

Ev-IDID: Enhancing Solutions to Interactive Dynamic Influence Diagrams through Evolutionary Algorithms

Demonstration Track

Biyang Ma

School of Computer Science and Engineering
Minnan Normal University, China

Yifeng Zeng

Department of Computer and Information Sciences
Northumbria University, UK

Yinghui Pan

National Engineering Laboratory for Big Data System
Computing Technology
Shenzhen University, China

Zhong Ming

College of Computer Science and Software Engineering
Shenzhen University, China
Email:zming@szu.edu.cn

ABSTRACT

Interactive dynamic influence diagrams (I-DIDs) are a general framework for multiagent sequential decision making under uncertainty. Due to the model complexity, a significant amount of research has been invested into solving the model through various types of either exact or approximate algorithms. However, there is no tool that allows users to specify the algorithm parameters and visualise the model solutions. In this demo, we develop an interactive I-DID system that implements most of the state-of-art I-DID algorithms and develops a new type of algorithms based on evolutionary computation. In particular, we propose a multi-population genetic algorithm for solving the I-DID models and automate the generation of behavioural models in the solutions. This demo will facilitate the I-DID research development and practical applications, and elicit a new wave of I-DID solutions based on evolutionary algorithms.

KEYWORDS

Interactive Dynamic Influence Diagrams; Evolutionary Algorithms; Interactive Decision Systems

ACM Reference Format:

Biyang Ma, Yinghui Pan, Yifeng Zeng, and Zhong Ming. 2022. Ev-IDID: Enhancing Solutions to Interactive Dynamic Influence Diagrams through Evolutionary Algorithms: Demonstration Track. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 3 pages.

1 INTRODUCTION

From the perspective of individual agents, interactive dynamic influence diagrams (I-DIDs) solve multiagent sequential decision making problems under uncertainty [5, 12]. Compared to its counterpart of interactive partially observable Markov decision process (I-POMDP) [7], the I-DID model exhibits computational advantages with a graphical representation. Other graphical representations for modelling multiagent decision making include multiagent influence diagrams (MAIDs) used to compute Nash equilibrium strategies [9] and networks of influence diagrams (NID) for recursively modelling other agents [6]. Both MAID and NID formalisms focus on a static,

single-shot interaction. In contrast, I-DIDs offer solutions over extended time interactions, where agents act and update their beliefs over others' models which are themselves dynamic.

An I-DID model mainly contains two components: one is a dynamic influence diagram modelling a subject agent's decision making process while the other is any decision or behavioural model that can predict what other agents act simultaneously in their common environment. Since the true models of the other agents are not known by the subject agent, the subject agent needs to hypothesize a large number of candidate models (in theory, the number is infinite) for the other agents so that it can predict their behaviours by solving the candidate models. Consequently, the I-DID complexity lies in solving the candidate models that grow exponentially with the number of agents' observations over the planning horizons.

A lot of I-DID algorithms, including exact and approximate ones, have been developed to solve I-DIDs by compressing the models space of other agents [3, 4, 13, 14]. They perform differently in terms of their parameter specification, e.g. the number of candidate models ascribed to other agents, the model similarity measurements, the planning horizon and so on. More importantly, what matters to a subject agent is an optimal policy (often represented by a policy tree) that is obtained from solving the I-DID model. It becomes useful if the solutions (including predicted behaviours of other agents and optimal policies for a subject agent) are to be visualized by users who can understand the agents' decisions in the decision evaluation stage. Hence we need an interactive tool to specify the algorithm parameters, evaluate the algorithms and visualize decision results when I-DIDs are used to solve a real-world problem.

In addition, the recent I-DID research still depends on hand-crafted decision models for other agents that are to be integrated into the I-DID models. This leads to a couple of limitations in the I-DID development: (a) the modelling process is tedious and subject to inputs of domain experts; and (b) behaviours (solutions from solving other agents' models) tend to be monotonic therefore limiting a subject agent's capability in exploring unknown behaviours of other agents. Hence the challenge is about automating and diversifying other agents' behaviours, which is still under exploration in the I-DID research. In this work, we propose an evolutionary algorithm [8] based framework to generate behaviours of other agents in the I-DID framework. The new approach generates new

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: The framework uses a MPGA method to generate new behaviours for other agents. It includes coding individual behaviours, randomising the behaviours and evaluating the behaviours in the evolutionary computation.

behaviours by exploring known models of other agents and randomising the behaviours through evolutionary operators. Our demo will glean the previous work on the I-DID algorithms and develop a new framework to enhance the I-DID solutions.

2 GENERATING NEW BEHAVIOURS FOR OTHER AGENTS

Figure 1 shows a general framework that uses a multi-population genetic algorithm (MPGA) to generate new behaviours for other agents. It includes several operators that evolve behaviours of other agents over iterations. Individuals in one population encode possible behaviours of other agents each of which is represented through a policy tree. Within the population, the individuals conduct conventional genetic operators [10], e.g. selection, mutation and crossover, therefore resulting in a set of new individuals. The difficulty lies in computing individual fitness values in the evolution. We evaluate individuals by computing expected rewards of the corresponding behaviours in other agents’ decision models. Since initial beliefs are unknown in the decision models, we use a grid search in the belief space and find the closest beliefs for individual behaviours.

In addition, we apply an emigration operator in every iteration. The emigration exchanges a number of individuals among multiple populations, which intends to improve the population diversity in the individual evolution. We also implement a number of sophisticated emigration operators, e.g. the selections based on a *sigmoid* function, besides a random selection. We shall note that we do not intend to build *decision models* for other agents, but provide predicted behaviours from a subject agent’s viewpoint.

