Emergence of Cooperation on Networks

Doctoral Consortium

Jacques Bara
University of Warwick
Coventry, United Kingdom
jack.bara@warwick.ac.uk

ABSTRACT

The emergence of cooperation is a major question in game theory and one under-studied aspect is the effects of networks on the emergent behaviour. My PhD asks this question over multiple collaborations and projects, using methodologies from (evolutionary) game theory, agent-based simulations, networks and complex systems. As a researcher, however, I am interested in an even wider and broader variety of topics and have such collaborated on other projects focusing on the effects of networks and hypergraphs on learning problems (namely agent learning and rating prediction).

KEYWORDS

Emergence of Cooperation; Networks; Dynamic Networks; Public Goods Games; Agent-based Modelling and Simulation; Game Theory; Opinion Dynamics; Voting; Interpretable Machine Learning; Hypergraphs

ACM Reference Format:

Jacques Bara. 2023. Emergence of Cooperation on Networks: Doctoral Consortium. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

Complex behaviours often emerge out of even the simplest intelligent multi-agent systems, such as the emergence of cooperation in evolutionary game theory. My Ph.D. at the University of Warwick investigates the effects of the underlying space or network on such emergent properties, by representing agent intelligence with parsimonious models and analysing their dynamic interactions with the environment. In particular my research focuses on the emergence of cooperation as well as other complex behaviours.

In this extended abstract I give abstracts of the multiple projects and collaborations I have undertaken under this large remit. I first discuss our work on the emergence of cooperation on dynamic networks under a variety of strategy-update and partner-update rules [8], particularly noting that core-peripheral structures also form. I find similarly that cooperation and mesoscale structures emerge out of spatial public goods games of pollution [7] towards which migration and density plays an important role. I then focus on predicting dynamic outcomes for biased political elections on gerrymandered social networks [4, 5]. Identifying the impacts of structure and space is a core and recurrent aspect of my work, as I outline the impacts of networks on agent learning in noisy information flows [6] as well as how hypergraph representations of structured data can improve interpretable rating predictions. I briefly turn to a data-based project identifying the impacts of infrastructural changes to road networks on social connectivity [1], before finally concluding with future avenues of research.

Cooperation on Dynamic Networks. With the increased importance of network-mediated interaction, researchers have shifted the attention to the impact of social networks and their dynamics in promoting or hindering cooperation, drawing various context-dependent conclusions. For example, some lines of research, theoretical and experimental, suggest the existence of a threshold effect in the ratio of timescales of network evolution, after which cooperation will emerge [12, 19], whereas other lines dispute this, suggesting instead a Goldilocks zone [16, 23].

In this project [8] we provide an evolutionary game theory framework to understand coevolutionary processes from a bottom up perspective - in particular the emergence of a cooperator-core and defector-periphery - clarifying the impact of partner selection and imitation strategies in promoting cooperative behaviour, without assuming underlying communication or reputation mechanisms. In doing so we provide a unifying framework to study imitation-based cooperation in dynamic social networks and show that disputes in the literature can in fact coexist in so far as the results stem from different equally valid assumptions.

Migration in a Spatial Social Dilemma of Pollution. Much attention has been devoted to understanding cooperation in populations where agents interact with random peers (well-mixed), interact over complex networks, or interact in fixed spatial positions [18, 20, 25]. In spatial settings with mobile agents, however, the effects of cooperation are circumscribed to arbitrary neighbourhoods and the stability of cooperation depends on individuals' capacity to move between sites and form dense clusters [9, 10, 21]. The existing interaction models, however, ignore the long-range effects of undesirable behaviour, which our project addresses.

In this project [7], we study spatial public goods games in which agents either pollute (defectors) in a large area or clean (cooperators) their local area and can migrate to empty sites within range. We ask whether migration promotes cooperation and reduces the negative impacts of defection. Analytically and through agent-based simulations, we show that migration ultimately reduces the pollution felt per-capita in at least two ways: 1) polluters encourage eco-friendly neighbours to migrate away, eventually clustering with other cooperators and 2) migration stabilises cooperation in dense population scenarios. Our results reveal a complex interaction between migration and density as key factors to promote cooperation in spatial social dilemmas.

