Understanding Language Evolution in Overlapping Generations of Reinforcement Learning Agents

Lewys G. Brace[1] and Seth Bullock[2]

[1]Institute for Complex Systems Simulation, University of Southampton, Southampton, United Kingdom.
[2]Department of Computer Science, University of Bristol, Bristol, United Kingdom.
L.G.Brace@soton.ac.uk

Abstract

Understanding how the dynamics of language learning and language change are influenced by the population structure of language users is crucial to understanding how lexical items and grammatical rules become established within the context of the cultural evolution of human language. This paper extends the recent body of work on the development of term-based languages through signalling games by exploring signalling game dynamics in a social population with overlapping generations. Specifically, we present a model with a dynamic population of agents, consisting of both mature and immature language users, where the latter learn from the formers' interactions with one another before reaching maturity. It is shown that populations in which mature individuals converse with many partners are more able to solve more complex signalling games. While interacting with a higher number of individuals initially makes it more difficult for language users to establish a conventionalised language, doing so leads to increased diversity within the input for language learners, and that this prevents them from developing the more idiosyncratic language that emerge when agents only interact with a small number of individuals. This, in turn, prevents the signalling conventions having to be renegotiated with each new generation of language users, resulting in the emerging language being more stable over subsequent generations of language users. Furthermore, it is shown that allowing the children of language users to interact with one another is beneficial to the communicative success of the population when the number of partners that mature agents interact with is low.

Introduction

The fact that children around the world are readily able to learn the language of their given social group, even though these languages are in a constant state of flux (Hopper, 1987) and exhibit high levels of variation, flexibility of usage, and are ever changing over time within dynamic populations (Christiansen and Kirby, 2003) indicates that cultural factors play a crucial role in the shaping of human language. The establishment of the meanings of lexical items and the subsequent change in these meanings over time is, in part, what led Lewis (1969) to work on the conventionality of meaning; how specific arbitrary signals establish themselves as referring to a specific meaning. He introduced a signalling game in order to explore how meaningful language might evolve from the use of initially random signals. Over the last decade, renewed interest in these ideas has led to a body of work that has explored the evolution of term-based languages through coordination games (Skyrms, 2004, 2009, 2010; Huttegger, 2007; Barrett, 2006, 2009; Argiento et al., 2009).

In this paper, we further develop this work by implementing a reinforcement learning (R-L) model involving a single Sender-Receiver pair, which is then extended into a population-based, multi-generational, simulation. This is both novel and necessary, given that human language persists in a complex social milieu, which is not captured by standard R-L models, and employing the R-L procedure in a population-based model could therefore offer insights into how lexical items become established within the context of the cultural evolution of human language in structured populations with overlapping generations.

This paper presents a model that demonstrates that, if agents only interact with a small number of other agents, then it is easier for these agents to establish a conventionalised system of language usage than in cases where they interact with a larger number of the population. However, by interacting with a smaller subset of the population, these individuals develop a more idiosyncratic language. Thus, it becomes difficult for the children of these agents, who learn from the interactions of their parents, to communicate with children of other mature agents during future epochs. In contrast, allowing individuals to interact with a larger proportion of the population does initially make it more difficult to establish agreed upon conventions of usage, but it does result in an increased amount of diversity within the language learner's training input data. This better enables the children of these mature agents to successfully interact with the offspring of other mature agents, previously unencountered by the agent in question; this aids the negotiation of conventional signalling in the population as a whole. This, in turn, leads to the development of a language that is more stable and consistent over generations of language users, compared to the case where individuals have to rene-

gotiate conventions of use in every generation.

The $n = 2$ game

In a Lewis signalling game there are two players, a Sender and a Receiver. A single bout of the game commences with the Sender knowing that the world is in some random state, t, but the Receiver being ignorant of this information. The Sender then selects a signal, s, with which to convey the world state to the Receiver; the Receiver observes s and has to pick an appropriate action, a. If the action chosen by the Receiver matches the world state (i.e., $a = t$), the bout is considered to have been a success. Here, t, s, and a are drawn from finite sets T, S, and A, respectively, which are all of size n; in Lewis' (1969) original model $n = 2$.

