictive performance compared to unimodal baselines by utilizing player gaze as predictive features (Min et al., 2017b).

In this work, we investigate the utility of reflection as evidence of player intent in open-world educational games. Reflection is both a backward and forward-looking process, where players reflect on their actions and prior learning and formulate future steps to achieve their learning goals (Cui et al., 2019). Prompting players to reflect on their learning and gameplay encourages them to actively think about their actions and make purposeful decisions that better align with their goals (Rogers, 2001). Written reflection prompts are often incorporated in educational game-based learning environments to encourage self-regulated learning in players (Villareale et al., 2020; Cloude et al., 2021). To date, there has been little work investigating written student reflection data as an input modality for dynamically recognizing players’ in-game goals. Reflection prompts encourage players to think about their strategy for successfully completing the game, making it a useful exercise that helps players ideate a plan that can be leveraged by goal recognition models. We hypothesize that this source of evidence is beneficial for goal recognition in open-world games in which players may not have well-defined subgoals and may discover subgoals while exploring the game.

Dataset

CRYSTAL ISLAND is an open-world educational game for middle school science (Figure 1). The objective of the game is to investigate a disease outbreak among a team of scientists on an island and present a diagnosis. In this work, we consider ten different key milestones in the game as subgoals accomplished by students: *speaking with the camp’s cook* to learn about the recently eaten food, *speaking with a sick patient*, *speaking with the lead scientist*, *testing an uncontaminated sample*, *testing a contaminated sample*, *submitting a diagnosis*, *speaking with the camp’s virus expert*, *speaking with the camp’s bacteria expert*, *speaking with the camp’s lab technician*, and *solving the mystery*, achievement of which collectively helps players complete the mission of the game. These milestone events, which we treat as

Reflection Response	Next Goals Achieved
<i>I am almost certain that I finally found the contaminant, bread/toast! It tested positive for viruses, so along with my other data, I'm almost certain the disease is the influenza virus!</i>	Submitted diagnosis Solved mystery
<i>What i learned is viruses and other things my next plan is to collect more information</i>	Speaking with the camp cook Speaking with the camp's lead scientist Speaking with the camp's bacteria expert Speaking with the camp's virus expert
<i>The sickness is passing from one person to the other</i>	Speaking with the camp cook Speaking with the camp's lead scientist

Table 1. Sample reflection responses and subsequent goals achieved.

in-game goals, were selected by considering the potential of goal recognition models to guide adaptive gameplay.

Students can accomplish the in-game goals in any order, and the way they navigate the environment may not be optimal for solving the mystery in the game. It is possible to successfully complete the game without accomplishing all the aforementioned goals. Completion of certain events, such as reading a microbiology book or obtaining a positive test result for the first time, may trigger a reflection prompt that asks students to write about what they have learned in the game thus far and what they plan to do moving forward. While some students' responses might be more specific than others, a majority of their reflections present a high-level idea of players' overall understanding of the game content and their strategy. The reflection responses may not refer to the in-game goals directly since players may not be explicitly aware of the existence of these in-game goals until they achieve them. In total, students receive up to five reflection prompts during gameplay, and each prompt is spaced out at least five minutes after the previous one. Responses to these prompts are logged as their written reflection responses.



Figure 1. CRYSTAL ISLAND game-based learning environment

In this work, we used a CRYSTAL ISLAND dataset containing logs from student interactions with the game during two studies conducted in 2018 and 2019, respectively. The combined dataset consists of game interaction logs of 156 students with ages between 13 and 14 years. A total of 729 written reflection responses with an average of approximately 20 words each were submitted by the students during gameplay. The students played the game until they successfully completed the game or 100 minutes of gameplay had elapsed. CRYSTAL ISLAND logged students' game interaction data as well as their written reflection responses during gameplay.

