How language learning and language use create linguistic structure

Kenny Smith

Centre for Language Evolution, School of Philosophy, Psychology & Language Sciences,
University of Edinburgh, UK

Kenny.smith@ed.ac.uk

Accepted for publication in *Current Directions in Psychological Science*

Abstract

Languages persist through a cycle of learning and use: you learn a language through immersion in the language used in your linguistic community, and in using language to communicate you produce further linguistic data which others might learn from in turn. We know that languages change over historical time as a result of errors and innovations in these processes of learning and use; this paper reviews experimental and computational methods which have been developed to test the hypotheses that those same processes of learning and use are responsible for creating the fundamental structural properties shared by all human languages, including some of the design features that make language such a powerful tool for communication.

Introduction

Human languages persist through the cycle of learning and use: you learn a language through immersion in the language used in your linguistic community, and in using language to meet your communicative goals and get stuff done in the world you produce further linguistic data which others in turn might learn from (Figure 1). Language learning and language use are inherently creative processes: when learning we extend and sharpen incipient patterns in the linguistic data we encounter, and in use we push constantly at the expressive envelope provided by our language, finding ways to express the concepts and distinctions we care about conveying.

As a result of these creative processes, everyone deviates subtly from the language of their parents and peers, and passes on a slightly altered version of the language to the next generation. Anyone who studied Shakespeare or Chaucer in school will of course be familiar with the idea that languages change over time as a result - changes in the words we use and the rules for combining them make texts composed even a few hundred years ago quite strange to modern ears. In this paper I'll review a body of work which uses experimental and computational methods to demonstrate how this cycle of learning and use not only leads to language change on a historical timescale, but might also be responsible for creating the fundamental structural properties shared by all human languages, including some of the design features that make human language such a powerful tool for communication.

Figure 1: Languages are passed from person to person via a cycle of learning and use. Individuals are exposed to linguistic data, and through a process of learning develop a grammar, a mental representation of how the language works. They then use that grammar to meet their communicative goals in the world; as a side-effect of using language to fulfill their communicative needs, they generate linguistic data which forms the basis for language learning in others.

Language design

Communication is widespread in the natural world, but human language stands out, due to the open-ended expressive power it provides: any thought we can entertain in our heads can be shared through language. Language achieves this expressive power through a double layer of combinatorial structure (see Figure 2 and also e.g. Hockett, 1959). In the first layer of combinatorial structure, meaningless units (contrastive speech sounds in spoken languages, or parameters like handshape and location in sign languages) are combined and recombined to produce a far larger set of meaningful units (words and the building-blocks of words known as morphemes, for example the past tense marker written -ed in English). In the second layer of combinatorial structure those meaningful morphemes and words are combined and recombined in various permutations to produce an unbounded set of complex signals which convey complex meanings. This double layer of combinatorial structure is an extraordinarily powerful means of building a flexible communication system from a relatively meagre set of articulatory resources: a small inventory of meaningless contrastive units can yield a much larger set of meaning-bearing units, which can in turn yield an infinite set of well-formed multi-word expressions. For instance, most varieties of English have around 40 contrastive speech sounds, the average English speaker knows tens of thousands of distinct words, and there is in principle no limit on the number of meaningful sentences which any English speaker can produce or

understand; while other languages might have fewer or more contrasting speech sounds (e.g. Rokotas, spoken in New Guinea, might have as few as 11; Archi, spoken in the South Dagestan region of Russia, might have as many as 90), they all have infinite expressive power thanks to this double layer of combinatorial structure.

These two combinatorial processes operate according to different sets of principles, as shown in Figure 3. The combinatorial process by which morphemes and words are built from meaningless contrastive units simply uses differences in the component parts to differentiate one word from another: the English words “log”, “dog”, “dig”, and “dim” differ minimally in the elements they are composed from, but have completely different meanings which can’t be predicted from their component sounds or letters. As a result it doesn’t make sense to ask what the meaning of *l* or *d* is in English. In contrast, the second combinatorial process by which morphemes and words are combined to form complex expressions is **compositional**: the meaning of a complex expression depends on the meaning of its parts and the way in which they are combined ensuring (as shown in Figure 3) that changes to the building blocks making up a complex expression have discrete and predictable effects on the overall meaning. It’s this compositional structure that makes the infinite combinatorial potential provided by language’s double layer of combinatorial structure useful for communication: as long as you know the meaning of the building blocks and the semantic effects of combination, thanks to compositionality you can seamlessly produce and interpret completely novel sentences, to convey completely novel meanings.