Biased Elections on Social Networks. A recently proposed graph-theoretic metric, the *influence gap* [24], has shown to be a reliable

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

predictor of the effect of social influence in two-party elections, albeit only tested on regular and scale-free graphs. In this project, we investigate whether the influence gap is able to predict the outcome of multi-party elections on networks exhibiting community structure, i.e., made of highly interconnected components, and therefore more resembling of real-world interaction. To encode communities we build on the classical model of caveman graphs [27], which we extend to a richer graph family that displays different levels of *homophily* [3, 26], i.e., how much connections and opinions are intertwined.

In the initial conference paper [4], we study the predictive power of the influence gap in the presence of communities. We show that when there is no clear initial majority the influence gap is not a good predictor of the election outcome. When we instead allow for varying majorities, although the influence gap improves as a predictor, counting the initial partisan majority does consistently better, across all levels of homophily. Moreover, we study the combined effect of the more predictive metrics, as function of the homophily levels. Using regression models, we demonstrate that the influence gap combined with the initial votes count does increase the overall predictive power for some levels of homophily.

Finally we extend our work in a journal paper [5], studying elections with more than two parties. Specifically, we extend the definition of the influence gap to any number of parties, considering various generalisations, and show that the initial votes count has an even higher predictive power when compared to influence gap than it did in the two-party case.

Noisy Information Flow in Networked Learning. We study the problem of noisy information propagation in networks - building on the Grapevine model [11] where a small number of sources send messages across the network - and agents use Bayesian updates to make inferences about the state of the world from the received messages (similar to [13]). We provide upper bounds on the total number of sources necessary for learning on a given network and refine the bound in the case of small-world networks. We then extend the model to include an adversarial attacker, who can corrupt some of the information sources.

We find that there is an optimal greedy attacking strategy in the case of a single learner, while the multi-learner case is not always solved optimally using greedy approaches. However, despite the influence function not being submodular, we show that the greedy algorithm performs well in practice. We also show that much simpler heuristics, which only look at centrality measures, can also provide a good basis to calculate successful attacking strategies. Finally we analyse the loss of optimality in the case when the attacker has incomplete information about the network and has to estimate the influence of source corruption heuristically. We use real-world social networks, as well as random network models, to empirically evaluate the effectiveness of attacking strategies and suggest a variety of measures to counteract them.

Hypergraph-based Interpretable Machine Learning. Given potential synergies or discords between team members, predicting future group ratings from past aggregate data is not always trivial. Will an individual belonging to a high-performing group contribute positively to a new collective? How to infer future ratings based on graphs representing previous groups' composition? While these challenges can be tackled with state-of-the-art graph representation learning approaches, in this project we focus on equipping interpretable methods (e.g., linear regression) (see for example discussion in [22]) with data engineered from the underlying group formation hypergraph.

We show that explicitly including data on the temporal group formation hypergraph improves future rating prediction. We compare three major approaches to predicting hyperedge quantities (i.e., group ratings): a classical linear regression approach which treats group composition as dummy variables; a hypergraph centric approach which calculates and aggregates intermediary quantities for constituent nodes and a line graph centric approach which projects the hypergraph into weighted directed line graphs. In essence, the two latter models use the inherent structure of the data (the hypergraph) to inform more intelligent feature engineering. Using the Internet Movie Database as a dataset, we find that the hypergraph methods significantly improve accuracy, with the line graph centric method having comparative - at times, better - accuracy than the hypergraph centric one.