Goal Recognition Framework

Deep learning frameworks do not necessarily require manual feature engineering and thus are readily generalizable across different domains (Min et al., 2016b). In this work, we use an LSTM-based deep learning model that effectively models students' low-level action sequences and natural language-based reflection responses simultaneously to predict the next goal that the student will accomplish in the game. At each timestep of gameplay, goal recognition models are trained to predict the immediate next sub-goal that will be achieved by the student, so the prediction task is cast as a 10-class classification problem.

Data Preprocessing

Our feature set consists of actions, locations, previously achieved goals, and written reflection responses. There are 24 distinct locations on the map that the player can visit. We consider 9 action types: movement, editing a worksheet, accomplishing a goal, conversing with a non-player character (NPC), scanning an object to test for contamination, reading books and articles, attempting a reflection prompt, interacting with a poster, and submitting a worksheet. For each acti-

Model	Features	RF Accuracy	LSTM Accuracy	RF SCP	LSTM SCP	RF CR	LSTM CR
GM	Actions (A)	34.65	37.4	1.049	0.9736	16.36	29.03
GM	Locations (L)	32.31	38.69	1.1007	0.9359	17.37	38.7
GM	Goals (G)	24.43	50.67	1.0511	0.8446	15.2	34.35
GM	AL	33.04	39.81	1.1087	0.8949	13.48	44.61
GM	AG	19.51	51.75	1.0982	0.8321	10.95	39.94
GM	LG	18.55	54.24	1.1033	0.8132	10.24	45.13
GM	ALG	17.83	52.67	1.106	0.8349	8.93	42.02
RM	Reflections (R)	33.77	44.7	0.9955	0.8789	19.45	25.82
CM	RA	35.2	54.96	1.0383	0.8232	12.78	39.82
CM	RL	30.53	57.89	1.0397	0.8066	18.81	47.34
CM	RG	16.7	55.35	1.0973	0.8083	8.58	39.14
CM	RAL	30.74	57.6	1.0379	0.8202	19.48	47.04
CM	RAG	17.83	53.4	1.0959	0.8282	9.59	40.18
CM	RLG	16.65	54.87	1.1027	0.8185	9.02	44.31
CM	RALG	18.79	56	1.0977	0.8254	9.53	44.2

Table 2. Comparison of performance metrics of top-1 goal prediction models (RF: Random forest; SCP: standardized convergence point; CR: convergence rate; GM: Gameplay model; RM: Reflection model; CM: Combined model).

on that the player takes in the game, we construct a 43-dimensional gameplay feature vector comprising a feature-level fusion of (1) an action vector representing the action taken by the player at that timestep, (2) a location vector representing the location at which they perform the action, and (3) a binary vector representing the goals achieved by the player up to that point of gameplay. A sequence of such gameplay feature vectors for the player’s past 20 actions is considered as gameplay input to our LSTM models at each timestep of prediction.

Each natural language written reflection response of the student is represented as the average of their word-level ELMo embeddings (Peters et al., 2018). For our experiments, we use an ELMo model that’s pretrained on the 1 Million Word Benchmark, a dataset comprising approximately 800M tokens of news crawl data from WMT 2011. These ELMo embeddings are of 1,024 dimensions each and can have our goal recognition models overfit to the reflection data, given our dataset has a total of 729 unique written responses only. We thus applied Principal Component Analysis (PCA) on the ELMo embeddings to reduce their dimensionality to 32 dimensions, preserving the dimensions that best capture game context and the variation in the written

reflection response dataset. The written reflection responses are passed as input to our model by constructing a sequence of such ELMo embeddings for each response previously submitted by the player in the game.

