Towards a Model-Learning Approach to Interactive Narrative Intelligence for Opportunistic Storytelling

Emmett Tomai and Luis Lopez

Department of Computer Science
University of Texas Rio Grande Valley, Edinburg, TX, USA
emmett.tomai@utrgv.edu, luis.a.lopez02@utrgv.edu

Abstract. Opportunistic storytelling is an approach to interactive narrative where the AI challenge is to tell a story about what the player is doing, with game play providing the ordinary activity that underlies story events. In this preliminary work, we motivate the need for event prediction within this framework, and describe a machine learning approach to the problem. We report results showing how different feature models can be learned and compared in this context, towards automating model selection.

Keywords: Intelligent Narrative Technologies, Artificial Intelligence, Machine Learning, Game Design

1 Introduction

One vision for interactive narrative seeks to combine the interactivity of video games with the immersion of narrative. However, the strength of game play as a model of interaction includes the freedom to master consistent systems through repetitive exploration [1]. This is at odds with the carefully curated web of meaningful events found in traditional narrative. But this conflict suggests a promising approach: to embrace game play as a stream of ordinary activity – the mundane events that are left out of stories – and recasts the AI challenge as telling a story, that meets authorial goals, about what the player is doing. In this *opportunistic storytelling* approach, the player has a clear role as game player, and the AI is constrained to the rules of the game simulation. We believe this can mitigate the conflict between player freedom and authorial control [2], and enable the AI to understand the domain in which the story is being told. We describe a game and narrative AI to implement this strategy, discuss the importance of event prediction, and report preliminary results in that direction.

2 Opportunistic Storytelling: Game and Narrative AI

Prior work in interactive narrative has used consistent simulation of agents [cf. 3], of story-specific behaviors [4] and of social interactions [5] to create interactive experiences with different degrees of emergent vs. directed story. The *Marlinspike* system

F. Nack and A.S. Gordon (Eds.): ICIDS 2016, LNCS 10045, pp. 428–432, 2016.
DOI: 10.1007/978-3-319-48279-8_42

[6] uses *Inform7* [7] game mechanics, and opportunistically selects scenes that make prior player actions significant to the story. Our approach is similar, but we have chosen graphical, open-world game play simulation, to lessen the problem of players feeling limited or led along. We are also combining the simple robustness of low-level believable agents with high-level story direction, as done in *IN-TALE* [8]. In our case, believable refers to meeting player expectations for the game mechanics, and the high-level direction is opportunistic guiding towards desired story states.

We have created a procedurally generated, infinite world survival game, similar to Klei Entertainment's *Don't Starve* [9] and others. The player follows an upward cycle of collecting resources to craft and build in order to survive the environment. This provides everyday goals, actions and experiences. Autonomous agents in the game have the same capabilities as the player, and choose behaviors that are consistent with surviving. The goal is not for agents to act just like players, but to have behaviors that make sense – that players can explore and master. Within this world, stories emerge from simple interactions between the player and agents over gathering resources and avoiding threats. The job of the opportunistic storytelling system is to maximize the probability of those stories being experienced. It can only intervene on unseen entities and internal goals and attributes, similar to the *alibi generation* problem [10]. To do this, it must be able to predict outcome probabilities with and without intervention.

Beyond integrating in a graphical real-time game world, the novel aspect of this work is using agent simulation to build a predictive model. Of course the game itself is already a perfect model, but simulating every possible future to find the narratively interesting ones is not tractable on-line. By running the game off-line, the AI can train models for estimating features and relationships among events. Typical machine learning work in video games [cf. 11] has focused primarily on models for optimal action selection. We believe that narrative intelligence will require learning and composing many heterogeneous models under higher level control abstractions. At this time, we are focused on automated learning of event prediction.

3 Learning Models for Event Prediction

This preliminary experiment explores the impact of different models on event prediction. The data is a set of gameplay traces where an agent is assigned to complete gathering tasks in a randomly generated environment. Anywhere from zero to dozens of gathering nodes and enemies could be in sight of the agent at any time. The system applies a feature model to the traces to create standard classification data, then trains and tests its ability to predict the next agent action: Attack, Gather, Flee or Explore.

Three feature models were evaluated. The *individual* model transforms every entity within visual range of the agent into a feature vector including its distance from the agent, health, combat power and the resources that can be gathered from it. Agent attributes such as risk aversion (randomly assigned) are also included. The *paths* model includes a numerical estimate of how clear the path from the agent to the entity is – a higher-level abstraction that should be relevant to predicting what will happen. The *regions* model clusters entities into spatial groups and aggregates their individual attributes. Again, this higher-level abstraction would be relevant to a human evaluation of the agent's next move. For each model, the experiment was run with 25, 50, 75 and 100 traces, randomly selected and split in half for training and testing. Each condition was averaged over 1000 runs. Both Naïve Bayes and Random Forest classifiers were used, with no notable difference. The Random Forest results are reported.

Results for Attack and Flee events were very poor across the board (precision and recall < 0.1), and degraded with more training samples. Precision and recall for predicting Explore and Gather events are shown in Figures 1-4. The horizontal axis shows number of game play traces used. The simple individual model generally improved with additional samples, but was erratic in its recall for Explore events. The paths model was the most consistent performer across both tasks, suggesting that it is an appropriate level of abstraction for these tasks. The more complex regions model showed no advantage, and significantly underperformed in recall for Gather events.

