

Evolutionary Synthesis of Stable Normative Systems

(Extended Abstract)

Javier Morales¹ and
Michael Wooldridge²
University of Oxford, Oxford, UK
¹javier.morales@cs.ox.ac.uk
²mjlw@cs.ox.ac.uk

Juan A. Rodríguez-Aguilar
Artificial Intelligence Research
Institute (IIIA-CSIC)
Campus UAB, Bellaterra, Spain
jar@iiia.csic.es

Maite López-Sánchez
Math & Comp. Science Dept.
Universitat de Barcelona
Barcelona, Spain
maite_lopez@ub.edu

ABSTRACT

Normative systems are a widely used framework to coordinate interdependent activities in multi-agent systems. Most research in this area has focused on how to compute normative systems that effectively accomplish a coordination task, as well as additional criteria such as synthesising norms that do not over-regulate a system, and the emergence of norms that remain stable over time. We introduce a framework for the synthesis of stable normative systems that are sufficient and necessary for coordination. Our approach is based on ideas from evolutionary game theory. We simulate multi-agent systems in which useful norms are more likely to prosper than useless norms. We empirically show the effectiveness of our approach in a simulated traffic domain.

Keywords

Norms - Norm synthesis - Evolutionary game theory

1. INTRODUCTION

Within the area of multi-agent systems, normative systems have been widely used as coordination mechanisms [5, 16]. Most research in the area has focused on the problem of how to compute a normative system that effectively coordinates a system, which was proven to be NP-Complete in [17]. Two alternative approaches have been proposed to solve this problem: *off-line* and *on-line*. Off-line approaches aim at designing normative systems at design time [17, 8, 1, 6]. On-line approaches aim at creating norms at run-time, either by emerging from the agent society [3, 18, 22, 2, 14], or by being created by a norm synthesis mechanism based on the observation of the system [11, 12]. On-line approaches are considered to be more flexible than off-line approaches, since norms can be adapted as a system evolves.

Most work in normative systems has focused on synthesising *effective* normative systems that successfully achieve certain goals [17, 1, 10]. Other works consider further criteria, such as the synthesis of *minimal* normative systems that, while being effective, do not over-regulate a system [7, 8, 12, 13]. Additionally, some works take inspiration on the framework of evolutionary game theory, and study how *sta-*

ble norms emerge as behavioural patterns of agent societies [3, 15, 18, 4, 9, 21]. They consider a game-theoretic setting [20, 19] in which agents iteratively play a game, learning the best strategies. Eventually, agents converge to stable behavioural patterns that can be considered as stable norms.

In this paper, we incorporate ideas from evolutionary game theory (EGT) to synthesise evolutionarily stable normative systems that effectively regulate a system without lapsing into over-regulation. We simulate a multi-agent system in which agents repeatedly play a collection of games over time, using norms to coordinate. Norms that are proven to be useful for coordination prosper and spread within the agent society, while useless norms are eventually discarded by the agents. The outputs of these simulations are normative systems whose norms are sufficient and necessary for coordination, and are evolutionarily stable.

2. EVOLUTIONARY NORM SYNTHESIS

We consider a simulated traffic junction scenario where agents are cars. The interactions between cars are regarded as *games* that can lead to *collisions*. Figure 1 illustrates two game examples. Game 1 (left of Figure 1) depicts two cars perceiving one another with available actions “*go*” (move forward) and “*stop*” (give way). Below each game is its payoff matrix. Both cars can avoid collisions (getting payoff 1) only if at least one of them stops. Otherwise, they collide and get payoff 0. In game 2 (right of Figure 1), both cars are required to stop in order to avoid collisions.

Norms are *soft* constraints aimed to regulate cars’ behaviours in games in order to avoid collisions. As an example, consider norms n_1 and n_2 designed for game 1.

$$n_1 : \langle \varphi, (“car1”, obl(stop)) \rangle \quad n_2 : \langle \varphi, (“car2”, obl(stop)) \rangle$$

These norms have a sentence φ that describes the game played between cars 1 and 2. Both norms oblige either car 1 or car 2 to stop, respectively. Note that collisions can be avoided if either car 1 applies n_1 , or car 2 applies n_2 . If both cars apply their respective norms they would remain stopped needlessly, and we would say the game is over-regulated.