Model Architecture

Our multimodal goal recognition framework consists of three separately trained submodels: a gameplay model, reflection model, and a combined model. The gameplay model was trained to predict the immediate next accomplished goal of the player at each timestep of gameplay using only low-level action logs (i.e., action performed, current player location, and previously accomplished goals) as input features. The reflection model was trained separately on previously submitted written reflection responses to predict the next goal to be accomplished by the player. Both gameplay and reflection models output probabilities for each of the 10 possible goals indicating how probable they are to be the next accomplished goal. Both of these probability predictions are then passed as input to a third model, freezing the model parameters for the gameplay and reflection models. This combined model performs a decision-level fusion of these

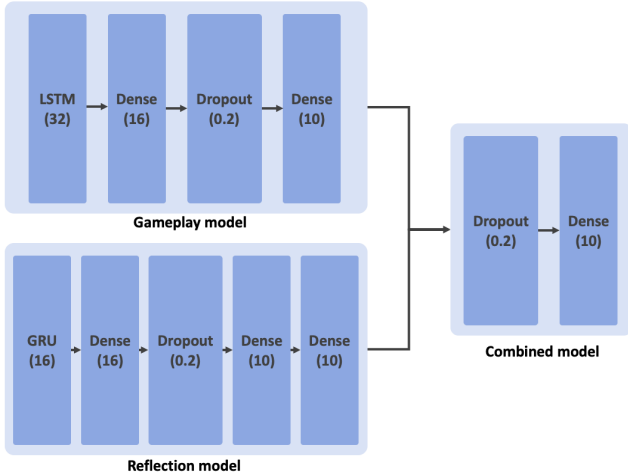


Figure 2. Model architecture for the gameplay, reflection and combined models. Numbers in parentheses denote the number of units in each layer, except in the case of the dropout layers where they denote dropout rates.

class probabilities to output a final prediction on the next goal to be accomplished (Figure 2). We use these final predictions from the combined model to evaluate our models. Since each goal can be accomplished once during gameplay, the goal prediction probabilities generated by the gameplay, reflection, and combined models are each post-processed to set the probabilities of previously accomplished goals to zero to improve goal recognition models’ predictive performance (Min et al., 2014). Each of these submodels was trained using the Adam optimizer with categorical cross entropy as the loss function.

Gameplay model. This model takes the past 20 actions, locations and previously accomplished goals as sequential input at each timestep of gameplay and passes it through an LSTM layer (32 hidden units, 0.1 dropout, 0.1 recurrent dropout, 0.01 L2 kernel/recurrent/bias regularization factor), a dense layer (16 hidden units, ReLU activation), a dropout layer (0.2), and a final dense layer (10 hidden units, sigmoid activation) that outputs the final probabilities for each goal.

Reflection model. This model takes the ELMo embeddings of the previous written reflection responses of the student as sequential input at each timestep of gameplay and passes it through a GRU layer (16 hidden units, 0.1 input-layer dropout rate, 0.1 recurrent-layer dropout rate, 0.01 L2 kernel/recurrent/bias regularization factor), dense layer (16 hidden units) and a dropout layer (0.2). The output is then passed through a dense layer (16 hidden units), and finally, another dense layer (10 hidden units, sigmoid activation) that outputs the final probabilities for each goal.

Combined model. This model accepts a concatenation of the prediction probability outputs of gameplay and reflection models at each timestep and passes this through a dropout layer (0.2) and a dense layer (10 hidden units, sigmoid

activation) that outputs the final probabilities for each goal. The model architecture was informed by previous work on student knowledge assessment in game-based learning environments (Gupta et al., 2021) and preliminary analysis based on the training set used in this work.

Evaluation

We performed a player-level nested cross-validation to evaluate our models with 5 inner folds and 3 outer folds of evaluation. In each inner fold, we performed a student-level 80-20 split on the training set to create training and validation sets. In each fold of the nested cross-validation for both the gameplay and reflection models, we explored different values of the number of LSTM hidden units (16, 32, and 64 hidden units) and dropout rates (0.2 and 0.5) as hyperparameters. For the combined model, we explored different dropout rates (0.1, 0.2, 0.5) as a hyperparameter. We evaluate our multimodal LSTM-based goal recognition approach compared to random forest-based goal recognition models, which serve as a competitive baseline. The sequential input used for the LSTM models were concatenated to create a contiguous feature vector as static input representations for the random forest models. The number of trees (50, 100, 200) in the random forest model was also optimized in a similar setup, with 5 inner folds and 3 outer folds of nested player-level cross-validation. For fair comparisons, we used the same split for the outer and inner folds between LSTMs and random forests. For all models, results are reported as the average of the 3 outer fold results.