The Impact of Structural Changes on Social Connectivity. Large infrastructural changes may result in lower travel times between regions and thus reduce their mutual access times or spatial distance. As a consequence, social connectivity is expected to increase among such regions with decreasing spatial distance of access times [2, 15, 17]. To explore this phenomenon, we use a unique geolocated, timestamped database from a Hungarian online social network called iWiW, containing data from 2002 through 2012, covering around 40% of the total population throughout its life-cycle [14]. To address infrastructural changes, we compile a historical road network dataset from OpenStreetMap coupled with data about the year of construction and average speed limit.

Focusing on those pairs that experienced a substantial change in their physical connection, we compute a variety of measures, and compare them to spatially-informed null models of social connectivity. Our results confirm that a decrease in travel times between pairs of settlements has a lagged increasing effect on social connectivity.

Future Research. Although we have found that migration and spatial density promotes cooperation (stylised as eco-friendly behaviour) [7] we will extend our model to capture *costly* migration - a more realistic assumption given the time cost and financial fees tenants must spend in order to move house. Moreover there are interesting questions regarding wealth disparity and to what extent this impacts the wider population.

Regarding the Ph.D. as a whole I will be developing further the connective tissue that ties together the variety of projects I have already undertaken, to synthesise into a coherent thesis. In particular I will focus on the themes of underlying structures and spaces, and how they impact and are impacted by multiagent dynamics.

ACKNOWLEDGMENTS

This work was supported by Engineering and Physical Sciences Research Council grant EP/S022244/1 for the MathSys CDT.

REFERENCES

[1] Yasaman Asgari, Jacques Bara, Paul Bouman, Eszter Bokanyi, Giulio Burgio, Franziska Franeck, Matteo Mazzamurro, and Liubov Tupikina. 2022. Poster: The impact of infrastructural change on social connectivity. Conference on Complex Systems.

[2] Michael Bailey, Patrick Farrell, Theresa Kuchler, and Johannes Stroebel. [n.d.]. Social Connectedness in Urban Areas. 118 ([n. d.]), 103264. https://doi.org/10/gnbpzg

[3] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.

[4] Jacques Bara, Omer Lev, and Paolo Turrini. 2021. Predicting Voting Outcomes in Presence of Communities. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Virtual Event, United Kingdom) *(AAMAS '21)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 151–159.

[5] Jacques Bara, Omer Lev, and Paolo Turrini. 2022. Predicting voting outcomes in the presence of communities, echo chambers and multiple parties. *Artificial Intelligence* 312 (2022), 103773. https://doi.org/10.1016/j.artint.2022.103773

[6] Jacques Bara, Charlie Pilgrim, Paolo Turrini, and Stanislav Zhydkov. 2023. The Grapevine Web: Analysing the Spread of False Information in Social Networks with Corrupted Sources. In *Proceedings of the 22nd International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '23)*. (in press).

[7] Jacques Bara, Fernando P. Santos, and Paolo Turrini. 2023. The Role of Space, Density and Migration in Social Dilemmas. In *Proceedings of the 22nd International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '23)*. (in press).

[8] Jacques Bara, Paolo Turrini, and Giulia Andrighetto. 2022. Enabling imitation-based cooperation in dynamic social networks. *Autonomous Agents and Multi-Agent Systems* 36, 2 (31 May 2022), 34. https://doi.org/10.1007/s10458-022-09562-w

[9] Dirk Helbing and Wenjian Yu. 2008. Migration as a mechanism to promote cooperation. *Advances in Complex Systems* 11, 04 (2008), 641–652.

[10] Dirk Helbing and Wenjian Yu. 2009. The outbreak of cooperation among success-driven individuals under noisy conditions. *PNAS* 106, 10 (2009), 3680–3685. https://doi.org/10.1073/pnas.0811503106 arXiv:https://www.pnas.org/content/106/10/3680.full.pdf

[11] Matthew O. Jackson, Suraj Malladi, and David McAdams. 2020. Learning through the Grapevine: The Impact of Message Mutation, Transmission Failure, and Deliberate Bias. In *EC*. ACM, 645.

