

Approximating Spatial Evolutionary Games using Bayesian Networks

Extended Abstract

Vincent Hsiao, Dana Nau, Xinyue Pan

University of Maryland
College Park

{vhsiao,nau}@umd.edu, xypan@terpmail.umd.edu

Rina Dechter

University of California
Irvine

dechter@ics.uci.edu

ABSTRACT

Evolutionary Game Theory is an application of game theory to evolving populations of organisms. Of recent interest are EGT models situated on structured populations or spatial evolutionary games. Due to the complexity added by introducing a population structure, model analysis is usually performed through agent-based Monte-Carlo simulations. However, it can be difficult to obtain desired quantities of interest from these simulations due to stochastic effects. We define a framework for modeling spatial evolutionary games using Dynamic Bayesian Networks that capture the underlying stochastic process. The resulting Dynamic Bayesian Networks can be queried for quantities of interest by performing exact inference on the network. We then propose a method for producing approximations of the spatial evolutionary game through the truncation of the corresponding DBN, taking advantage of the high symmetry of the model. This method generalizes mean-field and pair approximations in the literature for spatial evolutionary games. Furthermore, we show empirical results demonstrating the capability of the method to obtain much better accuracy than pair approximation with respect to stochastic simulations.

KEYWORDS

Spatial Evolutionary Games, Bayesian Networks, Moment Closure

ACM Reference Format:

Vincent Hsiao, Dana Nau, Xinyue Pan and Rina Dechter. 2021. Approximating Spatial Evolutionary Games using Bayesian Networks: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*, London, UK, May 3–7, 2021, IFAA-MAS, 3 pages.

1 INTRODUCTION

Evolutionary Game Theory (EGT) was initially developed to model biological evolution [19] but has found additional use in research on the evolution of cultural phenomena [6], and a variety of multi-agent systems topics [2, 23, 26, 27, 31]. While far from an exact description of human interactions, EGT models can be used to find trends that capture essential characteristics of modeled interactions [5, 7]. One can rely on agent-based stochastic simulations [1] to obtain insights. However, these simulations come with their own limitations in validation [18] and variability [21]. A commonly used alternative is pair approximation [13, 17] which has been used to

obtain qualitative insight into EGT models [12, 13, 20, 24] even if it is not very accurate to the underlying stochastic model [15, 16, 25, 29].

We propose a framework for exact modeling of spatial evolutionary games using a Dynamic Bayesian Network (DBN) [4], thus making the whole toolbox of probabilistic inference algorithms applicable to such stochastic games [3, 8–11, 14, 22]. Then we develop a method for producing approximations of stochastic spatial evolutionary games, by exploiting symmetry, through the truncation of the corresponding DBN. This provides a flexible framework for the exploration of higher order approximations beyond pair approximation that allow for better accuracy with respect to the underlying stochastic model. Finally, we provide preliminary empirical results illustrating the potential of our approach in modeling stochastic simulations and its advantages over existing approximations.

2 DBN EVOLUTIONARY-GAME MODEL

We consider a population of M agents $\{1, \dots, M\}$ placed on evenly-spaced points in a grid with circulatory boundary conditions. An evolutionary game consists of T iterations, each having an interaction phase followed by an update phase. In the interaction phase, each agent i chooses some action $s_i \in S$ and receives a payoff π_i as the sum of the payoffs received from playing a normal-form game with payoff matrix U against each of its neighbors $N(i)$. Each time the update phase occurs, a percentage of agents γ in the population use an update rule to decide how to change their strategies. For example, in the Fermi rule agents choose a random neighbor to compare their payoff with. Agents may also have a probability μ of mutating to a random strategy during the update phase [30].

We next define a model that fully captures our spatial evolutionary game using a Dynamic Bayesian Network (DBN). The DBN is $(X(t), D(t), P(t))$. The variable set $X(t) = X$ is split into two sets of variables $X = A \cup \text{Pay}$, where

- $A_{i,j}(t) \in A$: the strategy of the agent at coordinate (i, j) on the grid at the start of each iteration t .
- $\text{Pay}_{i,j}(t) \in \text{Pay}$: the payoff received by the agent at (i, j) during the interaction phase at time t .

CPT for Payoff Nodes Each node $\text{Pay}_{i,j}(t)$ has $d + 1$ parents consisting of $A_{i,j}(t)$ and its d neighbors $N(A_{i,j}) = \{A_{k,l} | (k, l) \in N(i, j)\}$. The conditional probability function $P(\text{Pay}_{i,j}(t) | \text{parents})$ is constructed from the payoff matrix U , expressed as a CPT:

$$\begin{aligned} & \Pr(\text{Pay}_{i,j}(t) | A_{i,j}(t), N(A_{i,j}(t))) \\ &= \begin{cases} 1 & \text{if } \text{Pay}_{i,j}(t) = \sum_{(k,l) \in N(i,j)} U(A_{i,j}(t), A_{k,l}(t)) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

CPT for Strategy Variables. Each $A_{i,j}(t+1)$ has $2(d+1)$ parents: $A_{i,j}(t)$, $Pay_{i,j}(t)$ and the $A(t)$ and $Pay(t)$ variables for each of the d neighboring $N(i, j)$ agents. Our goal is to define $\Pr(A_{i,j}(t+1) \mid \text{parents})$. We break up this probability into cases controlled by three variables:

- update: did an update happen (yes, with probability γ)?
- mut: did mutation happen (yes, with probability μ)?
- rand: which neighbor was chosen ((k, l) , with probability $\frac{1}{d}$)?

