Achieving Emergent Governance in Competitive Multi-Agent Systems

Doctoral Consortium

Michael Pernpeintner
Institute for Enterprise Systems (InES), University of Mannheim
Mannheim, Germany
pernpeintner@es.uni-mannheim.de

ABSTRACT

Our PhD research is concerned with the task of achieving cooperation in a system of competitive agents which cannot be explicitly controlled. To this end, it examines the problem from the system's point of view, without restricting the agents' behavior or requiring specific knowledge about their decision-making.

The governance of the MAS will be achieved via a dynamically adaptive *governing policy* based on a set of rules, which leaves full autonomy to the individual agents, but reacts to their actions via suitable changes of the environment. The mechanism is designed to lead to system-level cooperation while only assuming that the agents follow their own self-interested motives.

KEYWORDS

Multi-Agent System, Self-interested Agents, Competition, Emergent Governance, Normative System, Cooperation, Mechanism Design

ACM Reference Format:

Michael Pernpeintner. 2020. Achieving Emergent Governance in Competitive Multi-Agent Systems. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

1.1 Competitive Multi-Agent Systems

Acting successfully in a non-cooperative Multi-Agent System requires multiple skills: Besides the lack of predictability regarding the environment, one might face strategic and even destructive behavior from other agents, whereas cooperation and trust can never be taken for granted. Now imagine being in charge for governing such an MAS without being able to directly control the agents: You can only observe them, try to understand the rationale behind their actions and give incentives for constructive behavior—but what does constructive even mean in this context?

In contrast to cooperative MAS settings, where equilibria can often be reached through effective information sharing and coordination since all participants are aiming for the same common goal [2, 7], non-stable behavior and strategic adaptation are inherent in competitive systems [13], and therefore are necessary on the governance level as well.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

We are specifically working towards an *emergent MAS governance* given by a set of rules and a player who constructively changes the environment, such that the system fosters robust cooperation between the agents without relying on manual rule adaptation or human intervention at run-time. While the decision-making process and reasoning techniques of an individual agent in such a scenario have been analyzed in detail in recent years—using for example Markov decision processes, game theory, swarm intelligence and graph theoretic models [10]—we want to regard the agents as autonomous black boxes and focus on the "community level" of the MAS instead. Of course, a central assumption for the entire work is that the agents' actions can in principle be influenced by some sort of incentive scheme.

1.2 Emergence in MAS

The general concept of emergence derives from nature, where certain properties of systems only arise as the result of interactions of the system's components, but not on the individual components' level [4]. In our context, emergence refers to the fact that a *governing player* and corresponding *action policy* are built and adapted at run-time from the interplay of the rule set and the agents' actions. The agents do not purposefully construct this emergent policy, and yet it is formed through their involvement.

2 MODEL

2.1 Terminology

Based on the agent definition of [11], we consider a finite set $P = \{p_1, ..., p_n\}$ of agents (or players) which, at every time step $t \in \mathbb{N}$, perceive the current state $s_t \in S$ of a dynamic, temporally discretized environment and then act within this environment by each performing an action $a_i \in A_i, i \in \{1, ..., n\}$, following their individual (and confidential) action profile $\pi_i : S \to A_i$. The environmental state changes from one time step to the next according to the combination of actions taken by the agents, i.e., using a transition function $\delta : S \times A \to S$.

Additionally, the system contains a finite set $\mathcal{R} = \{R_1, ..., R_k\}$ of *rules*, where $R \subset S \;\forall R \in \mathcal{R}$. Here, the elements $s \in R$ are precisely the environmental states that are allowed by a rule R.

2.2 Governance Model

The rules describe the desired states of the environment, such that the task of governing the system essentially consists of minimizing the number of time steps t where $s_t \notin \cap_{R\in\mathcal{R}} R$.

In order to steer an agent without making any assumption about the right "stimulus" to which it reacts, we can only impact the environment in a way that plays back to the agent. We model this mechanism by using a *governing player* p_0 which differs from the other players insofar as it can observe their actions before choosing its own, while still acting in the same time step. Therefore, the action space A in the domain of the transition function is $A = \prod_{i=0}^{n} A_i$, including the action space A_0 of p_0.

The action policy $\pi_0 : S \times \prod_{i=1}^{n} A_i \rightarrow A_0$ of p_0 (also called the *governing policy*), which depends on approximating the other players' action profiles and acting accordingly, is the core of the emergent governance, and its deduction, adaptation and analysis is the ultimate goal of our PhD research.

3 RESEARCH QUESTIONS

When dealing with self-interested agents and conflicting goals, any governance scheme needs to be concerned with manipulative and non-cooperative behavior. As a consequence, our research questions are focused on strategic and evolutionary topics:

RQ1 Under which conditions does a governing policy as outlined in Section 2.2 converge towards a stable function?
RQ2 How can an agent manipulate the mechanism, and how can this be effectively prevented by p_0?
RQ3 Is full agent autonomy generally compatible with an effective governance mechanism? If not, which restrictions on agent behavior are necessary?

4 RELATED WORK

Our approach to emergence borrows loosely from [6] whose definition of emergent MAS is based on the three perspectives *Subject*, *Condition* and *Method*. In their setting, emergence amounts to autonomously finding a function which is considered "adequate" by a "relevant user". However, this implies that a user is able to (manually) evaluate and judge the quality of such a function. Emergent collaboration in the smaller context of *graphical collaboration games* is examined in [12], introducing a scoring mechanism which maximizes social welfare. Here, the authors define overall social welfare as the sum of individual utilities and model collaboration as a one-dimensional decision between 0 and 1, which requires that all components have similar utility scales and agree on a common option space. As [1] point out, emergent governance mechanisms become crucial when systems are too big or evolve too fast to be monitored by humans, as for example algorithmic trading systems, networks of business process services or autonomous cars.