Over successive bouts of the game, both players are expected to adapt their behaviour in order to increase the chance of achieving communicative success, typically through some kind of reinforcement learning. The easiest way to conceptualise this is in terms of urns and balls. At the outset of the simulation run, an unbiased Sender will have n urns, one for each state of the world, each of which contains n balls, one associated with each of the n possible signals. Let's suppose that during the first bout of the game, $t = $ "red". The Sender picks a random ball from their *red* urn. The symbol on this ball dictates the signal to be made, s; in this case, suppose $s = $ "fah". Likewise, the Receiver observes $s = $ "*fah*", and picks a random ball from their *fah* urn, which indicates the action to be taken, a. Both balls are then returned to their respective urns. If $a = t$, the interaction was a success, and in accordance with the principles of Roth-Erev reinforcement learning (Roth and Erev, 1995), the Sender adds extra balls of type s to urn t and the Receiver adds extra balls of type a to urn s. The number of extra balls, u, added to the urns corresponds to the utility associated with the outcome of the signalling bout; in Lewis' (1969) original game $u = 1$ if a bout is successful and $u = 0$ otherwise. More formally, at any point in time, $b(t, s)$ is the number of balls for signal s in the Sender's urn for state t, and accordingly, $b(s, a)$ is the number of balls corresponding to act a in the Receiver's urn s. Thus, the behavioural strategies for Sender (σ) and Receiver (ρ) are as follows:

$$\sigma(t, s) = \frac{b(t, s)}{\sum_{s' \in S} b(t, s')} \qquad \rho(s, a) = \frac{b(s, a)}{\sum_{a' \in A} b(s, a')} \tag{1}$$

There are a number of possible signalling equilibria that can arise in such a game. Perfect signalling strategies result in optimal pay-offs for the players by mapping each world state onto a unique signal and each signal onto the unique appropriate action (Figure 1). This behaviour constitutes an evolutionarily stable strategy (ESS) because when it is played by the whole population there is no incentive for any individual to change their strategy.
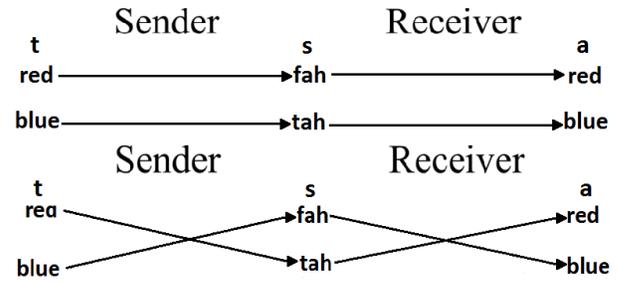

Figure 1: Optimal strategies for the $n = 2$ game.

However, players may spend significant time playing sub-optimal 'pooling' strategies in which Senders employ the same signal for multiple world states (pooling these world states together), making it impossible for Receivers to determine the state of the world from the signals that they receive (Figure 2). Pooling strategies are not ESSs since adjacent strategies often achieve equal fitness (e.g., the two pooling strategies in fig 2). The expected pay-off for such a pooling strategy is 0.5 when $n = 2$.
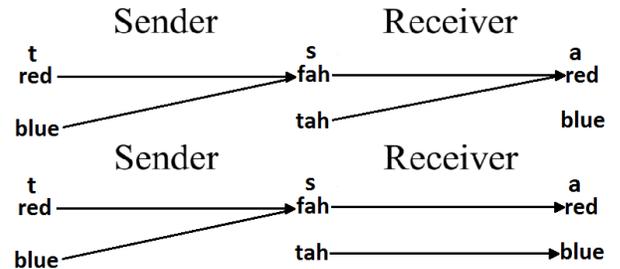

Figure 2: Two of the possible sub-optimal pooling strategies for the $n = 2$ game.

It has been shown by way of both computational simulation (Barrett, 2006, 2009; Skyrms, 2010) and mathematical modelling (Huttegger, 2007; Argiento et al., 2009) that the $n = 2$ game will nearly always converge upon an optimal signalling system. Indeed, Skyrms (2009) went on to demonstrate that this behaviour also holds in a case where there are two Senders and one Receiver. These results are further supported by Table 1, where it can be seen that a computational model of the $n = 2$ game being played for 10^6 bouts will almost always reach a perfect signalling equilibrium.

