Replicating Oliphant's Saussurean Simulations

Richard Littauer

"A Saussurean communication system exists when an entire communicating population uses a single language that maps states unambiguously onto symbols and then back into the original states" (Oliphant 1996: 31) Oliphant constructed simulations in order to illuminate the conditions necessary for the evolution and maintenance of such Saussurean systems. I replicate his first two simulations that involved direct pressure for good transmission. I go on to produce a third, novel simulation, which looks at gradations of reward for successful communication, in order to ensure that the first two simulations are not guilty of begging the question. I then further describe the benefits of communication to the sender and the receiver in evolutionary game theory, theoretically discussing reciprocal altruism and spatial organisation. Finally, I explore some theories involving indirect pressure on transmission, which would explain the generation and maintenance of optimal Saussurean communication systems.

Richard.Littaur@gmail.com
Saarland University

Introduction

A Saussurean sign is a "bidirectional mapping between meaning and symbol." (Oliphant 1996, p.31) This kind of sign is manifested in human communication: a meaning in the mind of a `speaker' is mapped onto a signal, which is then sent, and the `hearer' maps the signal back onto an identical meaning, in their own respective internal language system. This is arguably the simplest form of communication – transfer of meaning from one mind to another. There is a growing body of research, called `evolutionary linguistics' that looks into how such a system might have evolved, in order to explain human language. Often, evolutionary linguists use computer models to test theories about the evolution of language, including the evolution of Saussurean signs.

Computer simulations can help where science has no direct evidence – language has left no traces, fossils, or writings to help with understanding how it might have evolved in prehistory. However, by making simple computational models, and carefully selecting variables and parameters that might help us model language simply and accurately, it is possible to look at how processes change, and to draw connections between the simulations and reality. In this vein, many researchers (such as Yanco and Stein (1992) and Werner and Dryer (1991)) have looked into how a Saussurean system – where everyone is able to map symbols onto meanings and back again – might have evolved.

Where researchers have used computational models, it is scientifically important to go beyond peer review and to seek to replicate their studies in order to check their results and analyses. Computational models are especially in need of replication, as there are normally many variables and parameters that might need to be adjusted by the researcher, as well as many places in the code where there might have been a fault, leading to errors in any analysis of the results of the model.

This paper looks specifically at such a model, from Oliphant's paper 'The Dilemma of Saussurean Communication', which is an early, seminal work that looks at language modeling. Although Oliphant's paper now has a strong literature citing it, the code he used is not available to test; replication of his paper, then, becomes an important issue, as his claims must be backed up and able to be replicated with the information he gave, for his conclusions to hold. As more literature grows in this field, it becomes increasingly important to continually check and replicate seminal pieces of work. This paper seeks to do just that.

Oliphant's Models

Following groundbreaking evolutionary simulation work on Saussurean signs by Hurford (1989), Oliphant uses computer simulations in the same way to inquire how a Saussurean system can develop. His simulations use 'agents,' which are the simplest possible representations of actual people. In essence, an agent is an entity that can make a choice or perform an action. Here, they are not full-blooded simulations of people, but rather bits of code which have a pre-decided ability to send signals (as simple as ‘0’ or ‘1’), and to receive them – recognize a ‘0’ as a ‘0’, and respond according to their pre-coded nature. For instance, one agent may be a piece of code that, when it receives a ‘0’ from another agent, responds with a ‘1’. In Oliphant’s code, the agents are hearers and speakers, both of which can ideally communicate using these signs. The agents communicate with each other in a generational model (or a genetic algorithm, following Mitchell’s (1996) work), where new agents are introduced and old agents are removed (thus mirroring population changes in human society). A genetic model means only that the agents can clone themselves to create new agents, with the same encoded information, much like in reality, where parents passes on genetic information to their offspring. This can happen sexually, where that information mixes with another agents, or asexually, where the child has the same genetic code as the parent. Oliphant’s simulations involve communication between agents with different sending and receiving mechanisms, where only those agents that manage to communicate effectively in both ways – showing the clearest case of a Saussurean bidirectional mapping – are selected for producing the new population.

His first simulation ‘rewards’, or benefits, both the sender (speaker), and receiver (hearer) of signals, such that mutualistic cooperation is optimal, meaning that if a speaker and a hearer both understand each other, then they will do better in the simulation. Their offspring will then be present in the next generation. This process simulates Darwinian natural selection, where the ‘most fit’ individuals are able to reproduce (often termed ‘survival of the fittest’ in popular culture). In this model, communication leads swiftly to a near-optimal system (meaning, an almost perfect system, where everyone understands everyone else). However, it is worth noting that there is no intentionality in this system, as an agent cannot choose to manipulate signals – for instance, by withholding information, or lying, or only giving part of the information. This means that the agent’s talked about here are not the same as in real life – however, the important feature of the agent being analysed, communication, is retained, without any other features. This is what makes the simulation a simulation, and not a direct recreation of the natural linguistic environment of humans.