Figure 2: Language’s double layer of combinatorial structure allows us to go from a small inventory of meaningless contrastive elements (e.g. speech sounds, shown here using spoken English transcribed in the International Phonetic Alphabet as an example) to a much larger set of meaning-bearing units (morphemes and words), which can in turn yield an infinite set of well-formed multi-word expressions.

Non-compositionality in English

Word	Meaning
log	"Noun; an unhewn portion of a felled tree"
dog	"Noun; A domesticated carnivorous mammal"
dig	"Verb; To work in making holes or turning the ground"
dim	"Adjective; Faintly luminous"

Compositionality in English

Sentence	Meaning
A dog spots a dog	$dog'(x) \ \& \ dog'(y) \ \& \ spots'(x,y)$
A dog spots a cat	$dog'(x) \ \& \ cat'(y) \ \& \ spots'(x,y)$
A dog bites a cat	$dog'(x) \ \& \ cat'(y) \ \& \ bites'(x,y)$
A cat bites a dog	$cat'(x) \ \& \ dog'(y) \ \& \ bites'(x,y)$

Compositionality in Hungarian

Phrase	Meaning (English gloss)
a táská-n	"on the bag"
a táská- ban	" in the bag"
a zongorá-ban	"in the piano "
a zongorá- k -ban	"in the pianos"

Figure 3: An illustration of the distinct properties of the two combinatorial processes of language. In the first layer of combinatorial structure (top panel), illustrated here with English written words, meaning-bearing morphemes and words are built from meaningless units; differences in the meaningless building blocks simply serve to create distinct meaning-bearing units (changes in word form and corresponding change in meaning highlighted in red). In a compositional system (lower panels), changing one morpheme or word in a complex expression, or their order of combination, results in predictable changes in the meaning. Note that this is true both of English (lower left) and e.g. Hungarian (lower right), even though Hungarian conveys concepts using morphology (e.g. in vs on) which are conveyed by combining separate words in English.

As far as we know, this characteristic double layer of combinatorial structure is not seen in the communication system of any other species on Earth. At some level the uniqueness of human language must have its basis in unique features of human biology, and some authors have attempted to build an evolutionary explanation which links these kinds of structural features of language to evolved innate predispositions (e.g. Pinker & Bloom, 1990). I'll instead focus on work which seeks to understand what features of language can be explained "for free" as a consequence of the processes of learning and use that we see operating in the historical record and in the present day. Figuring out which features of language can be explained in terms of these known processes is parsimonious, because we don't have to appeal to any additional harder-to-evidence evolutionary mechanisms; furthermore, understanding how learning and use produce language-like systems allows us to home in on the features of human cognition and sociality that are responsible for our uniquely powerful communication system.

How learning and use produce structure: compositionality as a test case

Testing this hypothesis is challenging, and data from multiple sources and methods can provide valuable insights. Analysis of corpora from multiple languages can reveal recurring principles in how languages are organised: for example, the structure of lexicons in natural languages, the shape and composition of linguistic categories for talking about things like colour, number and kin, and the syntactic rules governing the orders in which words can be combined, all seem to be designed to make language easy to learn and useful for communication (see Kemp et al, 2018; Gibson et al. 2019 for review). However, these corpus methods focus on the outcomes of the processes of learning and use, rather than the processes themselves, and therefore cannot demonstrate mechanistically **how** this optimal design comes about. Observing the emergence

of new sign languages in deaf communities (see e.g. Meir et al. 2012) allows us to witness the creation of a linguistic system in real time, but lacks the control required to disentangle competing explanations for the crucial factors shaping those new languages. Methods from historical linguists provide valuable insights into how pressures in learning and use can reshape languages, including processes where new grammatical patterns are created and extended through processes of reanalysis and analogy (see Trask, 2015 for an introduction), but are restricted to studying languages from the recent past, which share the same fundamental structural features as present-day languages. An important innovation has therefore been to augment these other approaches with experimental and computational methods to test hypotheses about how learning and use shape **artificial** linguistic systems, in particular combining **artificial language learning** methods to study biases in learning, **communication games** to study how languages are reshaped in use, and **iterated learning** paradigms to simulate multi-generational transmission. The key finding of this work has been that **structure emerges as a result of the right blend of pressures from learning and use**; in this section I'll show how this argument works for compositional structure, a key feature of language design that, as described above, is central to the expressive power of human languages.