Each car has a normative system containing different norms to coordinate in different games. In a game, each car can identify which norms apply to it out of its normative system, and can decide whether or not to comply with these norms. Each norm has an *utility* to avoid collisions, computed in terms of whether it is *effective* and *necessary*. A norm is effective if it avoids collisions to the cars that *apply* it, and it is necessary if cars collide after *violating* it. In game 1,




Figure 1: Games in a traffic junction together with their payoff matrices. Game 1 has two cars perceiving one another. Game 2 is similar, but both cars have a stopped car in front.

both norms n_1 and n_2 are effective, but n_1 is necessary only if car 2 does not have norm n_2 (and the other way around).

Our aim is to synthesise a normative system that contains effective, necessary and evolutionarily stable norms to coordinate in each possible game. In this way, no car could benefit from switching to an alternative normative system. In our example, this would equal to choosing a useful and stable combination of norms to regulate game 1 (either n_1 or n_2), and so on for each game. With this aim, we take inspiration on EGT, which provides a model for the evolution of agent societies, and their convergence to evolutionarily stable strategies. We run simulations of our traffic scenario, and simulate norm evolution in terms of their utilities.

Our simulations consider a population of cars that initially have different normative systems. Each norm has a *frequency* that stands for the proportion of cars whose normative systems contain the norm. Each simulation runs in rounds of 1,000 ticks. In each round, cars interact playing games, applying and violating their applicable norms. Once a game is played, the utility of each norm is updated. At the end of each round, norms are *replicated* in terms of their *fitness*. The fitness of a norm stands for its average utility to a car that *has* the norm when it interacts with cars that have other norms. Such fitness is computed in terms of: (1) the utility of a car with such a norm when interacting with cars that have alternative norms; and (2) the probability to interact with cars that have these alternative norms (i.e., their frequencies). Norms whose fitness is above the average increase their frequency proportionally, whereas norms whose fitness is below the average decrease their frequency.

After norm replication, a new population of cars is created whose normative systems contain each norm according to its frequency. Thus, if a norm is 70% frequent, then 70% of cars will have it in their normative systems. In this way, the frequency of fittest norms increases after each round, and that of less fit norms decreases.

3. EMPIRICAL ANALYSIS AND RESULTS

We performed an empirical evaluation to analyse the performance of our approach. We ran 500 simulations of our traffic scenario, and analysed the normative systems synthesised upon convergence. In each game, cars have 10% probability to violate the norms that apply to them (and 90% to apply norms). Each simulation stops when the frequency of each norm remains stable (unchanged) for 20 rounds.

Remarkably, 100% of simulations converged to a norma-




Figure 2: Normative systems synthesised upon convergence. Each circle represents a normative system. Blue normative systems are right-hand side priority. White normative systems are left-hand side priority.




Figure 3: Competition of Ω_1 against mutant normative systems. The x axis shows different generations of normative systems created over time. The y axis depicts the identifier of different normative systems.

tive system whose norms are effective and necessary to avoid collisions. Figure 2 graphically represents the normative systems synthesised upon convergence. Each circle represents a normative system (Ω). The square on top of each circle stands for the proportion of simulations that converged to that normative system. The subset relationship between circles represents a normative system that is a subset of another (e.g., Ω_2 is a subset of Ω_1). Out of 500 simulations, our system converged to twelve different normative systems. Normative systems Ω_1 to Ω_7 contain different combinations of norms to give way to the right, and the remaining normative systems (Ω_8 to Ω_{12}) contain different combinations of norms to give way to the left. In particular, 67% of simulations converged to Ω_1 (right-hand side normative system), and 18% of simulations to Ω_8 (left-hand side normative system). The remaining simulations converged to similar normative systems (subsets of Ω_1 and Ω_8 that differ in 1-2 norms).

Stability analysis. We analysed the stability of the normative systems synthesised upon convergence. We performed 100 extra simulations that started with a population whose cars unanimously abided by Ω_1 (Figure 2). Each simulation lasted 400 rounds. In each round, we created mutant cars that abided by normative systems different from Ω_1 . At the end of each round, norms were replicated as described in Section 2. All simulations converged to a population whose agents abided by Ω_1 . Figure 3 illustrates one of these simulations. The x -axis shows different rounds, and the y -axis depicts the identifiers of the normative systems created over time. Black points represent mutant normative systems created in each round, and the red line represents the id of the most frequent normative system. For the sake of clarity, we represent Ω_1 as the normative system with id 1,000. After 200 rounds, the simulation created 2,500 mutant normative systems. Upon round 400, normative system Ω_1 remained stable most of the time. In punctual rounds, a big amounts of mutants were created, making the frequency of Ω_1 to go below stability. But, after a few rounds, Ω_1 replicated and became again the most frequent normative system. Upon round 400, the simulation converged to Ω_1 .