There is also no `cost of signaling’, a feature absent in all of Oliphant's (1996) simulations. This concept is difficult to understand within language alone: the core concept is that in communities where individuals communicate over time, different types of signals – talking, grooming (as in apes), buying presents (as in humans) and so on all have different costs. Buying an expensive present costs the giver (in time, energy, money etc.) more than saying `happy birthday’, which goes a long way towards explaining why saying `happy birthday’, and speech in general, are used more often – it is cheap. In Oliphant’s simulations, the agents can talk forever and not get tired – there is no cost to signaling. Although technically non-communicative outcomes in his simulations, where no one understands each other, lead to a predilection to not be chosen for reproduction in the generational simulation, there are still no adverse effects to the speaker or the hearer immediately for failing to communicate. This may not be the case in real language, where a misheard or misspoken ‘look out!’ can have dire, immediate consequences.

Oliphant's first simulation was fairly straightforward. "Where the interests of the signaler and the receiver are coincident (both signaler and receiver get some benefit), the evolution of signaling seems fairly unsurprising — individuals who signal will be selected for, as will individuals who respond to signaling." (Smith 2003, p.152) However, as Smith and Oliphant go on to state, language may not act this way. They both mention alarm calls, a textbook case where the sender does not benefit from communication, as he already knows about the danger, while the hearer benefits enormously. The sender may even endanger his position, having made clear to a predator where he is, as well as using up energy, even if vocal communication is somewhat cheap.

Oliphant's second simulation therefore gives rewards only to the speaker when communication is successful, with the result that an optimal system fails to evolve. This is the expected result, as a benefit for the speaker alone would not satisfy the condition of cooperation, which he was trying to prove was a precursor to linguistic communication. If all humans were orators, but no one listened, language would soon die out in a similar fashion – no one would bother learning it, and there would be no children to learn it in any event.

The third and fourth simulations (replication of which is beyond the scope of this paper) involve Reciprocal Altruism and spatial organisation. Reciprocal altruism is where an agent does something that disadvantages himself, in expectation that the favor will be returned – this can be summed up with the idioms `tit-for-tat', or `you scratch my back, I'll scratch yours.' Spatial organisation is merely where agents only talk to other agents that are next to them, just as in small communities with their own languages. Both of these methods of choosing communication styles lead on to near-optimal systems, and are intended to show alternative justifications for communication on the part of the speaker, for whom communication in its most basic form would be a waste of resources and data. This is presented in opposition to the hearer, who would derive benefit from any information without having to give information in return. The overall goal for these simulations is to show the conditions necessary for language. The conclusion is that "there must be pressure to select for good transmission systems." `Transmission systems' here means human language. Furthermore, Oliphant assumes that such pressure is unlikely to be directly worked into the new population, since it requires effort on the part of the speaker, without providing direct benefit (de Boer 2001, p.113) and so simulation three and four are constructed to show alternative means of arriving through indirect pressure at near-optimal systems (Oliphant 1996, p.37).

Simulation

The code

The basic genetic algorithm in the replications used in this paper – a generational model where populations changing over time – reflects both Mitchell's (1996) systems for genetic evolutionary algorithms, and Oliphant's (1996) language-specific model, with a few differences.

Oliphant's agents' genomes, or genetic makeup, were binary 'strings' – such as `0101', where only two different letters or symbols are available, '0' and '1. These strings could produce a symbol in a particular environment, and a response to said symbol. The production and reception 'strings' were different within each agent, and an agent could not examine (or understand) its own system, as they are only coded to send and receive. In the simulations in this paper, instead of having strings, each agent has a send matrix and a receive matrix, such that any one meaning or signal has multiple signals or meanings corresponding to it, limited by the amount previously specified. A matrix is merely a grid-representation, meaning that, if an agent received the signal '0', it would look up the row labeled '0', and send the appropriate response (which could be '1' or another '0'). Thus, as in Tonkes et al. (1999, 2000), each individual has both signaling and receiving matrices. Each agent utilises two separate binary strings for the send or the receive matrix, the size of which was predetermined by a set number of signals and meanings. The maximum size, then, of an agent's matrix is 4 cells: two signals, and two set responses to those signals. The terminology and the code that makes up each agent is a bit different than in Oliphant, but the idea is basically the same – an agent produces a response for a given signal. Note that there is no gradation of meanings in this system – a symbol cannot be near or like another one, although it can be exactly the same. There is also no weighting of signals, nor polysemy or homonyms in this model. To reiterate, the functional differences between Oliphant's model and the one used here are negligible.

As in Oliphant (1996), there were two signals with two corresponding meanings. In Oliphant's code, an agent would be a four-letter string, while in the code presented here, it would be a two-value matrix – both encode the same internal meanings. Having only two `words' is by no means an oversimplification, as significant results can still be gathered from such a small amount of signs used in the interactions. The population size was set at 100. In each new turn, this would be replaced by the next generation of agents. The mutation rate for each new generation was 0.01 per locus (meaning that, for each possible meaning in each agent, once every hundred generations a meaning would randomly change, to simulate random change in real evolution). Crossover of genes was not used in production of a new agent –there was no sexual reproduction in this model.