We use indicator functions to define

$$\Pr_{\delta} = \mathbb{1}_{A_{i,j}(t+1)=A_{k,l}(t)}, \Pr_{\emptyset} = \mathbb{1}_{A_{i,j}(t+1)=A_{i,j}(t)}$$

For example, if (update = 1) and (mut = 0), we can write:

$$\Pr(A(t+1)_{i,j} = s_{t+1} \mid A_{i,j}(t) = s_t, \text{other parents}) = \sum_{(k,l) \in N(i,j)} \frac{1}{d} \Pr_u(Pay_{i,j}, Pay_{k,l}) \Pr_{\delta}(1 - \Pr_{\emptyset}) + \Pr_{\emptyset}$$

where $\Pr_u(Pay_{i,j}, Pay_{k,l})$ is the probability that the agent (i, j) switches to strategy of the neighboring agent (k, l) .

The resulting DBN formulation fully encodes the stochastic process of the spatial evolutionary game. However, it is well known that exact inference is exponential in the tree-width of the network. To address this computation issue we propose a novel method for truncating the full DBN taking inspiration from moment-closure methods in the mean-field approximation literature.

3 TRUNCATION APPROXIMATION

We construct a Bayesian Network for each iteration that takes the states of each agent from input nodes at time t to output nodes at time $t+1$. Since the marginal distributions for each agent node in the exact model are identical, we can exploit symmetry by looking at the distribution of a few representative nodes. The idea is to truncate the DBN around a single focal agent and some of its neighbors. The input nodes consist of agents in the truncation neighborhood and the output nodes will be a few (e.g. one or two) nodes. We then use the transition probabilities in the exact model to link the input nodes to the output nodes. Since there are less output nodes than input nodes at the next time step, we must also approximate the joint distribution of the truncation neighborhood at the start of each iteration. There are three steps:

Truncation Neighborhood First, we choose some subset of agent nodes $B \subset A$ (examples shown in Fig. 2). We construct the




Figure 1: Slice of Dynamic Bayesian Network for the Fermi update rule centered at the agent located at position (1,1)

Bayesian Network from time t to time $t+1$ for just the nodes in B . On the truncated DBN, we can now run exact inference algorithms in a reasonable amount of time in order to compute the marginal and pair-wise conditional distributions of the output nodes at $t+1$.

Output query Second, we query a selection of lower order distributions from the output nodes at time $t+1$: $P_{s_i}, P_{s_i|s_j}$.

Input definition Finally, we approximate the joint distribution over the truncation neighborhood at the next iteration using a function of the single variable and pair distributions obtained during the output query step. We can approximate this using a tree-like Bayesian Network. We refer to our full paper for details.




Figure 2: Example Truncation Neighborhoods

4 EMPIRICAL EVALUATION

We run empirical experiments on several games from the evolutionary game theory literature. We test four different cases (BN-MF, BN-PA, BN-Medium, BN-Large) with increasing truncation neighborhoods. Approximation results are compared to the average of 20 simulations on a 50x50 grid. Figure 3 shows two examples.




Figure 3: Proportion of agents playing the second strategy in a Prisoner's Dilemma game (left) and a Snowdrift game (right) for different approximations.

Our results show that larger approximation neighborhoods reduce the error in the time evolution graphs, even in games such as Snowdrift (see [28, Section 3.8]) where pair approximation does not have good quantitative agreement with simulation results.

Acknowledgment. This work supported in part by NSF grant IIS-2008516 and AFOSR grant 1010GWA357. The information in this paper does not necessarily reflect the position or policy of the funders, and no official endorsement should be inferred.