The idea of using norms and rules to describe and control an agent's behavior is used in [5]. The authors define a formal language to communicate the normative position, i.e., obligations, permissions and prohibitions of an agent, to the other participants of an MAS. As opposed to our approach, this method is agent-centric and does not take the perspective of an outside observer.

A framework based on [5] is [8], which aims to create a system of norms at run-time without the involvement of agents. While considering not only effectiveness, but also compactness and liberality as qualitative measures for norm synthesis, the framework's governance mechanism relies on the regulative power of norms (i.e., a norm can prohibit certain actions or combinations of actions)—an instrument that we explicitly want to avoid.

When looking at the system-level perspective of MAS instead of the agents themselves, most approaches still require specific agent models: [14] split the total reward into intrinsic and extrinsic reward and restrict their analysis to solving *Intertemporal Social Dilemmas* (ISDs). [3] use norms and sanctions to control agent behavior; they propose heuristic strategies to revise sanctions when a norm is seen to be ineffective. However, they can only deal with a static set of norms and rely on assumptions on the agents' preference types. [9] also tackle the problem by explicitly restricting the agents' option space in order to enforce cooperation.

4.1 Research Gap

In existing research, we observe that individual agents are mostly at the center of attention, and even when they are not, there are still restrictions regarding their decision-making and control mechanisms. A central assumption in all cited works is that an external utility function can be fed to an agent such that this agent will react to rewards and sanctions as they are provided. Moreover, many approaches rely on explicit control of agents' actions in order to reach a desired outcome.

Our perspective does not make such restrictions, but allows for general action policies π_i which the governing player approximates from observations. In addition, the interaction of agents and governance is carried out directly in the environmental domain S, without the need for the agents to listen to external control.

5 OUTLOOK

After some further refinement of the model, our next steps will be the creation of a minimal working prototype for testing and evaluation, and the subsequent extension and adaptation of this prototype.

5.1 Evaluation

To evaluate our approach, we will need to tackle a number of challenges, amongst them: (a) How can the rules be expressed in a way that allows for fast evaluation of $s \in R$ for any given state s and rule R, and for easy identification of the violating agent(s)? (b) How can agents be designed which are complex enough to account for a realistic setting, yet simple enough to allow for precise analysis of the results? (c) Is real-time adaptation of the governing policy feasible and necessary? What would be other options?

5.2 Model Extensions

Starting from the basic model, there is a number of interesting tweaks that can make the setting more realistic and more powerful. Some ideas are: (a) Continuous degree of rule satisfaction instead of binary values, (b) Partial observability of the environment by the governing player, (c) The effect of trust, both between agents and between governance and agents, (d) Taking into account past actions, i.e., equipping π_0 with a memory, and (e) Allowing for a set of (distributed) governing players.

ACKNOWLEDGMENTS

This work is supported by the German Federal Ministry of Transport and Digital Infrastructure (BMVI).

REFERENCES

[1] N. Bulling and M. Dastani. 2016. Norm-based Mechanism Design. *Artif. Intell.* 239 (2016), 97–142. https://doi.org/10.1016/j.artint.2016.07.001

[2] C. Claus and C. Boutilier. 1998. The Dynamics of Reinforcement Learning in Cooperative Multiagent Systems. In *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence (AAAI '98/IAAI '98)*. American Association for Artificial Intelligence, 746–752.

[3] D. Dell'Anna, M. Dastani, and F. Dalpiaz. 2019. Runtime Revision of Norms and Sanctions Based on Agent Preferences. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, 1609–1617.

[4] J. Fromm. 2005. Types and Forms of Emergence. (2005). https://www.researchgate.net/publication/2147516_Types_and_Forms_of_Emergence

[5] A. García-Camino, J. Rodríguez-Aguilar, C. Sierra, and W. Vasconcelos. 2006. A rule-based approach to norm-oriented programming of electronic institutions. *SIGecom Exchanges* 5 (2006).

[6] J.-P. Georgé and M.-P. Gleizes. 2005. Experiments in Emergent Programming Using Self-organizing Multi-agent Systems. In *Multi-Agent Systems and Applications IV (Lecture Notes in Computer Science)*, M. Pěchouček, P. Petta, and L. Varga (Eds.). Springer, Berlin, Heidelberg, 450–459.

[7] T. Kemmerich and H. Kleine Büning. 2012. A Convergent Multiagent Reinforcement Learning Approach for a Subclass of Cooperative Stochastic Games. In *Adaptive and Learning Agents (Lecture Notes in Computer Science)*, Peter Vrancx, Matthew Knudson, and Marek Grześ (Eds.). Springer, 37–53.

[8] J. Morales. 2016. *On-line norm synthesis for open Multi-Agent systems*. Ph.D. Dissertation. Universitat de Barcelona.

[9] G. Perelli. 2019. Enforcing Equilibria in Multi-Agent Systems. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*. 188–196.

[10] Y. Rizk, M. Awad, and E. Tunstel. 2018. Decision Making in Multi-Agent Systems: A Survey. *IEEE Transactions on Cognitive and Developmental Systems* PP (May 2018), 1–1.

[11] S. Russell and P. Norvig. 2009. *Artificial Intelligence: A Modern Approach, 3rd edition*. Pearson.

[12] A. Sinha and M.P. Wellman. 2019. Incentivizing Collaboration in a Competition. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, 556–564.

[13] N. Vlassis. 2007. A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence. In *A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence*.

[14] J. Wang, E. Hughes, Ch. Fernando, W. Czarnecki, E. Duéñez-Guzmán, and J. Leibo. 2019. Evolving Intrinsic Motivations for Altruistic Behavior. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, 683–692.