In an artificial language learning experiment, participants attempt to learn and reproduce a miniature linguistic system: for example, participants might be asked to learn novel labels for novel objects or coloured geometrical shapes, and then subsequently attempt to produce their own labels when prompted with more objects/shapes. Artificial language learning has a long history as a tool for studying learning in both adults and children (see e.g. Culbertson & Schuler, 2019 for review), and allows us to explore how learners deviate from their data during learning. To extend this technique to study the cycle of learning and use shown in Figure 1, we can simply connect up a series of participants, using the language (re)produced by one participant as the learning input for the next participant in a chain of transmission, meaning that any errors or innovations in learning can accumulate over 'generations' of transmission. This method is sometimes known as *iterated learning*, and has a long history both in the psychological sciences (used e.g. by Bartlett, 1932 to study biases in memory) and as a children's game (where the accumulation of errors in a whispered message is exploited for comic rather than scientific effect).

Figure 4 depicts a simple iterated artificial language learning experiment from Kirby, Cornish & Smith (2008), where participants were asked to learn and reproduce labels for colored moving shapes. Participants were asked to learn a target language through exposure to a set of picture-label pairs, then tried to reproduce that language in a simple recall task, with the language produced by participant 1 providing the input to participant 2 and so on in a series of simulated generations of transmission. Since we were interested in whether compositionality would develop through this cycle of learning and reproduction, the first participant in each chain was trained on an artificial language which intentionally lacked compositionality (known as a **holistic** language, since each label is associated holistically with the picture it labels, rather than standing in a part-to-part compositional relationship). Crucially, in this experiment participants were simply asked to reproduce the labels they were trained on, rather than using the language communicatively. In this simple kind of learning-and-recall task the primary bias shaping the evolving artificial languages is therefore a preference for *simplicity*. Biases in perception and learning in multiple domains favour simpler solutions (e.g. Feldman, 2016); languages that are simpler (that have fewer labels, where those labels pick out related groups of pictures which share e.g. the same shape) are easier to learn than more complex languages (with more labels, or where those labels pick out random collections of pictures with no shared features); furthermore, mistakes participants make during learning and recall (e.g. forgetting a

label or extending the label for one picture to another similar picture) tend to make the languages simpler. As a result, Kirby et al. (2008) found that languages transmitted through this learning-and-recall cycle rapidly lost labels as they were passed from person to person, eventually ending up with a small number of highly ambiguous labels (see e.g. Silvey et al., 2015 for similar results). Simplicity facilitates learning and recall, but leads to a loss of *expressivity*, the ability to linguistically encode distinctions.

Figure 4: An illustration of the iterated artificial language learning method and indicative results, from Kirby, Cornish & Smith (2008), Experiment 1, Chain 4. Participants were asked to learn labels for coloured moving shapes (there were 3 shapes, 3 colours, 3 motions: only some are shown here). In this experiment, where the recall task is non-communicative, the language rapidly loses the ability to encode distinctions: in the initial language each picture is associated with a unique label, whereas by generation 10 the language shown here no longer codes distinctions of shape or colour among looping objects (anything looping is labelled “poi”), and within bouncing objects the shape distinction is still encoded but colour is not (bouncing circles are labelled “miniku”, bouncing triangles are “tupin”, regardless of colour). Other chains in this experiment ended up with even simpler languages, with as few as 2 labels.

Compositionality therefore doesn’t emerge for free from the iterated learning paradigm shown in Figure 4, because the simplest possible systems win out. But those learning-and-recall experiments lack a crucial ingredient operating in natural languages: *communication*. Languages are of course used to communicate our knowledge, desires, thoughts and feelings, and communication requires the ability to encode distinctions. Figure 5 shows a modified experimental paradigm from Kirby et al. (2015), where we added in communication: Pairs of participants learnt a miniature language, then used that language to play a communication game, taking turns to label pictures for each other and select the correct picture from an array of possibilities based on the label provided by their partner. The language produced during this communication game by one pair then provides the training input to the next pair in a chain, bringing us closer to the cycle of learning and communicative use operating on natural languages.

