

Speaking Our Minds is an enjoyable book, providing an excellent survey of some of the perennial and current issues in the field of language evolution, as well as providing a clear summary of Thom's position on the central role of ostensive-inferential communication in language origins. I hope neither author will mind if I say that it reminded me very strongly of Jim Hurford's recent (and rather more monumental) books on the same topic, which is perhaps not so surprising since Thom studied with Jim. It's also worth pointing out that it's a very bold book, since some of the crucial evidence that Thom would need to nail down his position isn't available - I can't help but be struck by how many cells in table 4.2 (which in reviewing the evidence for the human-uniqueness of capacities for ostensive-inferential communication is really the heart of the book's argument) are filled "Not (yet) directly studied". I hope the book stimulates some of that work.

I want to make two comments here, on sections of the book where Thom touches most closely on work that I am most familiar with, on the cultural evolution of language (Chapter 5) and (much more briefly!) how that changes our understanding of what the biological capacity for language actually is (Chapter 6).

### The role of communication in the cultural evolution of language

Chapter 5 of the book outlines Thom's account of how, once the capacity for ostensive-inferential communication was in place, symbols, compositional structure and grammatical function words might emerge from processes of language learning and language use. This chapter is framed in terms of cultural attraction, which I admit wasn't to my taste - I think it's a pity not to acknowledge the central role that Rob Boyd and Pete Richerson have played in establishing a scientific approach to studying cultural evolution, or the value of Tim Griffiths's important contribution in putting computational models of language evolution on a much sounder theoretical footing by making them Bayesian. But it's nonetheless a very useful, clear, coherent and up-to-date summary of the ongoing experimental work on the ways in which fundamental structural features of language arise from language learning and language use.

I was particularly intrigued by Thom's framing of the results of our 2008 paper (Kirby, Cornish & Smith, 2008), and the emphasis on the role of communication in that work, which coincides very nicely with my own current thinking. This is all very nicely described by Thom in the book, but to recap: in that paper we describe two experiments in which participants attempt to learn an artificial language which provides labels (typed words) for objects (coloured moving shapes), and are subsequently tested on their ability to recall labels when prompted with a shape. Participants in both experiments were arranged in transmission chains, where the first participant in each chain was trained on a random set of labels (an impossible learning task), and subsequent participants in a chain were trained on the language produced during recall by the previous participant in the chain (so the language produced during recall by participant 1 becomes the target language for participant 2, the language produced during recall by participant 2 becomes the target language for participant 3, and so on). Languages change as a result of this iterated learning process. In Experiment 1, the vanilla version of the experiment, we found that the languages rapidly became degenerate as they were passed from person to person, rapidly shedding labels and collapsing distinctions; after 10 'generations' of transmission, the languages generally had very few labels (in the most extreme case, only 2 distinct labels to describe 27 distinct pictures). These degenerate languages are easy to learn, because they are simple, but they aren't particularly language-like because they are too simple - they wouldn't allow a user of such a language to convey many (or in extreme cases, virtually any) distinctions between objects. Of course, considerations of communicative utility aren't relevant in this experiment, because the languages are never used—they simply have to be learned and recalled, and therefore evolve to maximise their learnability, at the expense of their communicative potential. We can nicely capture this result in models which assume that learners have a prior preference for simple languages, i.e. languages with a shorter

coding length; in line with established results in Bayesian iterated learning (e.g. Griffiths & Kalish, 2007) languages in transmission chains come to reflect this prior preference in learners (Kirby, Tamariz, Cornish & Smith, 2015)