Higher-n games

However, a successful outcome is not always achieved when the game is played with $n > 2$, i.e., with a higher number of states, signals and actions (Skyrms, 2010; Huttegger, 2007; Barrett, 2006, 2009). Here, we adopt the methodology of Barrett (2006, 2009). The R-L model that formed the basis of the population-based R-L model was run multiple times, for various values of n, with each run consisting

of 10^6 bouts, B, of the game, where a run of the simulation is considered to fail if the number of successful bouts is less than 90% of the total number of bouts. Table 1 (left) shows the results of these runs, which agree with those of Barrett (2006, 2009). Table 1 (right) shows the results of a smaller sample of 100 runs, with all other parameters being held constant, and results that are quite similar. It can be seen clearly from Table 1 that, in a $n=3$ game, the players fail to achieve a high enough rate of signalling success roughly 10% of the time, and that this increases to $\approx 20\%$ for $n=4$ games, $\approx 60\%$ for $n=8$ games, and so on. The comparison in Table 1 is important to show, as the extended model presented later is run for 100 runs due to limits on computational power.

	1000 runs		100 runs
$n = 2$	0.999	$n = 2$	0.99
$n = 3$	0.881	$n = 3$	0.87
$n = 4$	0.784	$n = 4$	0.84
$n = 8$	0.391	$n = 8$	0.33
$n = 10$	0.281	$n = 10$	0.22
$n = 20$	0.264	$n = 20$	0.23

Table 1: Table depicting the success rates of the R-L model after 10^6 bouts for various values of n, with 1000 runs (left) and 100 runs (right).

Generational Reinforcement Learning Model

Although interesting in their own right, the dyadic setting considered so far limits the conclusions that can be drawn. After all, human language persists in a highly complex social milieu, and it has been shown that the structure and composition of a population can influence the dynamics of language change over time (Brace et al., 2015). Thus, the original reinforcement learning model (R-L) was extended in a number of ways. First, whereas the original model focused on a single Sender and Receiver playing for B bouts, in the population model (R-L-P), there are a population of agents. This population is divided into a number of mature agents, N_M, and a number of immature agents, N_I; with all agents starting life as immature and then being promoted to mature status after the first epoch of their existence. Mature agents play bouts of the language game with one another, updating their language behaviour according to game outcomes. By contrast, while agents are immature they merely observe the language bouts played by their mature parent, and update their language behaviour on the basis of the outcomes of these observed games. The lifespan of agents is two epochs; the first as an immature agent and the second as a mature agent, after which they are removed from the simulation.

It is important to emphasise here that throughout the simulation, when new immature individuals are introduced to the population, as in the standard R-L model, they have *no knowledge of the language currently being used*. This is true for the initial population of mature agents, and also true for new immature agents born into all subsequent epochs. For each immature agent, each world state, t, is associated equally with each signal, s, when playing as Signaller, and each signal, s, is associated equally with each action, a, when playing as Receiver, i.e., each of an immature agent's n state urns contains a single ball for each possible signal, and each of their signal urns contain a single ball for each possible action. Thus, any change in communicative performance or language use over generations is the result of language evolution; there is no biological evolution on the part of the agents.

Secondly, instead of agents merely interacting B times, the R-L model is extended to include a generational aspect. In other words, the model is set up to run for a number of epochs, E, and during each epoch, every mature agent plays B bouts as the Receiver with other mature agents; with the amount of bouts it plays as the Speaker being the result of how many other agents it is partnered with divided by B. The number of different mature agents that a mature agent interacts with, P, is a key parameter of the model; with each mature agent's total number of interactions, B, being equally divided amongst its P unique partners, i.e., the number of interactions that an agent has with each partner is B/P (rounding up). A key feature of the model is therefore that varying P does not vary the number of bouts played, just the number of players that the bouts are played with.

The R-L-P model thus proceeds as follows. At the start of the simulation run, an initial population of $N_M = 15$ unbiased mature agents are created, with an equal chance of generating each signal for each world state. For each epoch, E, a fresh population of $N_I = 15$ unbiased immature agents is created, each having an equal chance of generating each signal for each world state. Each immature agent is assigned a randomly selected mature agent to act as their parent; with each mature agent acting as a parent to only one immature agent. Each mature agent is assigned P unique randomly selected mature partners with which to play the signalling game. Each of the mature agents then engages in B/P bouts with each assigned partner, with each participant updating their signalling or receiving strategy at the end of each bout through reinforcement learning. Each child will update their behaviour based on the outcome of the bouts that their parent are involved in; i.e., at the end of a successful bout, a Sender's child will add a ball of type s to urn t, and a Receiver's child will add a ball of type a to urn s. At the end of an epoch, all mature agents are removed, all immature agents are promoted to mature agent status, and a new set of unbiased immature agents is created.