Simple compositional systems gradually develop over generations in these experiments: languages emerge where each label consists of several meaningful parts, where each feature of the picture being described is encoded in a separate part of the label (e.g. the first part of the label might encode shape and the second part might encode colour). These compositional languages allow their users to encode the full set of distinctions they require; they are more complex than the small sets of labels that develop in learning-and-recall experiments like the

one shown in Figure 4, but are simpler than holistic systems, since they contain rules and regularities that learners can extract to form a compact representation of the language. Compositional languages therefore represent the optimal trade-off between the two pressures for simplicity and expressivity that arise from learning and communicative use respectively - they are the simplest of the expressive systems, and the most expressive of the simple systems. Other experimental studies simulate a pressure for expressivity without including actual communication, for example by simply deleting ambiguous labels from the data passed to the next generation (in Kirby et al. 2008 Experiment 2 or Beckner et al. 2017), with similar results.

Figure 5: An illustration of the combined learning and communication paradigm, and indicative results, from Kirby, Cornish & Tamariz & Smith (2015), Chain Aa. In this experiment, participants were asked to learn labels for patterned shapes (there were 3 shapes and 4 fill patterns; only some are shown here). The language does not degenerate as in the learning-and-recall experiment illustrated in Figure 4 but rather gradually develops a simple compositional structure: in the chain shown here, by generation 6, each label consists of two meaning-bearing morphemes: a prefix specifying shape (“ege-“ for the kidney bean shape, “mega-“ for the rounded star, “gamene-“ for the spiky leaf) and a suffix specifying fill pattern (“-wawu” for black, “-wawa” for checked, “-uwu” for spotted), with the meaning of each morphologically-complex word depending on the meaning of its component parts in a compositional fashion.

While in Kirby et al. (2015) we emphasized the role of both communication and transmission to naïve learners in driving the emergence of compositionality, other experimental work has shown that structure can develop over the course of interaction between a single pair of participants (i.e. without inter-generational transmission), if those participants are forced to extend their language and generalize it to convey new meanings (e.g. Selten & Warglien, 2007; Winters et al., 2018) or to communicate with new partners (Raviv et al., 2019). In the real world, learning and use are of course not as clearly separated as depicted in Figures 1, 4 and 5, and even in our experiments participants presumably learn more about the artificial language and the way their partner uses it as they interact. Generalization to new contexts therefore allows the simplicity preferences inherent in learning to have an effect in shaping participants’ productions during interaction; as a result, with a strong enough generalization pressure, compositional structure can develop during communication within a pair or small group, as a result of the same trade-off between simplicity and expressivity. An exciting related line of work shows how social factors operating during communication (for example, the desire to signal affiliation with certain linguistic groups), can counteract the pressures for simplicity and expressivity, leading languages to retain communicatively-inefficient features which would

otherwise be lost, if those features are associated with desirable social groups (e.g. Roberts & Fedzechkina, 2018). Real languages used in real communities are presumably subject to the same complex interplay of partially-conflicting pressures, and extending our experimental techniques to understand how social factors interact with biases for simplicity and expressivity is an important avenue of research.

Other modalities, other kinds of language structure

The experiments reviewed above use written labels, largely for experimental convenience. Perhaps surprisingly, experiments involving iterated learning with *spoken* miniature languages (a natural modality for language) are rare. Several studies have looked at the evolution of compositionality in the manual modality (another natural modality for human language, used by sign languages), using paradigms where participants communicate using their hands and bodies (e.g. Nölle et al., 2018; Bohn et al., 2019; Motamedi et al., 2019). These gesture-based experiments match the written experiments in showing that compositional structure develops when both learning and use are at play.

Similar experimental methods have also been used to look at how other important structural properties of natural languages emerge from learning and use. The first layer of combinatorial structure in human language (recall Figure 2), termed *duality of patterning* by Hockett (1959), is particularly interesting since, like compositionality, it seems to represent an optimal trade-off between expressivity (allowing many contrastive combinations to be built) and simplicity (distinct words are built of a smaller and more manageable inventory of reusable parts, providing regularities that learners can extract), and like compositionality it can emerge spontaneously from learning and use under the right mix of pressures for simplicity and expressivity (see e.g. Verhoef et al., 2014; Little et al., 2017; Roberts et al., 2015).