In a second experiment in our 2008 paper, we attempted to block the emergence of these degenerate languages by filtering the language: if a participant produced the same label for two or more objects during recall, we would pass on only one object paired with the wannabe-ambiguous label. This reduces the viability of degenerate languages, by eliminating the best evidence for degeneracy (multiple objects paired with the same label) from the data learners see. As a result, the languages in this second experiment (usually) evolved to be simple but not degenerate: they became structured, such that labels developed a compositional morphology, where parts of each label identified the colour, shape and motion of the object the label was associated with (e.g. labels for blue shapes might begin in *l-*, labels for black shapes might begin in *ne-*; labels for bouncing shapes might end in *-plo*, labels for looping shapes might end in *-pilu*, and so on). In coding length terms, such compositional languages are more complex than degenerate languages (in order to describe them, you have to write down all the morphemes in the language and the rules of their combination - in the simplest case for our experiment, this would require a list of 9 morphemes and 1 rule of combination), but substantially simpler than random languages (which can only be described by exhaustively enumerating all objects and their associated labels, in the case of our experiment requiring a clumsy dictionary of 27 entries). While these compositional languages therefore aren't the simplest possible languages (which are degenerate: to describe a degenerate language, you just have to write down a single label that can be used to label all objects), they are robust to the filtering pressure we imposed, since they provide a distinct label for every object; coincidentally, they would also be ideal for communication, allowing the maximal number of distinctions to be conveyed, although as in Experiment 1 of that paper, considerations of communicative utility aren't actually relevant in the experiment, since the participants never use the language to communicate.

We intended this filtering procedure to be a proxy for communication: rather than languages being transmitted via a process of learning and aimless recital, they are transmitted by learners learning from examples of language use, and we might reasonably expect language use to disfavour ambiguous signals. As Thom highlights, this experiment shows the crucial role that communication plays in the cultural evolution of linguistic structure: we only see the emergence of structured languages when there are communicative considerations at play. We have recently followed up on this work by replacing the filtering proxy with a far more satisfactory model of communication - Thom kindly cited a couple of proceedings papers outlining early stages of this work, the full version of which has now been published as Kirby et al. (2015). In that work we contrasted our Experiment 1 results from the 2008 paper (learning only, degenerate languages the result) with two conditions involving pairs of participants who learn a language and then use it to communicate, taking it in turns to describe pictures for each other and to identify the pictures their partner describes. In the Chain condition of this newer experiment, each pair of participants attempted to learn the language produced during communication by a previous pair of participants, mimicking the real way in which languages are transmitted by learning from language use during communication, and replacing the transmission filter with actual communication. In the Closed Group condition, the same pair of participants learnt an initial random language and then communicated with each other using that language over and over: this language is therefore under pressure to be useful for communication, but under reduced pressure from language learning (the language is only learnt by naïve learners once, right at the start of the experiment).

The kinds of languages that emerge in these two conditions are strikingly different, and also differ markedly from the degenerate languages, which emerge in Experiment 1 of our 2008 paper. In Chains, as one would hope, we see the emergence of structured languages; this fits the results of our

filtering experiment from 2008, and shows that languages which have to be learned and used to communicate will evolve to be structured, which fits with the story Thom gives in the book. However, the languages in the Closed Group condition are different: they remain essentially holistic, being composed of a set of largely arbitrary associations between meanings and idiosyncratic signals. These holistic languages are great for communication, since every object has a distinct label, but not particularly easy to learn (since they are quite complex) – but that doesn't matter in Closed Groups, because there is not much learning by naïve individuals going on. Again, these experimental results can be nicely captured by simulation models in which we assume that learners have a prior preference for simple languages (i.e. languages with shorter coding length) and that learners avoid ambiguous utterances when communicating (using a simple Bayesian model provided by Frank & Goodman, 2012).

What's rather striking about this result in reference to Thom's thesis is that it shows that communication alone is not enough to produce structure – transmission to naïve individuals is also required, because it imposes a pressure for simplicity, and structure only emerges when pressures for communication and simplicity are both at play. I don't actually think that's problematic for Thom's position, since he is careful to describe ostensive-inferential communication as one pressure acting on languages during their cultural evolution, but for me it's important to remember that the biases of language learners plays an equally important role.