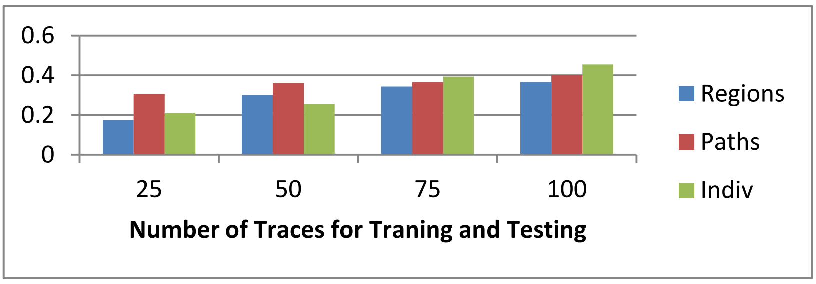

Fig. 1. Precision for predicting Explore events with three different feature models

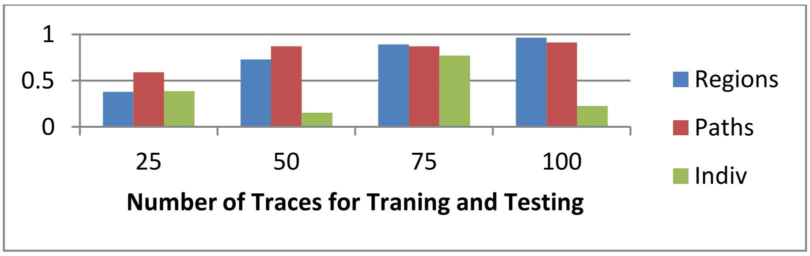

Fig. 2. Recall for predicting Explore events with three different feature models

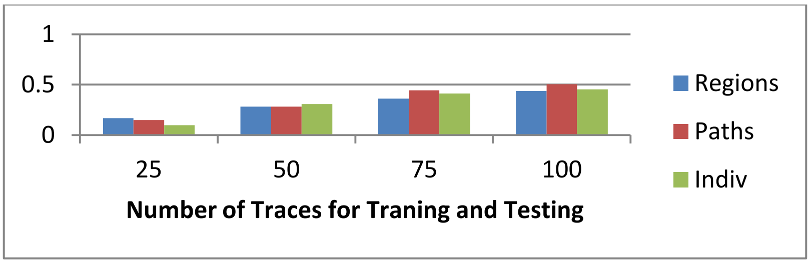

Fig. 3. Precision for predicting Gather events with three different feature models

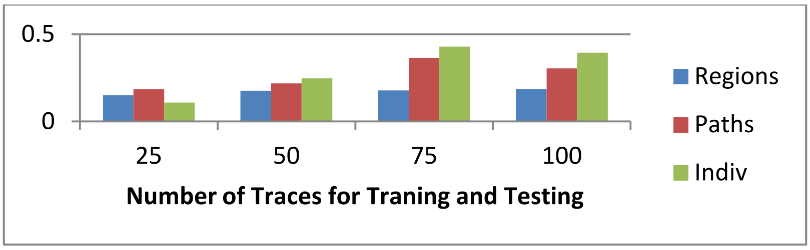

Fig. 4. Recall for predicting Gather events with three different feature models

Having established the experimental pipeline, we are encouraged by the ability to get coherent results from simple, fast learning models with under 50 training samples. The easy interchange of features and learning models is promising for the future step of automated model selection and composition.

Acknowledgements

This project was supported in part by the UTRGV Center of Excellence in STEM Education Undergraduate Research Program, Office of Naval Research contract W911NF-11-1-0150.

References

1. Koster, R. 2010. *A Theory of Fun for Game Design*. O'Reilly Media, Inc.
2. Young, R.M., 2000. Creating Interactive Narrative Structures: The Potential for AI Approaches. In *Proceedings of the AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment*, AAAI Press.
3. Louchart, S., & Aylett, R. 2004. Narrative theory and emergent interactive narrative. *Intl Journal of Continuing Engineering Education and Life Long Learning*, 14(6), 506-518.
4. Mateas, M. and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. Game Developers Conference, 2003.
5. McCoy, J., Treanor, M., Samuel, B., Reed, A., Mateas, M. & Wardrip-Fruin, Noah. "Prom Week: Designing past the game/story dilemma." *Proceedings of the 8th International Conference on Foundations of Digital Games*, Chania, Crete, Greece. May 14-17, 2013.
6. Tomaszewski, Z., and Binsted, K. 2009. "Demeter: An Implementation of the Marlinspike Interactive Drama System". In AAAI Spring Symposium: Intelligent Narrative Technologies II (p. 133).
7. *http://inform7.com*
8. Riedl, M. O., & Stern, A. 2006. Believable agents and intelligent story adaptation for interactive storytelling. In *Technologies for Interactive Digital Storytelling and Entertainment* (pp. 1-12). Springer Berlin Heidelberg.
9. *https://www.kleientertainment.com/games/dont-starve*
10. Sunshine-Hill, B. and Badler, N.I., 2010. Perceptually realistic behavior through alibi generation. In *Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.
11. Galway, L., Charles, D. and Black, M. (2008). Machine learning in digital games: a survey. *Artificial Intelligence Review*. Volume 29, Number 2, 123-161.