The same methods have also been applied to the study of the emergence of *arbitrariness*, another of Hockett's design features of language. Morphemes and words in language are broadly speaking arbitrary: the word “dog” does not look, sound or smell like the animal it refers to, and words and morphemes have the meanings they do purely by convention, rather than through any resemblance-based or *iconic* association with the meanings they convey (although there are some important exceptions: see Dingemanse et al., 2015). Arbitrary communicative symbols can develop from iconic origins when participants repeatedly learn and use signals in unfamiliar communicative mediums, including gesture (Bohn et al., 2019; Motamedi et al., 2019) and drawing (e.g. Fay et al., 2018; Caldwell & Smith, 2012), becoming simplified in form and eventually decoupled entirely from their iconic beginnings.

Finally, biases in learning and use may be responsible for other fine-grained properties of linguistic systems. For example, languages often offer their users multiple ways to express a given concept or function: in English, the first person pronoun sometimes appears as *I* and sometimes *me*; roughly the same event can be described as *Jessie gave Sam an apple* or *Jessie gave an apple to Sam*; the pronunciation of English *-ing* (as in *finding*, *sleeping*) takes one of two forms which can be conveyed in writing as *finding* vs *findin'*, *sleeping* vs *sleepin'*. However, these options aren't deployed randomly, but are instead predictable and conditioned on the context in which they are used: we use *I* when it is a subject (as in *I like bikes*) and *me* when it is an object (as in *They don't like me*); the choice between *Jessie gave Sam the apple* and *Jessie gave the apple to Sam* is probabilistically conditioned on parameters such as the novelty of *Sam* and *the apple* in the conversation to this point; variation in the form of *-ing* depends on the formality of the situation and the gender and social status of the speaker. Such conditioned and predictable systems of variation are ubiquitous in natural languages, and can

be shown to develop through the same processes of learning and use that produce compositionality, duality of patterning, and arbitrariness (see Smith et al., 2017, for review of experimental work on this topic).

Other learners

Most of these experiments use adult participants who know at least one natural language. This makes it hard to rule out the possibility that the emergence of these various features might depend on participants' sophisticated pre-existing knowledge of the structure of linguistic systems, in which case they can't really explain the **origins** of that same structure. It is therefore crucial to test whether these results generalize beyond adult human participants, for example by replacing adults with computationally simulated participants, with children, or with non-human animals.

There is a long tradition of simulating iterated learning (e.g. see e.g. Kirby et al., 2014 for review), substituting human participants for simulated language learners/users. For instance, Kirby et al.'s (2015) experimental results are accompanied by a computational model whose results match the results of the human experiments; furthermore, the simulations show that these results hold under fairly minimal assumptions about the biases learners bring to the task. The model assumes that languages which are simpler (in the formal sense of being representable in a compact fashion) are easier to learn, and that ambiguity is avoided or penalized somehow during communication; the combination of these two biases is sufficient to produce compositionality, suggesting that the experimental results might not be dependent on particularly sophisticated structure-building or pragmatic capacities of human participants.

Several papers have used iterated learning and communication paradigms with children; while even quite young children obviously have detailed knowledge of the structure and function of linguistic communication, it should at least be less developed than in adult participants. Adapting experimental paradigms designed to work with adults for children is challenging, and early results have been somewhat mixed: in written or auditory paradigms where adults would introduce compositional structure, children tend to converge on simple but relatively inexpressive systems (Raviv & Arnon, 2018; Kempe et al., 2019); however, in gestural communication, child participants generate compositionally-structured gestures to convey complex events (Bohn et al., 2019), suggesting that the creation of compositional structure is not dependent on an adult appreciation of the demands of communication or the structure of natural languages.

