Representation Learning in Game Provenance Graphs

Anonymous Author(s)
Affiliation
Address
email

1 Introduction

Capturing game session information is usually done using replay files, which have the disadvantage of being proprietary binary files. In [Weber and Mateas, 2009, Robertson and Watson, 2014] we have examples of information capturing works from such replay files with *specific* solutions for the source games. In this context, data provenance is suggested as an alternative for representing game session information. Provenance refers to the documented history of an object's life cycle[Group et al., 2005]. This documented history is structured in the form of provenance graphs. A provenance graph is an oriented, acyclic, causal graph enriched with annotations. Such a graph aims at mapping the flow of actions taken during an object's life cycle and the produced results. There are three types of nodes in a provenance graph: Artifacts, Processes, and Agents. In the context of games, Artifacts represent inanimate objects (such as weapons and items), Processes represent actions (such as running, jumping, attacking), and Agents represent players, enemies, and non-playable characters.

Game provenance was first conceived as the conceptual framework Ping and later implemented into the PinGU (Provenance in Games for Unity) plugin [Kohwalter et al., 2017, 2018]. The PingU plugin generates an XML file containing a provenance graph. Each node has a type according to the PinG framework and the game developer-defined attribute data.

Initially developed as a game analytics tool, since viewing the documented story through the enriched annotated graph could help to identify game design flaws and enhance the player experience, game provenance is also a strong candidate for machine learning applications. The large amount of information associated with each graph element (nodes and edges), which in turn is generated in abundance over a game session, is the key feature of this data structure for machine learning.

On the other hand, recent advances in the field of deep learning have been reflected in machine learning techniques applied to digital games relying on computer vision and reinforcement learning, and featuring high performance and generality[Volodymyr et al., 2013]. However, these techniques require a massive computational effort, since their training and validation steps are performed using deep networks fed by video and screen captures. Game provenance graphs, in turn, hold a wealth of information about the elements of a game and their relationships (interactions) that are not directly represented in videos or screenshots. Therefore, we believe that provenance graphs are useful for various machine learning tasks, not just for reinforcement learning and intelligent agent building.

2 PingUMiL: Experiments and Results

In the context of machine learning in graphs, we have devised a representation learning-based approach for dealing with game provenance graphs in the PingUMiL framework[1]. The idea behind this approach is to rely on vector representations to learn a mapping that encodes nodes, or (sub)graphs,

[1] `https://github.com/sidneyaraujomelo/PingUMiL`

as points in a vector space so that geometric relationships within that learned space reflect the structure in the original graph[Hamilton et al., 2017].

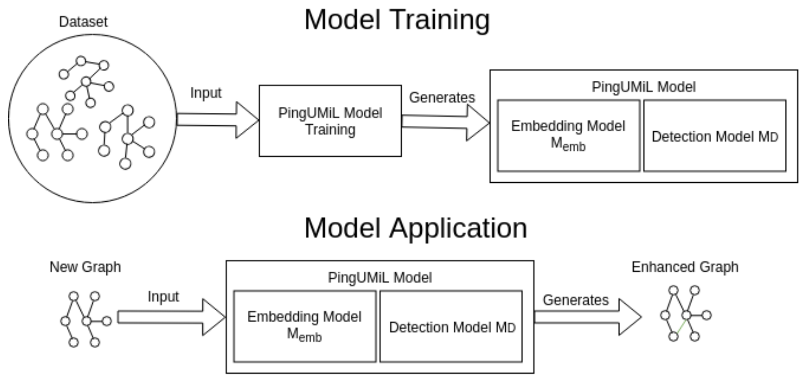

Figure 1: Overview of the PingUMiL framework for a link prediction task.

We applied a convolution-based representation learning approach, which determines the embedding of the node according to the attributes of its neighborhood. These methods are also called neighborhood aggregation methods. As the results of our investigation, we implemented a framework for machine learning tasks on game provenance graphs and conducted experiments on two racing games prototypes and a multiplayer airplane battle arena game[2]. In the racing games, we performed a link prediction task and achieved high-performance values for the metrics of precision and recall (above 70%). Besides, we also investigated the effect of learning representations using graphs from both racing games, which led to enhancement on recall metrics for specific types of edges. In the airplane battle arena game, preliminary results on a node classification task across graphs from multiple sessions show high micro f1-score (above 90%) on determining the player responsible for each recorded action.