A total of 5000 interactions were run for each generation between randomly chosen agents, using only horizontal communication; agents only ever spoke to agents in their own generation, never to their 'children' or 'parents'. This means that an agent would be the sender 50 times and a receiver 50 times. The initial population's matrices were randomly chosen, and subsequent generations were generated semi-randomly, with fitness – how good an agent was at speaking or hearing – being a weighting factor when choosing reproducing agents. The agents which did better were more likely to reproduce, meaning that they were more likely to have offspring based on their genome in the next generation. Fitness was determined by a ratio of successful iterations over the amount of iterations as both sender and receiver – so, an agent had to do well in both categories.

According to Smith, "The most basic models of genetic transmission are based around three processes: ontogeny, selection, and genetic transmission." (Smith 2003, p.150) Concerning ontogeny (sexual replication, in effect giving birth), neither the send- nor the receive- matrices of an agent are altered – however, for each agent, a score of how well they did in the simulation is kept in a separate list. This is altered after each communication. The values are changed if a successful meaning-signal transference is enacted while that agent is speaker or receiver, and the other values are changed for every iteration, regardless of communicative success or not. This provides an easy way of seeing the proportion of successful communications the agent engaged in.

Agents are chosen for procreation based on their performance (based on the extra list described above) – the chosen agents have their sending and receiving matrices copied, creating a new agent. Some of these copies will be mutated, as stated above. The original agent is then discarded –each generation has a complete turnover of agent population. As agents which have a better fitness score – which were better communicators – are chosen more often, many agents which do not have good fitness scores will have no offspring in the new generation. Because of this, it is harder for a single poor communicator to significantly skew future generations. However, it is worth noting here that if only transmitting correct signals, or only receiving correct signals, is the key to an agent's fitness score, the choosing function would not work as well to ensure only successful communicators procreate, as a perfectly good receiver might be bad at sending, and the model (the computer simulation being run) would have no way of knowing this, as it would only look at the score for receiving.

The total sum of all of the fitness scores of all of the agents in each generation can be used as an indicator of the frequency (per each generation) of an optimal, Saussurean system, as the average level of success for each agent can be readily seen. In the case of the replication produced here, an end result that hovers under a thousand would indicate an optimal system. A thousand is the maximal number here as each agent's fitness score ranges from 0 to 10, and there are 100 agents in each generation. Just under 1000 is actually only nearly optimal: a perfect system (called a Panglossian system by Saussure) is not necessarily the goal, as it does not reflect actual language. The reason that the result will `hover' instead of remaining at 1000 is the gene mutation, which insert randomness, and therefore wrong signals, into the agent's genomes.

The progress through time of these systems is the main area of interest, as it indicates to an extent the variable likelihood of optimal systems. By tracing a given simulation within certain parameters, how well that system would work in real life can be ascertained. The actual systems themselves are not of much use, as shown by Oliphant's findings: "In each case, the entire population quickly converged to a single transmission/reception system. Two such stable systems exist – the two Saussurean communication systems possible with this four-bit genome ("0101" and "1010").

Which of them the population converges to depends on the random seed given to the simulator." (Oliphant 1996, p. 33)

Finding that a population of agents can or cannot achieve an optimal system is interesting, because it says something about how humans themselves evolved language. By taking out many variables of human languages, and leaving just the ability to send and receive signals, it is possible to see what may or may not have happened in human evolution, as well as answer the question of why other species have not developed languages. Using a computer to speed up the process – hypothetically, these simulations could all be done by hand – allows a quick and easy way to test various hypotheses, such as 'Does it matter if a hearer understands?' or 'Does there have to be some sort of gain for the speaker in conversations to make speaking worthwhile?'