Finally, there is exciting potential to adapt these experimental techniques for use with non-humans, to entirely avoid participants with pre-existing linguistic knowledge. A number of studies have tested the ability of other species to learn language-relevant regularities from data, in paradigms that resemble human artificial language methods (see e.g. Milne et al., 2018 for review). A separate strand of work with non-human animals has studied iterated learning using either non-linguistic tasks (e.g. iterated learning of sets of visual patterns in Guinea baboons in Claidiere et al., 2014) or natural communicative behaviours (iterated learning of song in songbirds, Feher et al., 2009). Bringing these two methods together, i.e. combining an artificial language learning method designed for animals with iterated learning, would allow us to explore whether or not linguistic structure develops through iterated learning in other species as it does in humans, allowing us to home in on what unique features human cognition contributes to that process by identifying what needs to be added in order to see the same outcomes in non-humans. For example, in Claidiere et al. (2014), we provided the baboons with a visual pattern learning task which was complex enough to exhibit interesting structural

features, and we also facilitated the iterated learning process for them, using a computer-mediated experimental task to allow one animal's responses to be passed to another animal for reproduction; providing similar scaffolding for communicative interaction in other species might be particularly challenging if they lack the same cooperative and communicative motivations and capacities as humans (as argued e.g. in Tomasello, 2008).

Summary

In recent years a suite of experimental and computational methods have been developed to test hypotheses about how languages are shaped by their learning and use, with exciting potential to be extended to other species to test the cognitive and social prerequisites for the creation of linguistic structure. The artificial systems that develop in these experiments are obviously much simpler and less expressive than full-fledged human languages - they can be learned in a few minutes, rather than requiring years of immersion, and can only be used to encode a handful of distinctions, rather than offering the open-ended expressivity of natural languages. But these experiments provide empirical support for the hypothesis that at least some of the fundamental design features of natural language can be explained as a consequence of biases in the process of learning and use by which natural languages persist.

Recommended readings

- Christiansen, M. H., & Chater, N. (2008). Language as shaped by the brain. *Behavioral and Brain Sciences*, 31, 489–509. [Review paper setting out the general argument that language is shaped by learning and processing biases.]
- Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with Bayesian agents. *Cognitive Science*, 31, 441-480. [Seminal paper using Bayesian modelling to explain how learner biases influence language evolution via iterated learning.]
- Kemp, C., Xu, Y., & Regier, T. (2018). Semantic Typology and Efficient Communication. *Annual Review of Linguistics*, 4, 109-128. [Review of how the tradeoff between simplicity and expressivity shapes the category systems of natural languages]
- Tamariz, M. (2017). Experimental Studies on the Cultural Evolution of Language. *Annual Review of Linguistics*, 3, 389-407. [Provides a more detailed review of experimental methods.]
- Fehér, O., Wang, H., Saar, S., Mitra, P. P., & Tchernichovski, O. (2009). De novo establishment of wild-type song culture in the zebra finch. *Nature*, 459, 564-568. [Seminal study on iterated learning in songbirds.]

