The Emergence of Compositionality: Insights from Artificial Language Learning and Computational Modelling.

Leonidas Mylonadis

MA General Linguistics, 13863134

Supervisors: Prof. Enoch O. Aboh, Prof P. G. Paul Boersma

1 INTRODUCTION

Languages vary immensely across the globe and have done so throughout time: many sounds produced by a speaker of a Tamazight language may seem impossible to a British English speaker to recreate; a speaker of Latin would happily place the verb after the object of the sentence, while this arrangement would create undue emphasis of the object to a modern Italian speaker. Although languages differ in their phonology, morphological structure, syntax and pragmatics, it has been posited that all human, natural languages have some 'universal' factors in common. Many of these have been identified more accurately as macroscopic tendencies, such as word order harmony (Greenberg, 1963) or the case marking-word order trade-off (Comrie, 1989). One distinguishing factor of human communication, however, has been argued to appear in every natural language: compositionality.

Compositionality refers to the hierarchical structure of language and how it feeds into meaning: linguistic units tend to be composed of smaller parts. Phrases are often formed of words or morphemes and morphemes are often composed of multiple sounds. Moreover, these smaller units are combined and re-combined in new combinations to produce new meanings, often through a process known as recursion: a verb phrase such as *John sees Mary* can be expanded upon by recursively embedding the phrase in another, such as *Michael knows that John sees Mary*. The result is a communication system that can, theoretically, produce an infinite number of meaningful expressions through continued recursion and combination.

It is of interest to linguists and biologists that recursion arguably has not been identified in any other form of communication, such as that of animals (Hauser et al., 2002). How did recursion, and the resulting compositional structure of language, emerge in humans? What is it that drives its emergence in every successive generation of linguistically able humans? Theoretical accounts vary, some proposing that there is a biological factor unique to the human brain that contributed to the evolution of compositionality, and others rejecting that it is innate, or even that it even is truly universal at all (Evans and Levinson, 2009, pp. 440-444).

Empirical research into the potential origins and drivers of compositionality has taken many forms. Two in particular, artificial language learning and computational modelling, are the focus of the present study. Both present particularly interesting opportunities to investigate complex cultural phenomena such as language in empirical settings. Artificial Language Learning consists of experimental subjects learning a fake, toy language, and often then being given a certain task such as a communication or pattern recognition game. It has been used over the course of the past century in the fields of psychology, but in recent decades it has surged in popularity amongst linguists as a way of empirically testing models of language acquisition, use and transmission in a real setting.

This paper discusses the findings from several studies employing artificial languages computational models. Specifically, the implications of these findings are examined with respect to the debate on compositionality's universality and innateness. Where it is argued to not be innate, the proposed drivers of its emergence are examined and analysed. In Section 2, the perspectives of functionalism and nativism are introduced and their positions on this debate are discussed. In Section 3, arguments and evidence that compositionality may be a consequence of language being used within a community are discussed. In Section 4, an alternative view that compositionality may be an innate instinct of the language user is proposed and discussed. In Section 5, the proposal that language and compositionality emerge from a variety of domain-general systems is laid out.

2 COMPOSITIONALITY: INNATE OR LEARNED?

2.1 Nativism

On one side of the debate, nativists claim that language primarily functions as an instrument for thought (Keizer, 2015). Chomsky (2007) posits that language initially evolved as a way of evoking symbolic, cognitive images, and only later was externalised to be used as a form of communication. While the experiences of the language learner/user are crucial to the development of their specific language, communicative interactions are not integral to the existence of the instrument of language itself. The mental linguistic faculty is an innately endowed section of the language learner's mind. Communication, therefore, is a secondary aim to the nativist camp.

The early format of Chomsky's nativist solution, Universal Grammar (UG), often took the form of the Principles-and-Parameters format (PP), whereby certain 'principles' – universal properties or qualities of languages – exist cross-linguistically and each language is a combination of parameters that qualify or calibrate each principle (Chomsky, 1980, 1993, 1998).

2.2 Minimalism

More recently, the minimalist program has presented another solution to understanding the nature of the language faculty. Instead of a "top-down" approach of identifying all principles and parameters found in the world's languages, minimalism follows a "bottom-up" approach by identifying "how little can be attributed to UG while still accounting for the variety of internal languages attained" as a result of environmental, "third-factor" principles (Chomsky, 2007, p. 19). The result is the identification of one essential operation: Merge. This operation is the process by which the language learner combines two elements, X and Y, into one superset $\{X,Y\}$. The meaning of $\{X,Y\}$, which may be a sentence, phrase, or individual word, is therefore a compositional result of the constituent parts. Additionally, constituent parts can be re-combined repeated through the aforementioned process of recursion, which some nativists argue is the foundation that separates human communication from animal (Hauser et al., 2002) as it allows for unlimited expression from a finite set of building blocks. From this, minimalists hold, language learners are able to form words out of morphemes and phrases out of words.