REFERENCES

- [1] Christoph Adami, Jory Schossau, and Arend Hintze. 2016. Evolutionary game theory using agent-based methods. *Physics of life reviews* 19 (2016), 1–26.
- [2] Robert L Axtell. 2002. Non-cooperative dynamics of multi-agent teams. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 3*. 1082–1089.
- [3] B. Bidyuk and R. Dechter. 2007. Cutset sampling for Bayesian networks. *Journal of Artificial Intelligence (JAIR)* (2007).
- [4] Adnan Darwiche. 2009. *Modeling and reasoning with Bayesian networks*. Cambridge university press.
- [5] Soham De, Michele J Gelfand, Dana Nau, and Patrick Roos. 2015. The inevitability of ethnocentrism revisited: Ethnocentrism diminishes as mobility increases. *Scientific reports* 5, 1 (2015), 1–7.
- [6] Soham De, Dana S Nau, and Michele J Gelfand. 2017. Understanding norm change: An evolutionary game-theoretic approach. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. 1433–1441.
- [7] Soham De, Dana S. Nau, Xinyue Pan, and Michele J. Gelfand. 2018. Tipping Points for Norm Change in Human Cultures. *Lecture Notes in Computer Science* (2018), 61–69. https://doi.org/10.1007/978-3-319-93372-6_7
- [8] R. Dechter. 1999. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence* (1999), 41–85.
- [9] Rina Dechter. 2013. Reasoning with probabilistic and deterministic graphical models: Exact algorithms. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 7, 3 (2013), 1–191.
- [10] Rina Dechter. 2019. *Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms, Second Edition*. Morgan & Claypool Publishers. <https://doi.org/10.2200/S00893ED2V01Y201901AIM041>
- [11] Rina Dechter and Robert Mateescu. 2007. AND/OR search spaces for graphical models. *Artificial Intelligence* 171, 2-3 (2007), 73–106.
- [12] Stephen P Ellner. 2001. Pair approximation for lattice models with multiple interaction scales. *Journal of theoretical biology* 210, 4 (2001), 435–447.
- [13] Feng Fu, Martin A Nowak, and Christoph Hauert. 2010. Invasion and expansion of cooperators in lattice populations: Prisoner’s dilemma vs. snowdrift games. *Journal of theoretical biology* 266, 3 (2010), 358–366.
- [14] Vibhav Gogate and Rina Dechter. 2011. SampleSearch: Importance sampling in presence of determinism. *Artif. Intell.* 175, 2 (2011), 694–729.
- [15] Christoforos Hadjichrysanthou, Mark Broom, and Istvan Z Kiss. 2012. Approximating evolutionary dynamics on networks using a neighbourhood configuration model. *Journal of theoretical biology* 312 (2012), 13–21.
- [16] Christoph Hauert and Michael Doebeli. 2004. Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature* 428, 6983 (2004), 643–646.
- [17] Christoph Hauert and György Szabó. 2005. Game theory and physics. *American Journal of Physics* 73, 5 (2005), 405–414.
- [18] Benjamin Herd, Simon Miles, Peter McBurney, and Michael Luck. 2013. Verification and validation of agent-based simulations using approximate model checking. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation*. Springer, 53–70.
- [19] Josef Hofbauer, Karl Sigmund, et al. 1998. *Evolutionary games and population dynamics*. Cambridge university press.
- [20] Qing Jin, Lin Wang, Cheng-Yi Xia, and Zhen Wang. 2014. Spontaneous symmetry breaking in interdependent networked game. *Scientific reports* 4 (2014), 4095.
- [21] Gianluca Manzo and Toby Matthews. 2014. Potentialities and limitations of agent-based simulations. *Revue française de sociologie* 55, 4 (2014), 653–688.
- [22] Robert Mateescu, Kalev Kask, Vibhav Gogate, and Rina Dechter. 2010. Join-Graph Propagation Algorithms. *J. Artif. Intell. Res. (JAIR)* 37 (2010), 279–328.
- [23] Javier Morales, Michael Wooldridge, Juan A Rodríguez-Aguilar, and Maite López-Sánchez. 2018. Off-line synthesis of evolutionarily stable normative systems. *Autonomous agents and multi-agent systems* 32, 5 (2018), 635–671.
- [24] Mayuko Nakamaru, Hiroyuki Matsuda, and Yoh Iwasa. 1997. The evolution of cooperation in a lattice-structured population. *Journal of theoretical Biology* 184, 1 (1997), 65–81.
- [25] Hisashi Ohtsuki and Martin A Nowak. 2008. Evolutionary stability on graphs. *Journal of Theoretical Biology* 251, 4 (2008), 698–707.
- [26] Steve Phelps, Peter McBurney, and Simon Parsons. 2010. Evolutionary mechanism design: a review. *Autonomous agents and multi-agent systems* 21, 2 (2010), 237–264.
- [27] Marc Ponsen, Karl Tuyls, Michael Kaisers, and Jan Ramon. 2009. An evolutionary game-theoretic analysis of poker strategies. *Entertainment Computing* 1, 1 (2009), 39–45.
- [28] Carlos Pérez Roca. 2009. *Cooperation in evolutionary game theory: effects of time and structure*. Ph.D. Dissertation. Universidad Carlos III de Madrid.
- [29] György Szabó and Csaba Töke. 1998. Evolutionary prisoner’s dilemma game on a square lattice. *Physical Review E* 58, 1 (1998), 69.
- [30] Arne Traulsen, Christoph Hauert, Hannelore De Silva, Martin A Nowak, and Karl Sigmund. 2009. Exploration dynamics in evolutionary games. *Proceedings of the National Academy of Sciences* 106, 3 (2009), 709–712.
- [31] Karl Tuyls and Simon Parsons. 2007. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence* 171, 7 (2007), 406–416.