Simulation 1

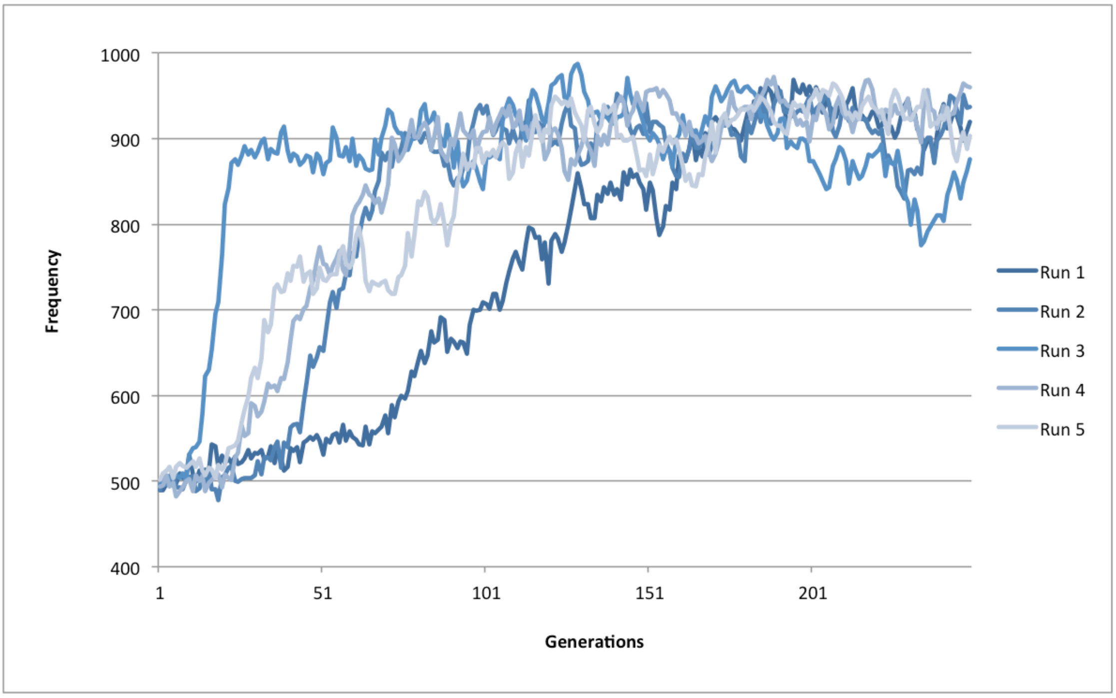

Figure 1. *Equal weighting reward for sender and receiver.*

Using the code described above, a simulation lasting 250 generations was run five times, with equal weighting for transmission/reception systems. The results are displayed in Figure 1. The frequent dips arise from genetic mutations - an algorithm that deals with mutation in a better fashion is able to rise faster, as the negative effect is less prevalent. Mutualistic cooperation (where there is equal reward for the sender and the receiver) is not meant to imply a theory of mind or altruism: merely that it is in both of the agent's best interest to communicate effectively.

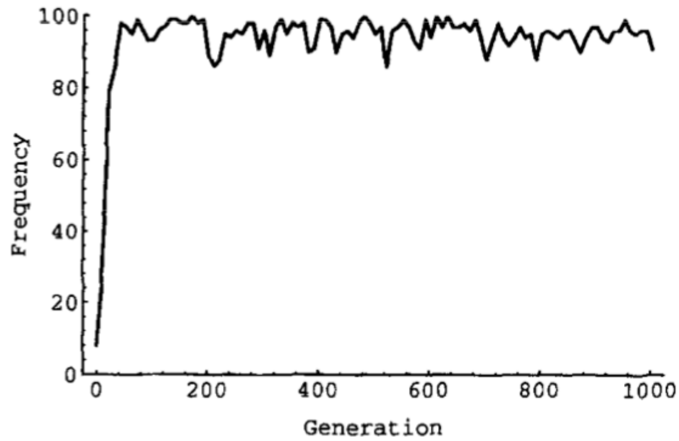

Figure 2. *Olphant's model with equal weighting reward for sender and receiver.*

The results in Figure 1 look slightly different from Oliphant's model, which can be see in Figure 2, most notably in the amount of time that it occasionally took to begin the climb to an optimal system. More will be said on this in the discussion section. As in Oliphant, the frequency amount reflects the sum of all of the agent's fitness scores.