2.3 Functionalism

On another side of the debate lie many functionalists who would disagree with nativists on several of the above points. Generally, functionalists view language primarily as a tool for expression (Keizer, 2015), rather than for thought. They hold that language only emerges in individuals through interaction, and therefore communication is the ultimate driver of organisation in a language. Functionalists may concede that compositional structure could be a linguistic universal, but they would attribute its emergence to domain-general cognitive abilities, rather than any specific linguistic ones. Moreover, many point to evidence that suggests that linguistic universals are impossible to find in the diversity of natural languages in the world.

Taking the example of hierarchical, compositional structure, Evans and Levinson (2009) refer to the large number of natural languages with a truly free word order as a counterexample to the perceived universality of constituent structure. They claim that free word order poses a problem to the generativist approach of grouping adjacent constituents into higher-order constituents, as the position of each lower-order constituent can vary without a change in its syntactic function. It is important to note here that the minimalist framework of generative grammar, however, allows for the movement of constituents, resulting in a different spell-outs without any change on the underlying level of meaning. Keeping consistuent parts together is only a surface-layer, linearised result of deeper hierarchical concepts, which can result in a number of different possible linearised expressions. Free word order, therefore, is not a convincing counterexample to the claim that compositionality is universal.

In addition, Evans and Levinson refer to Australian languages such as Jiwarli (Dench et al., 1988; Nordlinger, 1998) that do not seem to present constituent structure at all as a way of organising the language. They present examples of case suffixes matching up elements that in English would belong in two distinct clauses, such as in the sentence "the child (ERG) is chasing the dog (ACC) of the woman (DAT-ACC) who is sitting down cooking meat (DAT)". Instead, dependency relations are used to connect the "dog" and "woman" arguments across clauses, which Evans and Levinson claim is not possible with a constituent structure approach. However, minimalists are also able to parse this phrase by combining the morphemes (in this case, suffixes) that connect the arguments with Internal Merge, just as dependency relations do. Evans and Levinson's preference for dependency relations therefore is ultimately similar in its purpose of connecting the meanings of various constituent parts to form novel compositional meaning, serving the same purpose as generativist's conception of compositional structure.

Moreover, the concept of recursion, by which infinite outputs are created from a finite pool of inputs, has also had its universality called into question by Evans and Levinson. They refer to polysynthetic languages that limit the degree to which arguments and constituents can be embedded in one another due to their already complex morphology. They further refer to languages that cap the number of levels of recursion permitted (Kayardild, Evans (1995a,b)) or that disallow it entirely (Pirahã, Everett (2005)). Here, further research is required to understand the implications for the mental language faculty: are the levels of recursion limited due to pragmatic reasons (i.e. is it uncommon rather than impossible to apply recursion multiple times) in these languages? If so, it would not preclude the compositional structure of the language in the mind, but rather its acceptance amongst these language users.

When discussing languages as counterexamples to the universality of compositionality, one language is often cited: Al-Sayyid Bedouin Sign Language (ABSL), a very young language spoken by a community in southern Israel with a high incidence of profound prelingual deafness (Sandler et al., 2005). The community of ABSL speakers is highly endogamous, meaning the language has emerged largely independently from the primary spoken languages of the region, Hebrew and Levantine Arabic. In her extensive field research, Sandler describes ABSL as fully functional, systematic and comprehensive, but interestingly claims that it does not display duality of patterning. All human languages, it is believed, exhibit duality of patterning (the combination of meaningless acoustic units into meaningful linguistic units) (Hockett, 1960). Sandler, however, argues that in ABSL concepts are articulated individually, rather than comprising of smaller, meaningless constituent parts (Sandler et al., 2011). Does the example of ABSL pose a fundamental challenge to the nativist perspective, or is there something important context that may contribute to a lower incidence of clearly compositional signs? Possible answers to these questions are discussed later in this paper, in light of findings from artificial language learning experiments.