Model performance is evaluated on accuracy, standardized convergence point, and convergence rate. The accuracy of a model is measured as the micro accuracy of correct goal predictions for each action across all students in the test set. We evaluate the early prediction performance of our models by computing the standardized convergence point and convergence rate of the predictions for each goal. The predictions are considered to converge at a timestep once the model consistently predicts the correct next goal from that time step onwards. Standardized convergence point (SCP) is computed as $\sum_{i=1}^n (k_i/m_i)/n$, where n is the total number of action sequences corresponding to the next goal labels, m_i is the number of actions in the i^{th} action sequence and k_i depends on the model’s convergence for that action sequence; if the model predictions converge, k_i is the action at which the model successfully converges on the correct next goal prediction, while for predictions that do not converge, k_i is computed as $(m_i + p)/m_i$, where p is a constant penalty parameter (Min et al., 2016b). (We set the penalty parameter to 1 in this work). Convergence rate (CR) is computed as the percentage of action sequences for which the model’s prediction converges on the correct goal right before the player actually accomplishes their next goal (Blaylock and Allen, 2003). A model has better predictive performance if it has higher accuracy and convergence rates and

Model	Features	RF Accuracy	LSTM Accuracy	RF SCP	LSTM SCP	RF CR	LSTM CR
GM	Actions (A)	34.65	37.4	1.049	0.9736	16.36	29.03
GM	Locations (L)	32.31	38.69	1.1007	0.9359	17.37	38.7
GM	Goals (G)	24.43	50.67	1.0511	0.8446	15.2	34.35
GM	AL	33.04	39.81	1.1087	0.8949	13.48	44.61
GM	AG	19.51	51.75	1.0982	0.8321	10.95	39.94
GM	LG	18.55	54.24	1.1033	0.8132	10.24	45.13
GM	ALG	17.83	52.67	1.106	0.8349	8.93	42.02
RM	Reflections (R)	33.77	44.7	0.9955	0.8789	19.45	25.82
CM	RA	35.2	54.96	1.0383	0.8232	12.78	39.82
CM	RL	30.53	57.89	1.0397	0.8066	18.81	47.34
CM	RG	16.7	55.35	1.0973	0.8083	8.58	39.14
CM	RAL	30.74	57.6	1.0379	0.8202	19.48	47.04
CM	RAG	17.83	53.4	1.0959	0.8282	9.59	40.18
CM	RLG	16.65	54.87	1.1027	0.8185	9.02	44.31
CM	RALG	18.79	56	1.0977	0.8254	9.53	44.2

Table 2. Comparison of performance metrics of top-1 goal prediction models (RF: Random forest; SCP: standardized convergence point; CR: convergence rate; GM: Gameplay model; RM: Reflection model; CM: Combined model).