Simulation 2

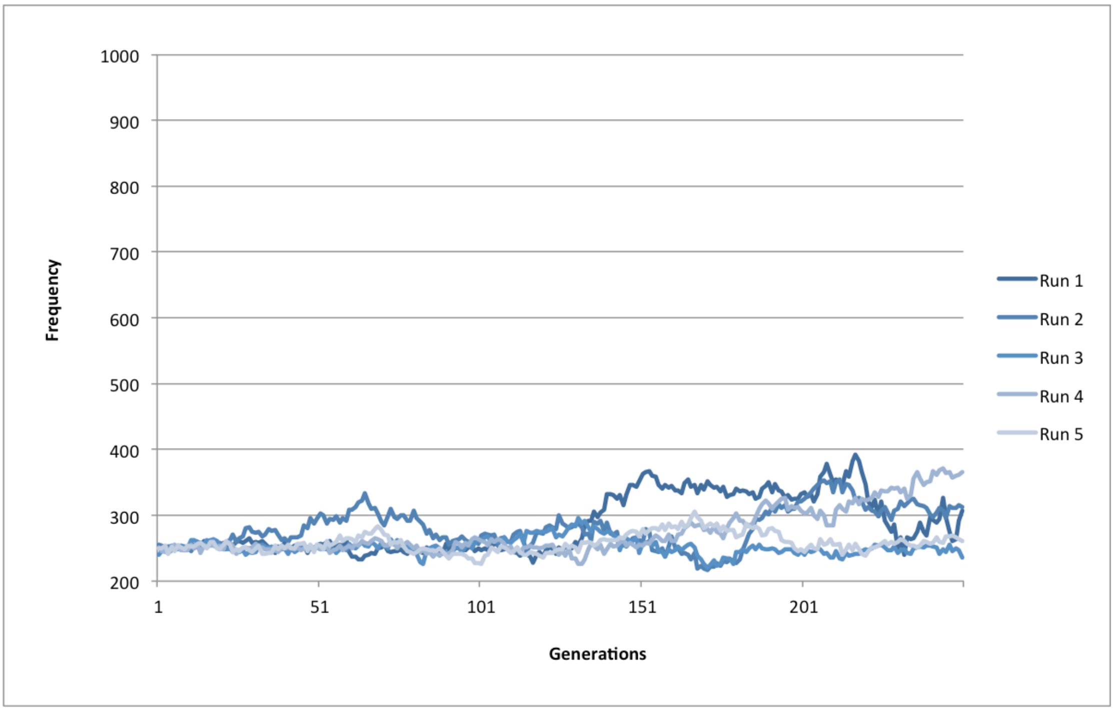

Figure 3. *Receiver reward only if transmission is correct.*

Oliphant's second simulation only rewards the hearer in correct communication. The results of the replication can be seen in Figure 3 – no optimal system emerges, although there is some small rise in comprehension for some runs, as might be expected due to the random nature of the code mutation. Oliphant claims that the reception systems organise towards a bi-stable resolution, while the transmission systems fail to converge on any particular one, due to the lack of exclusive selectivity in their reproduction. It is interesting to note that, given the identical original construction of the matrices, if the hearer is not rewarded, but the speaker is, there would be the same outcome.

Simulation 3

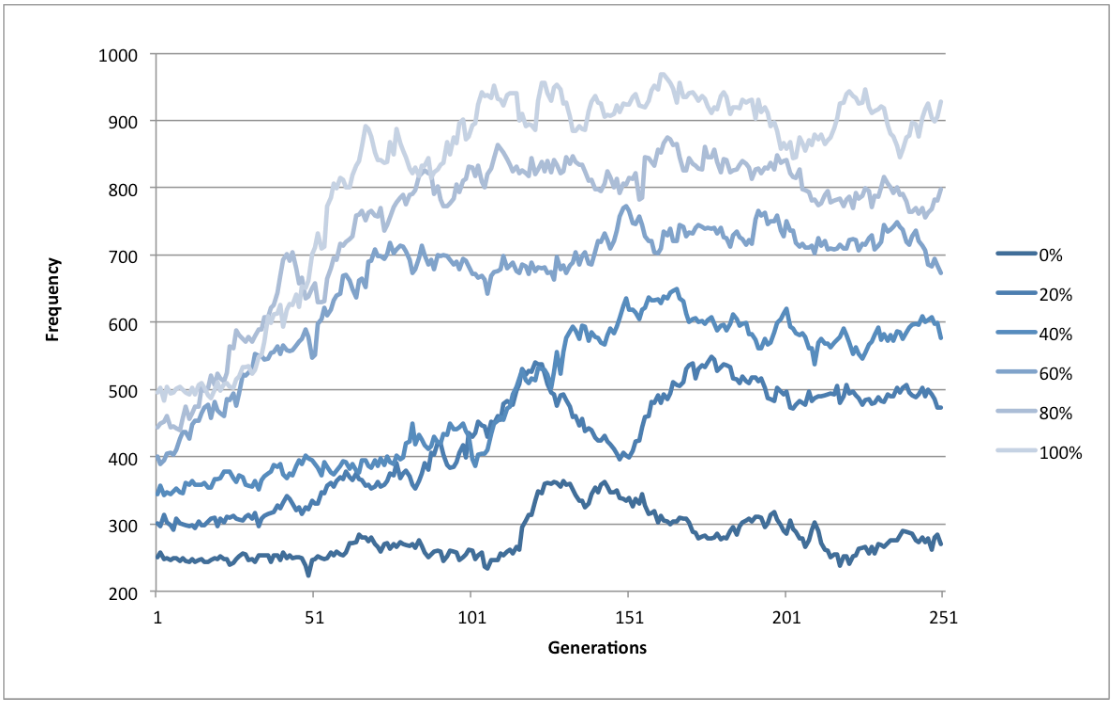

Figure 4. *Incremental reward for the sender.*

In order to understand the nature of Oliphant's first and second simulations, it is useful to look at gradations of reward. The simulation seen in Figure 4 does this, in order to see whether the stable result in the first replication and the unstable result in the second replication were results of the simplicity of the code – for instance, there are only four possible types of agents given a four-bit genome. By showing gradations of reward, the nature of the simulation itself can be understood better.

And so, in Figure 4, six lines of varying height can be ascertained. The weighting procedure for hearer and speaker fitness is changed fro each run: from the bottom, the speaker gets no benefit for good communication. However, at the next line the speaker gets twenty percent of the benefit that the hearer does – and forty and sixty and so on until 100. Here it is plain to see that any optimal solution to the problem of an altruistic speaker must confer the same benefits on him as the hearer, otherwise an optimal level is not gained. Oliphant does not run this simulation, but it does serve as a diagnostic towards understanding how the benefit system in the model directly affects optimal communication system frequency.