Results

The R-L-P model does not achieve a successful signalling system as often as the standard R-L model. Indeed, compar-
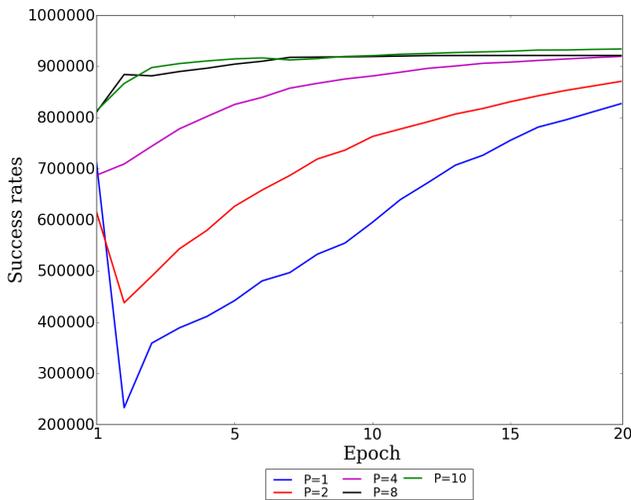
Figure 3: Graph depicting the average number of successful bouts across epochs for $P \in \{1, 2, 4, 8, 10\}$ for a $n = 20$ game with $N_M = 15$, $N_I = 15$, $B = 10^6$, and $u = 1$. Averaged over 30 runs.

ing Table 2 to Table 1 shows how, with $P = 1$, success rates are lower for all n-games than in the standard R-L model.

	Partners=1	Partners=2	Partners=5	Partners=10
$n = 4$	0.3	0.36	0.77	0.81
$n = 8$	0.0	0.3	0.48	0.62
$n = 10$	0.0	0.2	0.20	0.46
$n = 20$	0.0	0.0	0.29	0.57

Table 2: Language evolution success rates after 100 R-L-P model runs of $E = 20$ epochs each for various values of n and P, where $N_M = 15$, $N_I = 15$, $B = 10^6$, and $u = 1$. With success being measured using the aforementioned metric used in Table 1.

However, increasing the value of P does increase the rate of success (Table 2 and Figure 3). In Figure 3, we see that, with $P = 1$ or 2, there is an initial level of success, which corresponds to the number of successful bouts that would be seen in the normal R-L model for a n=20 game.

Low P values create a situation where mature agents form a communicative system based upon conventions agreed upon between themselves and only a small number of other agents. Thus, in subsequent epochs, when the children of a mature agent have to interact with the children of another mature agent, who has not previously interacted with the mature agent in question, the agreed upon conventions that both parties formulated during the first epoch are likely to be of little use, due to different agents forming conventions based upon their idiosyncratic experiences. This gives rise to sub-optimal behaviour at the population-level. However, any immature agents that are present learn from the success-
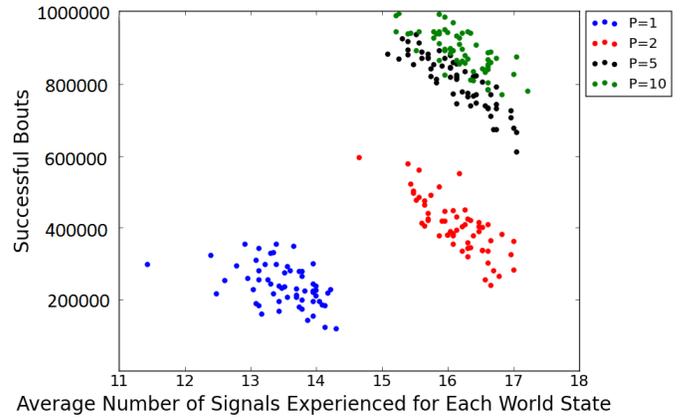

Figure 4: Graph depicting the average percentage of successful communicative bouts between all mature agents plotted against the number of unique signals presented to them during said bouts in the second epoch for a $n = 20$ game, for P=1, 2, 5, 10 and $N_M = 15$, $N_I = 15$, $B = 10^6$, and $u = 1$. Averaged over 60 runs.

ful bouts of their respective parents, hence the steady increase in success rates for these lower P values[1].