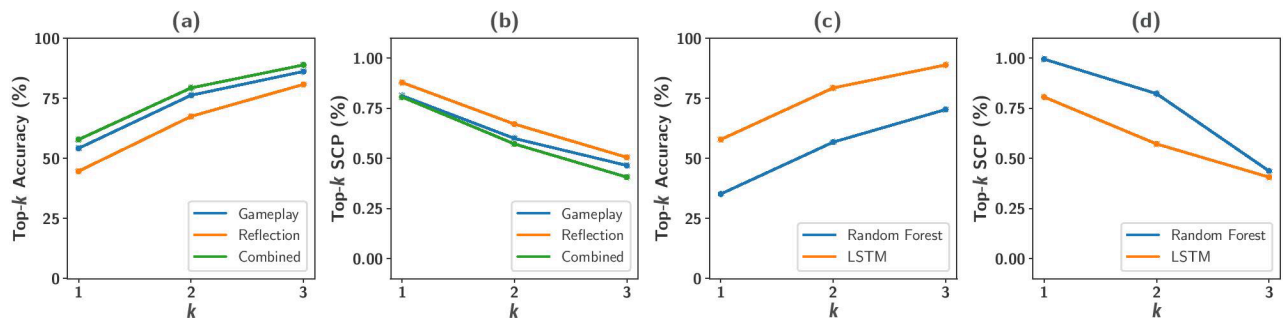


Figure 3. Comparison of predictive performance of (a,b) multimodal and unimodal LSTMs (c,d) LSTMs and random forests.

has lower standardized convergence points, which indicate better early prediction performance.

In the actual gameplay, players might have multiple goals in parallel at any given point. However, our dataset doesn't capture this aspect of interleaved or concurrent goals, introducing some uncertainty in the goal predictions. To account for this inherent uncertainty, we evaluate our models with top- k goal predictions, which has been widely used for image classification tasks. Identifying top- k goals of players can allow for more opportunities for player-adaptive gameplay, such as directing players towards a system-recommended goal selected from the set of top- k goals, or guiding

players to perform common action sequences for simultaneous progress towards multiple goals. In this work, we report the top- k performance of our models for k of 1 and 2. We evaluated model performances for top-3 predictions as well and observed 88.96% accuracy, 0.4069 SCP and 82.97% CR for our LSTM model with responses and gameplay features (full results have been omitted due to space constraints). We also computed the majority class baseline for our current dataset, yielding a top-1 prediction accuracy of 15% for the next goal prediction, which is significantly lower than the predictive performance of RF (35.2%) and LSTM (57.89%).

To evaluate the effectiveness of each feature in our

model, we performed ablation studies for both the deep learning and random forest models. For combined models (CMs) that leverage both gameplay and reflection features, we used gameplay and reflection submodels pre-trained with the game log features and reflection features, respectively, and used these submodels' predictions (i.e., freezing model parameters for the two submodels) to train and evaluate the CMs. The results for our ablation studies for top-1, top-2, and top-3 predictions are reported in Tables 2-3. The predictive accuracy and early prediction results of the best performing combined model are compared against those of the best performing gameplay submodel and reflection submodel in Figure 3(a,b) for top-1, top-2, and top-3 evaluations of goal recognition. Figure 3(c,d) illustrates a similar comparison between the best performing deep learning model and the best performing random forest model for top-1, top-2 and top-3 goal recognition evaluation.

Discussion

LSTM models outperformed the non-neural random forest baseline for all top- k goal recognition, in terms of both predictive accuracy and early prediction performance (Figure 3(c,d)). A possible explanation for this is that the objective of deep learning models is to extract high-level interpretations from low-level features, which aligns with goal recognition's objective of predicting the higher-level goals of a player in a game from their low-level game actions. From our ablation studies, we see that we obtain the best predictive performance using a deep learning model that is trained using location information and previous written reflection responses of players. This suggests that information about the players' recent movements on the map and what they said about their strategy in the game are most predictive of their next goals in the game.

From Figure 3(a,b), we can see that results using only gameplay features (trained on the gameplay submodel) are outperformed by combined models that utilize both written reflection and gameplay features (combined submodel) for top-1, top-2 and top-3 predictions, although written reflection features by themselves are not strong predictors. While the SCP and CR are similar for gameplay and combined submodels in our top-1 prediction results (Figure 3(a,b)), the combined submodel results outperform the early prediction performance of the gameplay model for top-2 and top-3 predictions. A possible explanation for this might be that the written reflection responses of players are not very specific about the immediate next goal but are more about a goal in the near future, which can be captured by top- k predictions. Moreover, players can also accomplish multiple goals since their last recorded written reflection response, given that 10 goals may be accomplished by the player while a maximum of only 5 reflection prompts can be encountered during gameplay. This could lead to better convergence for predicting the next k goals using written reflection responses with

gameplay features, as opposed to predicting the immediate next goal at each timestep.