I'm also excited about the potential for these same experimental and modelling methods to address more fine-grained questions about how communication interacts with learning to shape language, which I think has the potential to speak to, and should be informed by, the issues Thom raises in this chapter and the book more generally.

Our treatment of communication in Kirby et al. (2015) is rather minimal, and probably not representative of ostensive-inferential communication in the real world: participants have to uniquely identify an object for their partner, who must pick the correct object from an array; however, the speaker doesn't know what objects the hearer's array contains, and so the best they can do is provide a label which uniquely identifies the target object among all possible objects. In other words, it offers absolutely no consideration of how context might influence the structure of languages, and in particular how ambiguity in language might be tolerated or preferred as long as it isn't detrimental to communication in context. My student James Winters is doing excellent work on his PhD looking at how reliable features of the communicative context might work their way into the structure of the linguistic system: his first paper is out now (Winters, Kirby & Smith, 2014), although too late to make it into the book (maybe the second edition!). In the experiment reported in that paper, James trained participants on an artificial language which can be used to describe distinctive 'objects' (actually, little aliens) belonging to two categories: there are 4 distinct star-shaped aliens and 4 distinct blobby aliens. Participants then take turns describing aliens for their partner, and the language they produce during interaction becomes the target language for a fresh pair of participants, as before. However, unlike in our 2015 paper, James systematically manipulates context: the speaker's task is to produce a label which will enable the hearer to identify the correct alien from a set of two aliens, and the speaker knows the context in which the hearer will be making this selection (i.e. the speaker knows which two aliens the hearer will have to discriminate between). In one condition of the experiment, the two aliens that confront the hearer on every trial are always from the same category (i.e. two star aliens on one trial, two blob aliens on another); in a second condition, they are always from a different category (one star and one blob in every context); in a third condition, they are sometimes from the same category and sometimes from different categories (a mix of the first two conditions).

Using this paradigm, James was able to show that the reliable features of the context in which communication takes place ends up shaping the structure of the languages that evolve through

interaction and transmission. In particular, in conditions where participants are only ever confronted with contexts consisting of aliens drawn from different categories (one star, one blob), the languages tend to become partially degenerate: all 4 star aliens are associated with a single label, all 4 blob aliens with another distinct label. This language looks ambiguous when taken out of context (all the star aliens have the same label), but given the context in which communication takes place it is actually unambiguous - participants are never called upon to discriminate between star aliens or between blob aliens, so the 'ambiguity' of the language is never a problem for communication. In contrast, in the condition where participants are only ever required to discriminate within-category (i.e. differentiating between star aliens, or between blob aliens), the languages tend not to develop this ambiguity - rather, the languages retain 8 distinct labels for all 8 aliens, which allows within-category discriminations to be made. Finally, in the mixed condition where participants must sometimes differentiate within-category and sometimes between-category, the languages tend to develop an elegant structure in which the labels simultaneously encode category membership and individual identity within that category—for instance, the labels for the star aliens might be “hupa”, “hepa”, “hopa” and “hapa”, where the ‘-pa’ ending conveys that these are all stars (which makes discriminating star aliens from blob aliens easy) and the first syllable encodes the identity of the individual aliens (essential if you want to discriminate between a hupa and a hapa).

I think this work gives an exciting hint of how we can start to explore how communicative context shapes language, but reading *Speaking Our Minds* has made me consider how we need to expand our experiments and models to look at genuinely ostensive-inferential acts of communication— it may be that there are relatively minor tweaks that we can perform that will make them more informative, or it may be that we need a rather more radical rethink of how we approach communication in the lab, and what kinds of linguistic adaptations we should be seeking to explain.

