INFORMATION DYNAMICS OF LEARNED SIGNALLING GAMES

MATTHEW SPIKE, SIMON KIRBY, KENNY SMITH

Language Evolution and Computation Research Unit, Linguistics and English Language, University of Edinburgh, Edinburgh, UK

matthew.spike@ed.ac.uk

Signalling games involving agent learners exist in various guises, from the game-theoretic Roth-Erev learners of Skyrms (2010), to the Naming Game (Steels, 1997), and agents employing varieties of observational learning (e.g. Oliphant & Batali, 1996; Smith, 2002). The agent-based nature of this work means that the resulting dynamics have an inherently unpredictable character: individual simulations may or may not be representative of average behaviour, if such a thing exists at all. Typically, the best way of overcoming this problem is by running large numbers of simulations and observing the aggregate behaviour. This contrasts with other frameworks — for example, classical or evolutionary game theory. In these cases, there is some *macro-level* property of the model which drives the overall dynamic of the game. For example, fitness of individual agents in evolutionary models is evaluated using the global average communicative success. Because of this, it is possible to calculate the mean-field dynamic for any known mixture of strategies in the population, revealing any attractors or stable points. In the case of agent-based models, because overall dynamics are completely determined by individual pairwise interactions — at the *micro-level* (Mühlenbernd, 2013) — the likely result of any interaction is not a direct consequence of the global communicative success of a population, which as a result cannot serve to describe the overall dynamics. Hence, identifying attractors and stable points poses a much harder problem. In order to resolve this problem, we introduce a new information-theoretic measure of optimality which *can* describe the overall dynamics of signalling populations of learning agents.

Typically, information theory (Shannon, 1948) has proven difficult to apply to problems involving meaningful communication as it has no way of describing semantic or referential content. Although there have been attempts to address this (e.g. Corominas-Murtra, Fortuny, & Solé, 2014), these still include a problematic macro-level term such as described above. However, we are able to avoid this under the assumption that agent signalling production and reception behaviours are derived from a single shared set of signal meaning associations. In this case, we can use the signal production behaviour of individual agents to describe their *in-*

dividual optimality in terms of the conditional entropy of meanings given signals, $H(M|S)$, where low entropy represents low ambiguity. Employing this measure, we show that the overall entropy of a system has two components determined by the average *individual entropy* and average *alignment entropy*: individual entropy measures the optimality of a single agent's own signalling system, while alignment entropy is the extra uncertainty due to the divergence of any agent from the population mean. We draw on results such as (Xue, 2006) which show that any population of agents which *imitate* each other with positive probability will inevitably drive the alignment entropy to zero.

This allows us to dissect the overall dynamics of any signalling game involving associative agents, which we do by analysing the pairwise interaction defined by its model of learning. In particular, we can describe any population as a point in an entropy *state-space*. Certain points within this space represent final stable states of the population in terms of their optimality. As such, we are able to show that the way 'imitative' learning by itself causes populations to move around the state-space resembles a type of genetic drift. Moreover, we identify the features which must exist to ensure populations develop optimal signalling: firstly, the *imitative* property described above; secondly, the learning model must on average *reduce conditional entropy* in any pairwise interaction. Finally, there must be a way to prevent learning slowdown: i.e. agents must retain *plasticity*. Using these three factors as a diagnostic, we are able to determine the dynamics of any population model involving associative signalling agents without recourse to numerical simulation, including whether or not it will develop optimal signalling. This applies to not just modelling work, but any theory of the emergence of novel lexicons.

References

Corominas-Murtra, B., Fortuny, J., & Solé, R. V. (2014). Towards a mathematical theory of meaningful communication. *Scientific reports*, *4*, 4587.

Mühlenbernd, R. (2013). *Signals and the Structure of Societies.* Phd thesis, University of Tübingen.

Oliphant, M., & Batali, J. (1996). Learning and the Emergence of Coordinated Communication. *Center for research on language newsletter*, *11*(1), 1–46.

Shannon, C. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*(July, October), 379–423, 623–656.

Skyrms, B. (2010). *Signals: Evolution, Learning & Information.* Oxford: Oxford University Press.

Smith, K. (2002). The cultural evolution of communication in a population. *Connection Science*, *14*(1), 65–84.

Steels, L. (1997). Self-organising vocabularies. In *Proceedings of the 5th international workshop on artificial life: Synthesis and simulation of living systems.* ALIFE-96.

Xue, J. (2006). A Polya Urn Model of Conformity.