Discussion and Implications of Results

Oliphant's conclusion was that certain conditions must be met for the evolution and maintenance of an optimal, Saussurean system. He identified three conditions: either mutualistic cooperation, where both the sender and the receiver benefit; Reciprocal Altruism, where agents benefit each other and hope for a return of the favor; or spatial organisation, where agents talk only to their related kin. The first condition is seen to be a clear prerequisite, given the reproduced simulations above. In both the first and second replications above, the results were very similar to Oliphant's. (As stated previously, replicating Reciprocal Altruism and spatial organisation is beyond the scope of this paper.)

The discrepancy in results between Oliphant's first simulation and the replication may have been caused either by stochastic noise or by slight changes in the code that would influence the organisation of optimal systems. If the latter is the cause, then the model would be guilty of begging the question, where the results are worked into the question or the means to the answer itself. It would then not be an adequate model for language evolution. This charge has been leveled against language simulations before – for instance, having a predefined semantic space does not accurately reflect (most theories concerning) language (Gong 2009). The simplicity of the model presented here alleviates most of these concerns, but it would be fruitful (although not possible) to statistically compare the results above with Oliphant's code. This worry can be waived away without statistical comparison by looking at the results of the original simulation run in simulation three (above), where gradated levels of reward show clearly that the results are predictable and not the fault of a model, or certain set parameters. It is important to note that the replicated code in replication one and two was written solely based on the original article, and not from Oliphant's original code, which is not publicly available.

Oliphant identifies two other conditions in his paper for optimal systems, as stated above. Reciprocal Altruism (which is, in effect, cooperation) would colloquially be understood as "you scratch my back, I'll scratch yours", and needs agents to be able to retain a memory of their co-communicators, as they cannot predict cooperative or manipulative responses. Various systems can be weighed against each other for this to work, such as the so-called tit-for-tat system, where each agent remembers their interactions with other agents and cooperates with agents that have already cooperated with it (Hurford 2007). Spatial organisation, on the other hand, works within the altruistic Kin Selection theory (Trivers 1971), where agents might avoid their best interest when genetically related individuals are involved (Dawkins 1976).

However, as Hurford (2007, p. 255) states, "The central question surrounding all these theories is *cui bono?* – who gets any benefit from communication? ... We do not need to suppose that the use of language must always be either wholly for the benefit of the hearer, or alternatively wholly for the benefit of the speaker." Other alternatives to the altruistic or cooperative speaker are possible. Hauser (1996), looking at nonhuman communication system's signals, found that they focused only on three domains: mating, social interaction, and survival. As Hurford points out, Darwin spotted this first: "When we treat of sexual selection we shall see that primeval man...probably first used his voice in producing true musical cadences...this power would have been especially exerted during the courtship of the sexes, – would have expressed various emotions, such as love, jealousy, triumph, – and would have served as a challenge to rivals." (Darwin 1871) Bickerton runs with this idea, claiming that language coevolved with a previous communication system, verifying this by nature of the benefit non-human communication systems bestow on the speaker, but not on the hearer. (Bickerton 2008, p.27-28)

It is important to keep two things in mind concerning separation of production and reception systems. First, that "[a] characteristic of human language [is] that anyone can act as either producer or recipient of a message." (Hurford 2007, p.168) Second, that, among others, "Bates et al (1998) are incredulous about the possibility that linguistic knowledge is encoded in the genome [in a simulation]." (Turner 2002, p.47) This is a misunderstanding of a necessary simplification involved in linguistic simulations – the agents are not individuals themselves but rather the internal language system of an individual, represented in the brains of the population (Kirby 1999).

So, in order to answer the question of *cui bono?* (who benefits), which is the central goal of the simulations above, the view that communication is meant only for information transfer must be discarded. As Geoff Miller states, "As long as language is viewed purely in terms of information transmission, it will be seen as bringing more benefits to the listener than to the speaker." (Miller 2000, p.350-351) This view ignores the coupling of production and comprehension in a single agent, and is deficient in that it demands a temporal view that lasts no longer than a single sender/receiver interaction. Krebs and Dawkins (1984), in a similar vein to Miller, "do not define animal communication in terms of information transmissions but as a method whereby one animal exploits the muscle power of another … different kinds of communication may evolve under conditions of conflict and of cooperation." (Noble 2000, p.41-42) Taking such a Machiavellian view of speakers, where speaking's sole purpose is to make others do work at their own expense, isn't necessarily the only option, or a non-contested one. For instance, Oliphant uses Kin Selection and Reciprocal Altruism to emphasise the importance of communication beyond information transfer. Smith (2005), like Oliphant, when talking about syntactic evolution, notes that "issues of honesty, altruism and so on don't seem to have a great impact on the structure of language — people tell the truth and tell lies, and altruistically give and selfishly withhold information all the time, and it is not clear if this has any lasting structural consequences." (Smith 2003, p.153) These differing theories are not offered without merit – again, the purpose is to explore all of the theoretical options that might shed light on the evolution of optimal, Saussurean systems, where every agent ought to benefit.