We observe significant improvements in predictive accuracy (21.5%), SCP (23.42%), and CR (22.41%) in the best performing deep learning model when we compare top-1 and top-2 prediction results. This phenomenon can be explained by multiple concurrent goals players might have during the gameplay, where the model is not certain of what the immediate next goal achieved is. In future work, it would be interesting to perform error analysis and investigate which pairs of goals are often confused by the model and if there are common action patterns between these.

Conclusion

Goal recognition shows promise for dynamically identifying players' intentions and enabling player-adaptive games that provide coherent interactive narratives and adaptive support for individual players. While much of previous goal recognition work for digital games has utilized low-level in-game actions for predicting players' higher-level intentions during gameplay, other modalities of data available during gameplay may also serve as additional predictive features for enhancing player goal recognition. In this paper, we investigate the effectiveness of utilizing natural language reflection responses of players in an open-world educational game as an additional modality for top- k predictions of players' next goals. Results show that our LSTM-based multimodal goal recognition framework outperformed a random forest baseline. LSTM models utilizing a combination of written reflection responses and gameplay data as features outperformed unimodal models utilizing gameplay data only, both in terms of predictive accuracy as well as early prediction performance. These results demonstrate that reflection-based natural language data is a promising modality that can be included to improve the predictive performance of goal recognition frameworks. Written reflection responses could serve as a first-hand account of students' intentions in the game, complementing evidence of their past actions taken in the game. Directions for future work include investigating additional modalities for goal recognition, such as speech and body movement, that can be indicative of player engagement. It will also be important to incorporate multimodal goal recognition models into run-time game environments and investigate the fidelity of such models with respect to predictive accuracy and early prediction for driving interventions that enhance players' gameplay experiences.

Acknowledgements

This research was supported by funding from the National Science Foundation under Grant DRL-1661202. 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