References

- Bartlett, F. C. (1932). *Remembering*. Oxford, UK: Macmillan.
- Beckner, C., Pierrehumbert, J. B., & Hay, J. (2017). The emergence of linguistic structure in an online iterated learning task. *Journal of Language Evolution*, 2, 160–176.
- Bohn, M., Kachel, G., & Tomasello, M. (2019). Young children spontaneously recreate core properties of language in a new modality. *Proceedings of the National Academy of Sciences*, 116, 26072-26077.
- Caldwell, C. A., & Smith, K. (2012). Cultural evolution and the perpetuation of arbitrary communicative conventions in experimental microsocieties. *PLoS ONE*, 7, e43807.
- Claidière, N., Smith, K., Kirby, S., & Fagot, J. (2014). Cultural evolution of a systematically structured behaviour in a non-human primate. *Proceedings of the Royal Society B*, 281, 20141541.
- Culbertson, J., & Schuler, K. (2019). Artificial Language Learning in Children. *Annual Review of Linguistics*, 5, 353-373.
- Dingemanse, M., Blasi, D. E., Lupyan, G., Christiansen, M. H., & Monaghan, P. (2015). Arbitrariness, iconicity and systematicity in language. *Trends in Cognitive Sciences*, 19, 603-615.
- Fay, N., Walker, B., Swoboda, N. & Garrod, S. (2018). How to Create Shared Symbols. *Cognitive Science*, 42, 241-269.
- Fehér, O., Wang, H., Saar, S., Mitra, P. P., & Tchernichovski, O. (2009). De novo establishment of wild-type song culture in the zebra finch. *Nature*, 459, 564-568.
- Feldman, J. (2016). The simplicity principle in perception and cognition: The simplicity principle. *Wiley Interdisciplinary Reviews: Cognitive Science*, 7, 330–340.
- Gibson, E., Futrell, R., Piantadosi, S. T., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How Efficiency Shapes Human Language. *Trends in Cognitive Science*, 23, 389-407.
- Hockett, C. F. (1959). Animal "languages" and human language. *Human Biology*, 31, 32-39.
- Kemp, C., Xu, Y., & Regier, T. (2018). Semantic Typology and Efficient Communication. *Annual Review of Linguistics*, 4, 109-128.
- Kempe, V., Gauvrit, N., Gibson, A. & Jamieson, M. (2019). Adults are more efficient in creating and transmitting novel signalling systems than children. *Journal of Language Evolution*, 4, 44-70.
- Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: an experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences, USA*, 105, 10681-10686.
- Kirby, S., Griffiths, T. L., & Smith, K. (2014). Iterated learning and the evolution of language. *Current Opinion in Neurobiology*, 28, 108-114.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87-102.
- Little, H., Eryilmaz, K., & de Boer, B. (2017). Signal dimensionality and the emergence of combinatorial structure. *Cognition*, 168, 1– 15.
- Meir, I., Sandler, W., Padden, C., & Aronoff, M. (2012). Emerging Sign Languages. In *The Oxford Handbook of Deaf Studies, Language, and Education, Vol. 2* (Eds. M. Marschark & P. E. Spencer). Oxford: Oxford University Press.
- Milne, A. E., Wilson, B., & Christiansen, M. H. (2018). Structured sequence learning across sensory modalities in humans and nonhuman primates. *Current Opinion in Behavioral Sciences*, 21, 39-48.
- Motamedi, Y., Schouwstra, M., Smith, K., Culbertson, J., & Kirby, S. (2019). Evolving artificial sign languages in the lab: from improvised gesture to systematic sign. *Cognition*, 192, 103964.

- Nölle, J., Staib, M., Fusaroli, R., & Tylén, K. (2018). The emergence of systematicity: How environmental and communicative factors shape a novel communication system. *Cognition*, 181, 93-104.
- Pinker, S., & Bloom, P. (1990). Natural language and natural selection. *Behavioral and Brain Sciences*, 13, 707-784.
- Raviv, L., & Arnon, I. (2018). Systematicity, but not compositionality: Examining the emergence of linguistic structure in children and adults using iterated learning. *Cognition*, 181, 160-173.
- Raviv, L., Meyer, A. S., & Lev-Ari, S. (2019). Larger communities create more systematic languages. *Proceedings of the Royal Society B: Biological Sciences*, 286(1907): 20191262.
- Roberts, G. and Fedzechkina, M. (2018) Social biases modulate the loss of redundant forms in the cultural evolution of language. *Cognition*, 171, 194-201
- Roberts, G., Lewandowski, J., & Galantucci, B. (2015). How communication changes when we cannot mime the world: Experimental evidence for the effect of iconicity on combinatoriality. *Cognition*, 141, 52-66.
- Selten, R., & Warglien, M. (2007). The emergence of simple languages in an experimental coordination game. *Proceedings of the National Academy of Sciences, USA*, 104, 7361-7366.
- Silvey, C., Kirby, S., & Smith, K. (2015). Word meanings evolve to selectively preserve distinctions on salient dimensions. *Cognitive Science*, 39, 212-226.
- Smith, K., Perfors, A., Fehér, O., Samara, A., Swoboda, K., & Wonnacott, E. (2017). Language learning, language use, and the evolution of linguistic variation. *Philosophical Transactions of the Royal Society B*, 372, 20160051.
- Tomasello, M. (2008). *Origins of Human Communication*. Cambridge, MA: MIT Press.
- Trask, R. L. (2015). *Trask's Historical Linguistics*, 3rd edition. Edited by R. McColl Millar. London, UK: Routledge.
- Verhoef, T., Kirby, S., & de Boer, B. (2014). Emergence of combinatorial structure and economy through iterated learning with continuous acoustic signals. *Journal of Phonetics*, 43, 57-68.
- Winters, J., Kirby, S., & Smith, K. (2018). Contextual predictability shapes signal autonomy. *Cognition*, 176, 15-30.