There are further options to consider as possible influences, apart from Oliphant's view of entirely-receiver beneficial communication (not including Kin Selection and Reciprocal Altruism). For instance, as Noble (2000, p.41) points out, Dessalles (1998) argues that inter-group competition leads to a situation where honest information is given away freely to gain status. Here, both the speaker and the hearer benefit. Noble also points out that "Knight (1998) argues that the cooperative exchange of information that characterises speech involves a great risk of deception, and therefore that speech-like communication could only be evolutionarily stable if there was some mechanism that made it strategically sound to trust other members of the group." (Noble 2000, p.41) What this mechanism might be remains an open research question.

These views and theories must be dissected closely, just as Oliphant and Smith's concern that direct pressure may not influence the growth of optimal, Saussurean, communication systems. But there is no dearth of arguments for the benefit of communicating, and Oliphant's second simulation, where the receiver alone benefited, may not be the best representative of growth, either. Hurford sums up the argument for the speaker quite clearly: "A form of communication exists because the producer of a signal normally gets some benefit from it." (Hurford 2007, p.168) What benefit this might be is not built into the computational models provided here.

On an almost contradictory note, Oliphant's example, "if the receiver does not understand the signal, it stands a higher chance of getting eaten", stresses survival more than it might have been in a proto-linguistic setting. In this, and other ways, the model presented here is somewhat biased and overly simplified. If the success of the two simulations replicated is put aside temporarily, and what they do not take into account is considered, it can be seen that they are not clearly representative of language. Consider that there is no clear intentionality (the agents are not philosophically independent actors), no constraints on communicator selection (excepting the fitness score used in reproduction), no constraint on memory, no duality of patterning (where a whole utterance's meaning is the sum of the meaning of its parts – namely, syntax), no theory of mind (where agents can predict what the other agents will say), no cost for production, no *obverter* like mechanism (where production and reception are combined), no specified triadic communication (where two agents share focus on a third point or meaning), to name a few. Naturally, that was not the object of the simulation: the point was to minimise variables so that the conditions necessary for language growth could be manipulated and identified.

Even with this simplification in mind, the simulation, running as it does on life or death signals, fails to take into account that those are the situations in which agents would assumedly lean more towards altruistic behaviour and cooperation, as opposed to the daily grind of language use. Noble (1998) and Krebs and Dawkins (1984) go in the opposite direction by evoking costs, claiming that "where both sender and receiver receive a clear payoff then signals should be cheap, whereas if the signaler can expect to receive a large payoff but the receiver only receives a marginal benefit, receivers should be resistant to cheap signals and signalers should therefore evolve to use costly signals." (Smith 2003, p.152) The issue is not easily solved; indeed, a solution may be impossible given the simulations provided.

Conclusion

This paper described Oliphant's simulations which sought out the cooperative conditions necessary for the development of optimal, Saussurean, communication systems. It replicated his first two simulations, with new code, by separating an agent's sending and receiving matrices, and then running them with similar parameters to those which Oliphant chose, dealing first with both the sender and receiver benefiting from successful communication. Here, the same results were drawn, in that an optimal system managed to evolve and maintain itself. The same conclusion arose from the replication of Oliphant's second simulation, where only the receiver benefits; there was no optimisation of the system. The third, novel simulation showed the effects of a gradient scale of sender benefits, when the receiver still benefits. This supports Oliphant's conclusion that there needs to be mutualistic cooperation in order for an optimal, Saussurean system to occur.

Oliphant's conclusions were that three conditions must be met for the creation and maintenance of optimal, Saussurean systems: mutualistic cooperation, Reciprocal Altruism, or spatial organisation. However, in lieu of other theories about language, and in particular studies from nonhuman communication systems and manipulation studies which were discussed above, the speaker may benefit more than Oliphant originally figured into his models, as his second simulation discounted any benefit. The conclusions drawn from the models presented are also all products of the simulations they are encoded in; and it is important to remember, here and in Oliphant's studies, that simulations are not intended to be perfect models of real world linguistic systems. It is also possible that Oliphant fails to take into account a bias inherent in these simulations towards life-or-death situations; although the agents involved are not identities but language-users, the delineation between them is not clear. A caveat concerning the minimalist nature of simulations and the conclusions drawn from them must be posited.

Appendix

Acknowledgements: This paper would be nowhere near publishable if it weren't for: Simon Kirby and Andrew Smith, who assigned the original paper and taught Simulating Language well; an anonymous reviewer, who gave very useful advice; and of course the many proofreaders, who gave up their time for friendship, and hopefully not their friendship for proofreading. Any faults within the text are mine, naturally.