Acknowledgements. Work funded by the H2020-MSCA-IF project number 707688. Mike Wooldridge was supported by the ERC under Advanced Grant 291528 (“RACE”). Maite López-Sánchez and Juan A. Rodríguez-Aguilar were supported by projects Collectiveware (TIN2015-66863-C2-1-R, MINECO/FEDER), 2009-SGR-1434, and 2014 SGR 623.

REFERENCES

- [1] T. Ågotnes, W. Van Der Hoek, C. Sierra, and M. Wooldridge. On the logic of normative systems. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI '07*, pages 1175–1180, 2007.
- [2] S. Airiau, S. Sen, and D. Villatoro. Emergence of conventions through social learning. *Autonomous Agents and Multi-Agent Systems*, 28(5):779–804, 2014.
- [3] R. Axelrod. An evolutionary approach to norms. *American political science review*, 80(04):1095–1111, 1986.
- [4] O. H. Azar. What sustains social norms and how they evolve?: The case of tipping. *Journal of Economic Behavior & Organization*, 54(1):49–64, 2004.
- [5] G. Boella, L. Van Der Torre, and H. Verhagen. Introduction to normative multiagent systems. *Computational & Mathematical Organization Theory*, 12(2-3):71–79, 2006.
- [6] D. Corapi, A. Russo, M. de Vos, J. Padget, and K. Satoh. Normative design using inductive learning. *Theory and Practice of Logic Programming*, 11(4-5):783–799, 007 2011.
- [7] D. Fitoussi and M. Tennenholtz. Minimal social laws. In *AAAI*, pages 26–31, 1998.
- [8] D. Fitoussi and M. Tennenholtz. Choosing social laws for multi-agent systems: Minimality and simplicity. *Artificial Intelligence*, 119(1-2):61–101, 2000.
- [9] F. Germano. Stochastic evolution of rules for playing finite normal form games. *Theory and Decision*, 62(4):311–333, 2007.
- [10] M. Hjelmbloom. *Offline Norm Evolution*, pages 316–333. Springer International Publishing, Cham, 2015.
- [11] J. Morales, M. Lopez-Sanchez, J. A. Rodriguez-Aguilar, M. Wooldridge, and W. Vasconcelos. Automated synthesis of normative systems. In *Proceedings of the AAMAS'13*, pages 483–490. IFAAMAS, 2013.
- [12] J. Morales, M. Lopez-Sanchez, J. A. Rodriguez-Aguilar, M. Wooldridge, and W. Vasconcelos. Minimality and simplicity in the on-line automated synthesis of normative systems. In *AAMAS '14: Proceedings of the 13th international conference on autonomous agents and multiagent systems*, pages 109–116, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS).
- [13] J. Morales, M. Lopez-Sanchez, J. A. Rodriguez-Aguilar, M. Wooldridge, and W. Vasconcelos. Synthesising liberal normative systems. In *AAMAS '15: Proceedings of the 14th international conference on autonomous agents and multiagent systems*, pages 433–441. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 04/05/2015 2015.
- [14] R. Riveret, A. Artikis, D. Busquets, and J. Pitt. Self-governance by transfiguration: From learning to prescriptions. In *International Conference on Deontic Logic in Computer Science*, pages 177–191. Springer, 2014.
- [15] R. Sethi. Evolutionary stability and social norms. *Journal of economic behavior & organization*, 29(1):113–140, 1996.
- [16] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, 2009.
- [17] Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: off-line design. *Artificial Intelligence*, 73(1-2):231–252, February 1995.
- [18] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94(1):139–166, 1997.
- [19] B. Skyrms. *Evolution of the Social Contract*. Cambridge University Press, 1996.
- [20] J. M. Smith. *Evolution and the Theory of Games*, pages 202–215. Springer US, Boston, MA, 1988.
- [21] T. Sugawara. Emergence and stability of social conventions in conflict situations. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume One*, pages 371–378. AAAI Press, 2011.
- [22] D. Villatoro, J. Sabater-Mir, and S. Sen. Social instruments for robust convention emergence. In *IJCAI*, pages 420–425, 2011.