In contrast, with high P values, we see an obvious and immediate increase in communicative success. This is due to the way an increase in P leads to the children of the mature agents having more diversity in their training input, which better enables these individuals to communicate with a larger number of other agents upon being promoted to mature agent status (Figure 4).

Imagine a hypothetical mature agent from epoch one, who is partnered with ten other randomly selected agents, who in turn, are partnered with ten other agents. In the simulation, bouts are scheduled in such a way that $agent_1$ will have one of the allocated bouts with one of its randomly selected partners, then $agent_2$ will do the same; and so on, until we reach $agent_{M_N}$. At which point we go back to $agent_1$ and allow it to have its second bout, again with a randomly selected partner; and so on until each partner of every agent has played B/P bouts with the agent. In the P=1 case, unsurprisingly, we see higher levels of initial success during the first generation than in the P=10, due to the establishment of a convention involving fewer agents having to negotiate with one another (Figure 5, left).

In contrast, with $P = 10$, it is slightly harder to establish a conventionalised usage because each agent has to negotiate with an increased number of different agents, which results higher levels of signal diversity (Figures 4 and 5, left). However, when the offspring of the first epoch's mature agents are forced to interact with a different subset of the population in the second epoch, populations with higher

---

[1]Given enough epochs, it is likely that the agents would give rise to a successful communicative system.
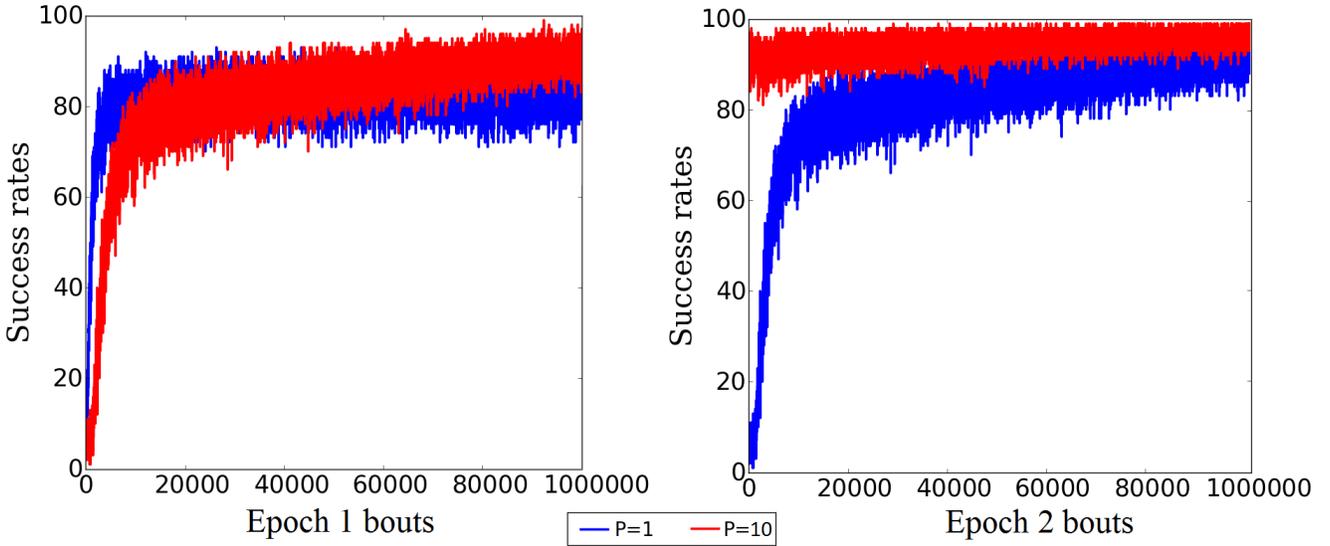
Figure 5: Graph depicting the number of successful bouts out of every 100 bouts, over all 10^6 bouts of a random agent during the first epoch (left) and the second epoch (right); with $P=1$ (blue line) and $P=10$ (red line). Where $n = 20$, $N_M = 15$, $N_I = 15$, $B = 10^6$, and $u = 1$.