Supplementary Code: All of the code used in this replication, written by the author, has been uploaded freely, under no license whatsoever. Python 2.6 or higher will be needed to run it. It can be downloaded here:

https://www.github.com/RichardLitt/Publications/blob/master/oliphant.py

References

BATES, E., et al. (1998) Innateness and Mergentism IN: BETCHEL, W. and GRAHAM, G., eds. *A Companion to Cognitive Science,* Oxford, UK: Basil Blackwell, 560-601.

BiICKERTON, D. (2008) Two Neglected Factors in Language Evolution IN: SMITH, A., SMITH, K. and CANCHO, R.F., eds. *The Evolution of Language*, London, UK. The Evolution of Language 7th International Conference (EVOLANG): World Scientific, 26-33.

DE BOER, B. (2001) *The Origins of Vowel Systems*, Oxford Linguistics: Studies in the Evolution of Language, Oxford, UK: Oxford University Press.

DARWIN, C. (1871) *The Descent of Man, and Selection in Relation to Sex*, London, UK: John Murray.

DAWKINS, R. (1976) *The Selfish Gene*, Oxford, UK: Oxford University Press.

DESSALLES, J-L. (1998) Altruism, Status, and the Origin of Relevance IN: HURFORD, J.R., KNIGHT, C. and STUDDERT-KENNEDY, M. eds. *Approaches to the Evolution of Human Language: Social and Cognitive Basis*, Cambridge, UK: Cambridge University Press, 130-147.

GONG, T. (2009) *Computational Simulation in Evolutionary Linguistics: A Study on Language Emergence*, Taipei, Taiwan: Institute of Linguistics, Academia Sinica.

HAUSER, M.D. (1996) *The Evolution of Communication*, Cambridge, MA: MIT Press.

HURFORD, J.R. (1989) Biological Evolution of the Saussurean Sign as a Component of the Language Acquisition Device, *Lingua* 77, 187–222.

HURFORD, J.R. (2007) *The Origins of Meaning: Language in the Light of Evolution,* Oxford, UK: Oxford University Press.

KIRBY, S. (1999) *Function, Selection and Innateness: The Emergence of Language Universals*, Oxford, UK: Oxford University Press.

KNIGHT, C. (1998) Ritual/Speech Coevolution: a Solution to the Problem of Deception IN: HURFORD, J.R., KNIGHT, C. and STUDDERT-KENNEDY, M., eds. *Approaches to the Evolution of Human Language: Social and Cognitive Basis,* Cambridge, UK: Cambridge University Press, 68-91.

KREBS, J.R., and DAWKINS, R. (1984) Animal Signals: Mind Reading and Manipulation IN: KREBS, J.R. and DAVIES, N.B., eds. *Behavioural Ecology: An Evolutionary Approach* 2nd ed. Oxford, UK: Blackwell.

MILLER, G. (2000) *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*, London, UK: William Heinemann.

MITCHELL, M. (1996) *An Introduction to Genetic Algorithms*, Cambridge, MA: MIT Press.

NOBLE, J. (2000) Cooperation, Competition and the Evolution of Prelinguistic Communication HURFORD, J.R., KNIGHT, C. and STUDDERT-KENNEDY, M., eds. *The Evolutionary Emergence of Language*, Cambridge, UK: Cambridge University Press, 40-61.

OLIPHANT, M. (1996) The Dilemma of Saussurean Communication, *Biosystems* 37(1-2), 31–38.

SMITH, K. (2003) *The Transmission of Language: Models of Biological and Cultural Evolution*, Theoretical and Applied Linguistics, The University of Edinburgh, Edinburgh, UK.

TONKES, B., BLAIR, A., and WILKES, J. (1999) A Paradox of Neural Encoders and Decorders, or, Why Don't We Talk Backwards? IN: MCKAY, B., YAO, X., NEWTON, C.S., KIM, J.H., and FURUHASHI, T. eds. *Simulated Evolution and Learning, Lecture Notes in Artificial Intelligence* 1585, London, UK: Springer, 357-364.

TONKES, B., BLAIR, A., and WILKES, J. (2000) Evolving Learnable Languages IN: SOLIA, S.A., LEEN, T.K. and MULLER, K.R. eds. *Advances in Neural Information Processing Systems*, Cambridge, MA: MIT Press, 66-72.

TRIVERS, R.L. (1971) The Evolution of Reciprocal Altruism, *Quarterly Review of Biology* 46, 35–57.

TURNER, H. (2002) An Introduction to Methods for Simulating the Evolution of Language IN: CANGELOSI, A. and PARISI, D. eds. *Simulating the Evolution of Language*, London: Springer, 29-50.

WERNER, G. and DYER, M. (1991) Evolution of Communication in Artificial Organisms IN: LANGTON, C., TAYLOR, C., FARMER, J., and RASMUSSEN, S. eds. *Artifical Life II*, Redwood City, CA: Addison-Wesley, 659-687.

YANCO, H. and STEIN, L. (1992) An Adaptive Communication Protocol for Cooperating Mobile Robotos IN: MEYER, J., ROITBLAT, H. and WILSON, S. eds. *From Animals to Animals*, Cambridge, MA: MIT Press, 478-485.