[12] Jillian J. Jordan, David G. Rand, Samuel Arbesman, James H. Fowler, and Nicholas A. Christakis. 2013. Contagion of Cooperation in Static and Fluid Social Networks. *PLOS ONE* 8, 6 (06 2013), 1–10. https://doi.org/10.1371/journal.pone.0066199

[13] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. 137–146.

[14] Balázs Lengyel, Attila Varga, Bence Ságvári, Ákos Jakobi, and János Kertész. [n.d.]. Geographies of an Online Social Network. 10, 9 ([n. d.]), e0137248. https://doi.org/10/8fj arXiv:26359668

[15] Jure Leskovec and Eric Horvitz. [n.d.]. Geospatial Structure of a Planetary-Scale Social Network. 1, 3 ([n. d.]), 156–163. https://doi.org/10/gnbp2k

[16] Aming Li, Lei Zhou, Qi Su, Sean P. Cornelius, Yang-Yu Liu, Long Wang, and Simon A. Levin. 2020. Evolution of cooperation on temporal networks. *Nature Communications* 11, 1 (08 May 2020), 2259. https://doi.org/10.1038/s41467-020-16088-w

[17] David Liben-Nowell, Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. [n.d.]. Geographic Routing in Social Networks. 102, 33 ([n. d.]), 11623–11628. https://doi.org/10.1073/pnas.0503018102 arXiv:16081538

[18] Hisashi Ohtsuki, Christoph Hauert, Erez Lieberman, and Martin A. Nowak. 2006. A simple rule for the evolution of cooperation on graphs and social networks. *Nature* 441, 7092 (01 May 2006), 502–505. https://doi.org/10.1038/nature04605

[19] David G. Rand, Samuel Arbesman, and Nicholas A. Christakis. 2011. Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences* 108, 48 (2011), 19193–19198. https://doi.org/10.1073/pnas.1108243108 arXiv:https://www.pnas.org/content/108/48/19193.full.pdf

[20] David G. Rand, Martin A. Nowak, James H. Fowler, and Nicholas A. Christakis. 2014. Static network structure can stabilize human cooperation. *Proceedings of the National Academy of Sciences of the United States of America* 111, 48 (02 Dec 2014), 17093–17098. https://doi.org/10.1073/pnas.1400406111

[21] Carlos P. Roca and Dirk Helbing. 2011. Emergence of social cohesion in a model society of greedy, mobile individuals. *Proceedings of the National Academy of Sciences* 108, 28 (2011), 11370–11374. https://doi.org/10.1073/pnas.1101044108 arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.1101044108

[22] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (01 May 2019), 206–215. https://doi.org/10.1038/s42256-019-0048-x

[23] Hirokazu Shirado, Feng Fu, James H. Fowler, and Nicholas A. Christakis. 2013. Quality versus quantity of social ties in experimental cooperative networks. *Nature Communications* 4, 1 (14 Nov 2013), 2814. https://doi.org/10.1038/ncomms3814

[24] Alexander J. Stewart, Mohsen Mosleh, Marina Diakonova, Antonio A. Arechar, David G. Rand, and Joshua B. Plotkin. 2019. Information gerrymandering and undemocratic decisions. *Nature* 573, 7772 (2019), 117–121.

[25] Alex Stivala, Yoshihisa Kashima, and Michael Kirley. 2016. Culture and cooperation in a spatial public goods game. *Phys. Rev. E* 94 (Sep 2016), 032303. Issue 3. https://doi.org/10.1103/PhysRevE.94.032303

[26] Alan Tsang and Kate Larson. 2016. The Echo Chamber: Strategic Voting and Homophily in Social Networks. In *Proceedings of the 15th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (Singapore, Singapore) *(AAMAS '16)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 368–375.

[27] Duncan J. Watts. 1999. Networks, Dynamics, and the Small-World Phenomenon. *Amer. J. Sociology* 105, 2 (1999), 493–527.