P values exhibit higher communicative success because the increased signal diversity in the previous epoch, combined with the immature agents learning from the successful bouts of their parents, has resulted in these agents having established a conventionalised usage that requires less renegotiating when speaking to previously unencountered agents than in the $P = 1$ case, where agents have a more idiosyncratic language that requires them to renegotiate the conventions established by their parents (Figure 5, right).

This is why Figure 4 shows an increase in communicative success with higher values of P, while also indicating a negative trend in each of the data clusters for each specific P value; although it is harder to establish a language when negotiating meaning-signal pairs with more individuals, doing so makes it more stable across generations (Figure 5). Indeed, as Figure 5 (right) shows, the agreed upon convention of usage in cases of lower P values has to be renegotiated in subsequent epochs due to it offering little communicative success to agents when communicating with newly encountered individuals.

It is important to note that the increase in communicative success is the result of higher P values and not of another variable, such as B. Indeed, Figure 6 demonstrates the average level of communicative success over twenty epochs is significantly lower for $P = 1$ or 2, as compared to $P = 4$; a trend that continues as P is increased. Furthermore, it can be seen from Figure 6 that higher P values allow for an increased amount of communicative success even when agents have significantly fewer training sessions (lower B values).

However, in the real world, children are not just passive receivers of linguistic input. They interact with others, including other children, who may not yet be fully linguistically competent. Thus, a number of model runs were conducted whereby immature agents had B bouts with P other immature agents while witnessing their parents bouts (Figure 7). In other words, bouts are scheduled in a similar manner to that described above, in that we allow each agent to have one bout with a randomly selected partner, starting with $agent_1$ and cycling through to $agent_{N_M+N_I}$, before going back to $agent1$ again. In these runs, mature agents only interact with mature agents and immature agents only interact with other immature agents. Although, immature agents still learn from their parent's interactions, in the manner detailed above.

Figure 7 demonstrates how, performance in the $P = 10$ case is impeded by allowing interactions between immature agents. This is to be expected as linguistically underdeveloped individuals interacting with one another will add a degree of noise into the communicative system.

However, with $P = 1$, allowing immature agents to interact with one another dramatically increases communicative success. This difference in behaviour can again be attributed to signal diversity. While in the above results immature agents only learned from the interactions of their parents, meaning they got a degenerative sample of the language because they only ever witnessed the same two individuals communicating during their first epoch, here they are also interacting and learning with another individual who is likely to have witnessed two different mature agent's interacting with one another. This would increase the amount of signal diversity in the immature agents training data.
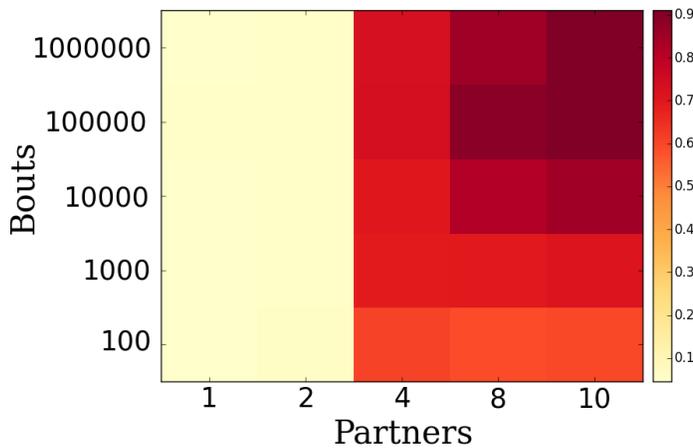
Figure 6: Graph depicting the average amount of communicative success over 20 epochs for various values of B and P. Where $n = 20$, $N_M = 15$, $N_I = 15$, $B = 10^6$, and $u = 1$. Averaged over 30 runs.

Figure 7: Graph depicting the average amount of communicative success over 20 generations for $P = 1$ where only N_M interact with one another (blue solid line) and where both N_M and N_I interact with other mature and immature agents, respectively (blue dashed line), and likewise for $P = 10$ (red solid line and red dashed line, respectively). Where $n = 20$, $N_M = 15$, $N_I = 15$, $B = 10^6$, and $u = 1$. Averaged over 30 runs.