- Aung, M.; Demediuk, S.; Sun, Y.; Tu, Y.; Ang, Y.; Nekkanti, S.; Raghav, S.; Klabjan, D.; Sifa, R.; and Drachen, A. 2019, August. The trails of Just Cause 2: spatio-temporal player profiling in open-world games. In Proceedings of the 14th International Conference on the Foundations of Digital Games: 1-11.
- Alexander, R.; and Martens, C. 2017, August. Deriving quests from open world mechanics. In Proceedings of the 12th International Conference on the Foundations of Digital Games: 1-7.
- Blaylock, N.; and Allen, J. 2003, August. Corpus-based, statistical goal recognition. In IJCAI (Vol. 3): 1303-1308.
- Cloude, E.B.; Carpenter, D.; Dever, D.A.; Azevedo, R.; and Lester, J. 2021. Game-Based Learning Analytics for Supporting Adolescents' Reflection. *Journal of Learning Analytics*, 8(2): 51-72.
- Cui, Y.; Wise, A.F.; and Allen, K.L. 2019. Developing reflection analytics for health professions education: A multi-dimensional framework to align critical concepts with data features. *Computers in Human Behavior*, 100: 305-324.
- Emerson, A.; Cloude, E.B.; Azevedo, R.; and Lester, J. 2020. Multimodal learning analytics for game-based learning. *British Journal of Educational Technology*, 51(5):1505-1526.
- Geib, C.W.; and Goldman, R.P. 2009. A probabilistic plan recognition algorithm based on plan tree grammars. *Artificial Intelligence*, 173(11): 1101-1132.
- Gupta, A.; Carpenter, D.; Min, W.; Rowe, J. P.; Azevedo, R.; and Lester, J. C. 2021. Multimodal multi-task stealth assessment for reflection-enriched game-based learning. In MAIED@ AIED: 93-102.
- Ha, E.Y.; Rowe, J.P.; Mott, B.W.; and 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.
- Ha, E.Y.; Rowe, J.P.; Mott, B.W.; and Lester, J.C. 2012, July. Goal recognition with Markov logic networks for player-adaptive games. In Twenty-Sixth AAAI Conference on Artificial Intelligence: 2113-2119.
- Ha, E. Y.; Rowe, J. P.; Mott, B. W.; Lester, J. C. 2014. Recognizing player goals in open-ended digital games with markov logic networks. *Plan, Activity and Intent Recognition: Theory and Practice*: 289-311.
- Henderson, N.; Rowe, J.; Paquette, L.; Baker, R.S.; and Lester, J. 2020, July. Improving affect detection in game-based learning with multimodal data fusion. In International Conference on Artificial Intelligence in Education. Springer, Cham: 228-239.
- Hooshyar, D.; Yousefi, M.; and Lim, H. 2019. A systematic review of data-driven approaches in player modeling of educational games. *Artificial Intelligence Review*, 52(3): 1997-2017.
- Min, W.; Ha, E.Y.; Rowe, J.; Mott, B.; and 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.
- Min, W.; Mott, B.W.; Rowe, J.P.; Liu, B.; and Lester, J.C. 2016a, July. Player Goal Recognition in Open-World Digital Games with Long Short-Term Memory Networks. In IJCAI: 2590-2596.
- Min, W.; Baikadi, A.; Mott, B.; Rowe, J.; Liu, B.; Ha, E.Y.; and Lester, J. 2016b, September. A generalized multidimensional evaluation framework for player goal recognition. In Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference: 197-203.
- Min, W.; Mott, B.; Rowe, J.; and Lester, J. 2017a, March. Deep LSTM-based goal recognition models for open-world digital games. In Workshop on Plan, Activity, and Intent Recognition at the Thirty-First AAAI Conference on Artificial Intelligence: 2590-2596.
- Min, W.; Mott, B.; Rowe, J.; Taylor, R.; Wiebe, E.; Boyer, K.E.; and Lester, J. 2017b, September. Multimodal goal recognition in open-world digital games. In Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference: 80-86.
- Min, W.; Rowe, J.P.; Mott, B.W.; and Lester, J.C. 2013, July. Personalizing embedded assessment sequences in narrative-centered learning environments: a collaborative filtering approach. In International Conference on Artificial Intelligence in Education. Springer, Berlin, Heidelberg: 369-378.
- Mott, B.; Lee, S.; and Lester, J. 2006, July. Probabilistic goal recognition in interactive narrative environments. In AAAI: 187-192.
- Pereira, R.F.; Oren, N.; and Meneguzzi, F. 2020. Landmark-based approaches for goal recognition as planning. *Artificial Intelligence*, 279, 103217: 1-32.
- Peters, M.E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume I (Long Papers). New Orleans, Louisiana. Association for Computational Linguistics: 2227-2237.
- Pynadath, D.V.; and Wellman, M.P. 2013. Probabilistic state-dependent grammars for plan recognition. *arXiv preprint arXiv:1301.3888*: 507-514.
- Riedl, M.O.; and Bulitko, V. 2013. Interactive narrative: An intelligent systems approach. *AI Magazine*, 34(1): 67-67.
- Rogers, R.R. 2001. Reflection in higher education: A concept analysis. *Innovative higher education*, 26(1): 37-57.
- Villareale, J.; F. Biemer, C.; Seif El-Nasr, M.; and Zhu, J. 2020, September. Reflection in game-based learning: a survey of programming games. In International Conference on the Foundations of Digital Games: 1-9.
- Yannakakis, G.N.; Spronck, P.; Loiacono, D.; and André, E. 2013. Player Modeling. In S. M. Lucas, M. Mateas, M. Preuss, P. Spronck, & J. Togelius (Eds.), *Artificial and computational intelligence in games*, DFU, 6. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik: 45-59.