3 Current Research

From those previous findings, we noticed that one of the significant limitations for combining provenance graphs and machine learning is the lack of support for handling heterogeneity of nodes and edges. In the context of digital games, it is natural that nodes representing different entities (an item, a player, or an enemy) have completely different attribute vectors. For the aggregation approach mentioned previously, only [Veličković et al., 2018] supports heterogeneous nodes as input, generating vector representations oriented to the learning task performed. On the other hand, works such as [Weston et al., 2011] and [Chang et al., 2015] present approaches for learning representations of entirely distinct datasets (image and text, for example) by mapping their data points into the same vector space. For both approaches, learned representations depend on local neighborhood and explicitly determined relationships (instantiated as edges); in the game provenance graph, however, learned representations should reflect the role of a node in spite of its neighborhood. We also believe that successfully solving this challenge will facilitate the transfer learning across multiple games, which is necessary for data reuse since obtaining provenance data is not trivial and the amount of data depends on the number of played sessions.

The general goal of the current research is to develop a *machine learning-based framework for heterogeneous graphs based problems, such as game provenance graphs*. To fully achieve this goal, we intend to develop **graph network architectures**[Battaglia et al., 2018] capable of **handling the intrinsic heterogeneity of nodes and edges of provenance graphs and tackle prediction and classification tasks** in edge, node and subgraph levels. To that, we are investigating not only graph-based deep learning approaches, but also the more general class of methods that combine neural networks and logical languages to learning and reasoning from relational data [Kazemi and Poole, 2018, Manhaeve et al., 2018]. We aimed at designing graph-based machine learning techniques that are capable of adapting state-of-the-art performance models, such as relational deep reinforcement learning [Zambaldi et al., 2018] for heterogeneous graphs, using, for example, attention mechanisms.

[2]A paper with the framework and its results is to be released in Elsevier's Entertainment Computing journal.

References

P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.

S. Chang, W. Han, J. Tang, G.-J. Qi, C. C. Aggarwal, and T. S. Huang. Heterogeneous network embedding via deep architectures. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 119–128. ACM, 2015.

P. W. Group et al. Data dictionary for preservation metadata: final report of the premis working group. *OCLC Online Computer Library Center & Research Libraries Group, Dublin, Ohio, USA, Final report*, 2005.

W. L. Hamilton, R. Ying, and J. Leskovec. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017.

S. M. Kazemi and D. Poole. Relnn: A deep neural model for relational learning. In *Proc. of the 32nd AAAI Conference on Artificial Intelligence*, pages 6367–6375, 2018.

T. C. Kohwalter, L. G. P. Murta, and E. W. G. Clua. Capturing game telemetry with provenance. In *2017 16th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*, pages 66–75. IEEE, 2017.

T. C. Kohwalter, F. M. de Azeredo Figueira, E. A. de Lima Serdeiro, J. R. da S. Junior, L. G. P. Murta, and E. W. G. Clua. Understanding game sessions through provenance. *Entertainment Computing*, 27:110–127, 2018.

R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, and L. De Raedt. Deepproblog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems*, pages 3749–3759, 2018.

G. Robertson and I. Watson. An improved dataset and extraction process for starcraft ai. In *The Twenty-Seventh International Flairs Conference*, 2014.

P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. *Proceedings of the International Conference on Learning Representations (ICLR).*, 2018.

M. Volodymyr, K. Kavukcuoglu, D. Silver, A. Graves, and I. Antonoglou. Playing atari with deep reinforcement learning. In *NIPS Deep Learning Workshop*, 2013.

B. G. Weber and M. Mateas. A data mining approach to strategy prediction. In *2009 IEEE Symposium on Computational Intelligence and Games*, pages 140–147. IEEE, 2009.

J. Weston, S. Bengio, and N. Usunier. Wsabie: Scaling up to large vocabulary image annotation. In *Proc. of the 22nd Int. Joint Conference on Artificial Intelligence*, pages 2764–2770, 2011.

V. Zambaldi, D. Raposo, A. Santoro, V. Bapst, Y. Li, I. Babuschkin, K. Tuyls, D. Reichert, T. Lillicrap, E. Lockhart, et al. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830*, 2018.