Discussion

The results presented here build upon a larger body of work; both in regards to signalling conventions (Skyrms, 2004, 2010; Barrett, 2006, 2009) and expression/induction models research in general (Hurford, 2002). It has been shown that a signal can acquire a conventionalised meaning without the Sender intending for it to do so. Moreover, the meaning of such simple signals is dependent upon the stabilisation of usage conventions, which emerge from functional historical signal production. Thus, even the most automatic or reflexive signals can acquire meaning, so long as the production and response mechanisms are co-adapted to coordinate their behaviours in accordance with such an arbitrary signal (Harms, 2004).

More interestingly, it has been shown that a population structure that allows for interaction between more of its members is beneficial in allowing it to evolve an efficient term-based language.

More specifically, it has been shown that, as intuition dictates, while it is harder to establish a conventionalised system of usage with larger numbers of individuals, doing so enables the emerging language to persist in subsequent epochs. This is due to the input into language learners being initially more diverse, which prevents these learners from developing a more idiosyncratic, communicative system that makes it harder to communicate with previously unencountered individuals.

In addition, the results and behaviour reported here can be seen to be linked to the concept of a linguistic bottleneck. This refers to how the input data for a language learner will only be a subset of the potentially large range of grammars of the Speaker from which it is learning.

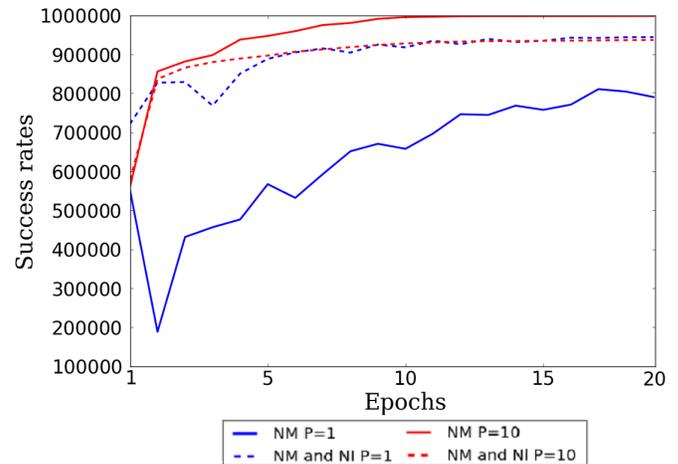A series of computer-based simulations that use the method of *iterated learning* have shown that the linguistic bottleneck is crucially important in regards to whether or not language can be successfully passed from one generation to the next and, in situations where this transmission can be achieved successfully, show that it is also crucial to the linguistic structure that arises (Kirby, 2002b,a; Kirby and Hurford, 2002; Kirby et al., 2014; Smith, 2002; Smith et al., 2003; Brace et al., 2015).

Although a similar effect to the bottleneck is seen in other types of uni-generational models, such as the *naming game* (Steels, 1995), the model presented here is novel in that it demonstrates the impact of bottleneck-like behaviour in a generational-based simulation that explores term-based languages. Here, this bottleneck-like behaviour takes the form of the way in which internal representations of individuals are induced from limited examples of the behaviour of other agents (Hurford, 2002).

This supports other work that has demonstrated a link between the linguistic bottleneck and the number of linguistic tutors (Brace et al., 2015). Indeed, the behaviour seen in Figure 7 indicates that the factors underpinning the cultural transmission of language change and linguistic variation are perhaps too complicated to be understood by analysing the nature of just inter- and intra-generational transmission; and that further research into linguistic change should focus on the nature of the social network the underpins linguistic populations (Wichmann and Holman, 2009; Lupyan and Dale, 2010; Reitter and Lebiere, 2010; Milroy, 2013).

This point becomes more important given that traditional

expression/induction models have largely ignored population dynamics so as to function on other aspects of language. Although, given the aims of such models, it made logical sense to opt for more simplistic population structures, it has been shown here that population dynamics can have a significant impact upon communicative behaviour.

Indeed, it would be interesting to explore how an expanding and contracting population size, with varying numbers of mature language users and immature language learners, could impact the emergence and form of a language (Johansson, 1997; Hurford, 2002). An expression/induction model geared towards this interest could provide valuable insights for a growing body of research that is interested in the nature of the relationship between language and population change; such as the impact of population size on linguistic forms or the way in which periods of linguistic simplification tend to coincide with periods where there are a higher number of language-learners within a population (Johansson, 1997; Nettle, 1999; Wichmann and Holman, 2009; Lupyan and Dale, 2010; Milroy, 2013; Trudgill, 2013). These themes will form the basis of future work.