Once the evolution reaches a pre-defined number of iterations, we select K individuals from one population that have the largest coverage of behaviours for other agents. The resulting behaviours will provide extra inputs to the I-DID models where a subject agent needs to predict other agents’ behaviours in the modelling process.

3 AN INTERACTIVE I-DID SYSTEM

Figure 2 presents an interactive system, namely Ev-IDID, that allows users to build and test I-DID models in solving a sequential multiagent decision making problem. The system streamlines all the processes from generating behaviours for other agents by solving their decision models (e.g. dynamic influence diagrams), building




Figure 2: Ev-IDID: An interactive system contains a number of components, including the state-of-art I-DID algorithms and newly developed GA-based solutions, for running, evaluating and visualising I-DID solutions.

I-DID models for a subject agent, to evaluating and visualising the resulting policies for other agents and the subject agent. We provide the .exe file to run the system ^{1 2}. The system is also integrated with the SMILE engine - a well-known Bayesian networks tool, and all the decision models can run through the GeNIe application ³.

Once providing necessary inputs to the system, e.g. the planning horizon, initial beliefs of agents and prior knowledge about a problem domain, users can select a set of the existing I-DID algorithms to solve the I-DID model while either solving DID models or generating new behaviours for other agents. One noticeable feature in this system is to provide a comprehensive tool to reasoning with other agents’ behaviours through policy trees in addition to decision models. This provides users ample channels on exploring research on modelling other agents. For example, users can learn behaviours from agents’ interaction data [2, 11]. Users can also develop different evolutionary algorithms, similar to a MPGA method in Fig. 1, to generate other agents’ behaviours.

4 CONCLUSION

Ev-IDID is the first system that provides a comprehensive set of I-DID algorithms and develops a new way of modelling other agents in the application. The implementation becomes a uniform platform to research I-DIDs and facilitate their applications in various domains [1]. We will continue to add more functions into this system, and provide users with I-DIDs models for a set of problem domains so that the users can experiment with the models and gain experiences of using the Ev-IDID system.

ACKNOWLEDGMENTS

Professor Yifeng Zeng received the support from the EPSRC New Investigator Award and Biyang is partially supported by the EPSRC project (Grant No. EP/S011609/1). This work is supported in part by the National Natural Science Foundation of China (Grants No.61772442, 61836005 and 62176225).

¹<https://github.com/lamingic/Ev-IDID>

²<https://pan.baidu.com/s/1QpIgU1s6BiA1VjA9EITLbA?pwd=s2N6>

³<https://www.bayesfusion.com>

REFERENCES

- [1] Muthukumar Chandrasekaran, Prashant Doshi, Yifeng Zeng, and Yingke Chen. 2014. Team behavior in interactive dynamic influence diagrams with applications to ad hoc teams. In *International conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1559–1560.
- [2] Ross Conroy, Yifeng Zeng, Marc Cavazza, and Yingke Chen. 2015. Learning Behaviors in Agents Systems with Interactive Dynamic Influence Diagrams. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence IJCAI*. 39–45.
- [3] R. Conroy, Y. Zeng, M. Cavazza, J. Tang, and Y. Pan. 2016. A Value Equivalence Approach for Solving Interactive Dynamic Influence Diagrams. In *Proceedings of the 15th International Conference on Autonomous Agents & Multiagent Systems (AAMAS)*. 1162–1170.
- [4] Prashant Doshi and Yifeng Zeng. 2009. Improved Approximation of Interactive Dynamic Influence Diagrams Using Discriminative Model Updates. In *AAMAS*. 907–914.
- [5] Prashant Doshi, Yifeng Zeng, and Qiongyu Chen. 2009. Graphical Models for Interactive POMDPs: Representations and Solutions. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)* 18, 3 (2009), 376–416.
- [6] Ya'akov Gal and Avi Pfeffer. 2008. Networks of Influence Diagrams: A Formalism for Representing Agents' Beliefs and Decision-Making Processes. *Journal of Artificial Intelligence Research (JAIR)* 33, 1 (sep 2008), 109–147.
- [7] P. J Gmytrasiewicz and P. Doshi. 2005. A Framework for Sequential Planning in Multiagent Settings. *Journal of Artificial Intelligence Research (JAIR)* 24 (2005), 49–79.
- [8] Kenneth A. De Jong. 2016. *Evolutionary Computation: A Unified Approach*. MIT Press, Cambridge, USA.
- [9] Daphne Koller and Brian Milch. 2003. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior* 45, 1 (2003), 181–221.
- [10] Mitchell Melanie. 1996. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, USA.
- [11] Yinghui Pan, Jing Tang, Biyang Ma, Yifeng Zeng, and Zhong Ming. 2021. Towards Data-driven Solutions to Interactive Dynamic Influence Diagrams. *Knowledge and Information Systems* <https://doi.org/10.1007/s10115-021-01600-5> (2021).
- [12] Yifeng Zeng and Prashant Doshi. 2012. Exploiting Model Equivalences for Solving Interactive Dynamic Influence Diagrams. *Journal of Artificial Intelligence Research (JAIR)* 43 (2012), 211–255.
- [13] Y. Zeng, P. Doshi, Y. Chen, Y. Pan, H. Mao, and M. Chandrasekaran. 2016. Approximating behavioral equivalence for scaling solutions of I-DIDs. *Knowledge and Information Systems* 49, 2 (2016), 511–552.
- [14] Y. Zeng, H. Mao, Y. Pan, and J. Luo. 2012. Improved Use of Partial Policies for Identifying Behavioral Equivalences. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1015–1022.