### Biology and culture in the evolution of language

It seems to me, based on Chapter 5, that Thom and I have very similar positions on how we should explain fundamental properties of design features of language (i.e. the fact that language is symbolic, combinatorial, compositional, employs ambiguity in just the right places etc): all these properties of language arise as a result of cultural evolution, as a consequence of people learning a language from the observable communicative behaviour of others. In Thom’s thesis all these features of language therefore follow “for free” from the crucial evolutionary breakthrough, unique to humans, of the capacity for ostensive-inferential communication. Similarly, I have previously suggested (e.g. in Smith & Kirby, 2008) that the uniqueness of human language may be due to uniquely human abilities to infer the communicative intentions of others, which is pre-requisite for the cultural transmission of meaning-signal mappings; once cultural transmission of sets of such meaning-signal pairs is possible, structure inevitably follows. These two positions seem entirely compatible, and as discussed above, I like Thom’s emphasis on the crucial role communication plays in shaping culturally-transmitted communication systems, which was absent from that 2008 paper.

Given this high level of agreement, I confess I was rather puzzled by sections of chapter 6, in which Thom takes issue with a nice quote from Simon Kirby which I would fully endorse, that “Cultural transmission ... provides an alternative to traditional ... adaptationist explanations for the properties of human language” (Kirby, Dowman & Griffiths, 2007, p. 5241, quoted in Scott-Phillips, 2015, p. 136). Thom responds that “cultural attraction does not provide an alternative to adaptationist explanations of design in nature, because this supposed contrast, between cultural evolution and natural selection, is in fact a false dichotomy” (Scott-Phillips, 2015, p. 136), then goes on to describe a co-evolutionary model (the model from Smith & Kirby, 2008) in which language evolution consists of two interacting evolutionary processes, cultural evolution of languages and biological evolution of the learning biases underpinning language learning. It’s great to see this model discussed here, but it seems rather tangential to the central claim that, as I understand it, Simon was making in that

quote: important properties of human language (e.g. the fact that language is symbolic, combinatorial, compositional, employs ambiguity in just the right places, etc.) are not biological phenomena, which need to be explained in terms of the selective advantages of these linguistic features that drive the genes ultimately coding for those linguistic features to fixation; rather, they are cultural phenomena, a consequence (through a very complex, indirect route) of the biological apparatus underpinning cultural transmission. This seems to me to be completely compatible with Thom's position, that the uniquely human adaptation behind language is the capacity for ostensive-inferential communication, and that the evolution of this capacity set in place the cultural evolutionary dynamic which lead to the emergence of languages which are symbolic, combinatorial, compositional, employ ambiguity in just the right places, etc. In other words, they both seem to agree (as do I) that we don't need to provide an evolutionary account of the fitness advantages associated with specific linguistic features, but an explanation for the evolution of the capacity for cultural transmission of communication systems in humans. In Thom's case, this boils down to providing an explanation for the evolution of the capacity for ostensive-inferential communication in humans. I might place the emphasis more on the evolution of the cognitive capacities underpinning social learning (if these are different from the capacities underpinning ostensive-inferential communication, which is something I'd like to think about more), Simon might place the emphasis elsewhere, but in all three cases there is agreement that the uniquely human biological capacity for language consists of some potentially quite high-level, abstract capacities, rather than a very detailed specification of the details of the structure of human language. We all therefore seem to be in agreement about the types of capacities, which evolutionary accounts of language have to explain; given that I'm not sure this is still a mainstream position, I think it makes sense to emphasise the commonalities between these closely-related accounts, rather than dwell on the differences of emphasis.

## References

- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336, 998.
- Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with Bayesian agents. *Cognitive Science*, 31, 441-480.
- Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *PNAS*, 105, 10681-10686.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015) Compression and communication drive the evolution of language. *Cognition*, 141, 87-102.
- Smith, K., & Kirby, S. (2008). Cultural evolution: implications for understanding the human language faculty and its evolution. *Philosophical Transactions of the Royal Society B*, 363, 3591-3603.
- Winters, J., Kirby, S., & Smith, K. (2014). Linguistic systems adapt to their contextual niche. [Language and Cognition](#)