References

Argiento, R., Pemantle, R., Skyrms, B., and Volkov, S. (2009). Learning to signal: Analysis of a micro-level reinforcement model. *Stochastic processes and their applications*, 119:373–390.

Barrett, J. (2009). The evolution of coding in signaling games. *Theory and decision*, 67:223–237.

Barrett, J. A. (2006). Numerical simulations of the Lewis signaling game: Learning strategies, pooling equilibria and the evolution of grammar. *University of California eScholarship*, page https://escholarship.org/uc/item/5xr0b0vp.

Brace, L., Bullock, S., and Noble, J. (2015). Achieving compositional language in a population of iterated learners. *Proceedings of the European conference on artificial life 2015*, 349-356.

Christiansen, M. and Kirby, S. (2003). Language Evolution: the Hardest Problem in Science? In Christiansen, M. and Kirby, S., editors, *Language Evolution*. Oxford University Press, Oxford.

Harms, W. (2004). Primitive content, translation, and the emergence of meaning in animal communication. In O., K. and Griebel, U., editors, *Evolution of communication systems: a comparative approach*. The MIT Press, London.

Hopper, P. (1987). Emergent grammar. *Berkeley Linguistics Society*, 13:139–157.

Hurford, J. R. (2002). Expression/induction models of language evolution: Dimensions and issues. In Briscoe, T., editor, *Linguistic evolution through language acquisition: formal and computational models*. Cambridge University Press, Cambridge.

Huttegger, S. (2007). Evolutionary explanations of meaning. *Philosophy of science*, 74:1–27.

Johansson, C. (1997). *A view from language: growth of language in individuals and populations*. Lund University Press, Lund.

Kirby, S. (2002a). Learning, bottlenecks and the evolution of recursive syntax. In Briscoe, T., editor, *Linguistic evolution through language acquisition: formal and computational models*, pages 173–204. Cambridge University Press., Cambridge.

Kirby, S. (2002b). Natural language from artificial life. *Artificial life*, 8(2):185–215.

Kirby, S., Griffiths, T., and Smith, K. (2014). Iterated learning and the evolution of language. *Current opinion in Neurobiology*, 28:108–114.

Kirby, S. and Hurford, J. (2002). The emergence of linguistic structure: an overview of the iterated learning model. In Cangelosi, A. and Parisi, D., editors, *Simulating the Evolution of Language*. Springer, London.

Lewis, D. (1969). *Convention: a philosophical study*. Harvard University Press, Harvard, USA.

Lupyan, G. and Dale, R. (2010). Language structure is partly determined by social structure. *PloS one*, 5(1):e8559.

Milroy, L. (2013). Social networks. In Chambers, J. and Schilling, N., editors, *The handbook of language variation and change*. Wiley-Blackwell, Oxford.

Nettle, D. (1999). Is the rate of linguistic change constant? *Lingua*, 108:119–136.

Reitter, D. and Lebiere, C. (2010). Did social networks shape language evolution?: a multi-agent cognitive simulation. In *Proceedings of the 2010 Workshop on Cognitive …*, pages 9–17. Association for Computational Linguistics, Stroudsburg, USA.

Roth, A. and Erev, I. (1995). Learning in extensive form games: experimental data and simple dynamical models in the intermediate term. *Games and economic behaviour*, 8(1):164–212.

Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge University Press., Cambridge.

Skyrms, B. (2009). Evolution of signalling systems with multiple senders and receivers. *Philosophical transactions of the royal society of London. Series B*, 364:771–779.

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

Smith, K. (2002). Natural selection and cultural selection in the evolution of communication. *Adaptive behavior*, 10(1):25–45.

Smith, K., Kirby, S., and Brighton, H. (2003). Iterated learning: a framework for the emergence of language. *Artificial life*, 9(4):371–86.

Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial Life*, 2:319–332.

Trudgill, P. (2013). Linguistic and social typology. In Chambers, J. and Schilling, N., editors, *The handbook of language variation and change*. Wiley-Blackwell, Oxford.

Wichmann, S. and Holman, E. (2009). Population size and rates of language change. *Human biology*, 81(2):259–274.