It is problematic to functionalists to explain the diversity of structures and phenomena in the world's languages as emerging from one single cognitive system as the nativists claim. Instead, functionalists describe the language learner's mental system as a hybrid "biological system tuned to a specific linguistic system, itself a cultural historical product." (Evans and Levinson, 2009, p. 446) The various 'universals' that are observed cross-linguistically have in fact evolved independently of one another, they claim, but the high frequency of certain phenomena (e.g. the role of "subject", or the distinction of nominal and verbal roots) represent "evolutionarily stable", and functionally more economical features of a communication system.

Nativists often respond to these claims by drawing a distinction between what languages are instantiated in the world and what the internal, mental language system operates on. Chomsky distinguished two distinct forms of language: "internal" or I-language and "external" or E-language (Chomsky, 1986). The former refers to the neural/cognitive system that processes the grammars learned by the individual language learner over their lifetime, while the E-language is an abstraction of the linguistic behaviours (i.e., the language uses) of a certain community in a certain place and time. The I-language is the form of language that is of most interest to the study of a theory of the emergence of language grounded in biology. By contrast, the E-language is more interest to other levels of linguistic analysis such as historical linguistics, for example. Examples of E-languages, therefore, with limited or no compositional structure do not pose a challenge to the notion that all I-languages are processed in a hierarchical, structured manner, as nativists claim.

3 COMPOSITIONALITY AS A RESULT OF LANGUAGE USE AND GENERATIONAL TRANSMISSION

3.1 Computational studies

Recent empirical work has drawn links between the interaction of the I-language with the E-language and the emergence of compositional structure. Across many studies, Kirby (2001; 2008; 2014; 2015; Kirby and Hurford (2002); Ren et al. (2020)) has argued that compositional structure only emerges as a result of the iterated process of language acquisition across generations (see Figure 1). While domain-general cognitive processes contribute to the development of structure over generations, Kirby does not require a biological or innate language-specific mechanism for the development of compositional structure in human language.

This has been demonstrated through computational techniques (Kirby, 2001; Kirby and Hurford, 2002) with the design of the Iterated Learning Model (ILM). The ILM was developed as a way of simulating the transmission of language across generations free from any biological or evolutionarily selective pressures and without any objective measure of communicative success driving the model's adaptation. The model was built with four key components:

1. a meaning space;
2. a signal space;
3. one or more learning agents;
4. one or more adult agents.

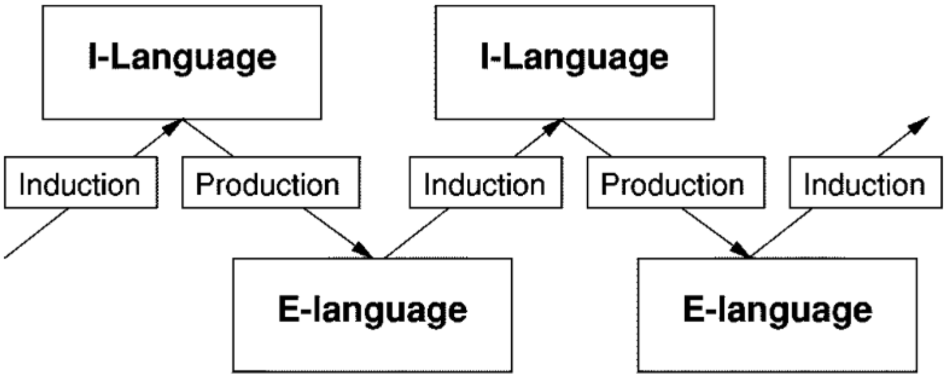

Figure 1. Process of linguistic transmission. From Kirby (2001).

The agents in the ILM learn the artificial language through an induction algorithm which consists of incorporation (learning new rules in the language) and generalisation (applying existing rules to new input). Additionally, the agents were defined with an invention algorithm which aimed to create new signals for meanings that did not already exist. The agents then communicated, passing on the language to the next 'generation' who then continued the cycle of learning and communication. It is important to note that the agents were not punished or incentivised depending on their communicative success: they do not receive any feedback on whether their communication has been successful or not.

The results of this modelling experiment showed that compositional structure emerged over the 'generations' the model cycled through. Kirby noted that the initial results were too uniformly regular, as there were no pressures on the agents other than for the language to be learnable (Kirby, 2001, p. 106). A second round of the experiment added in a production pressure, whereby the agents were coded to prefer to express themselves using the shortest string possible in a situation where there is a choice.

The 'language' the model developed over successive generations exhibited clear examples of compositional structure, with individual characters or series or strings corresponding to a certain aspect of the stimulus being described. The language evolved and varied over the generations, leading the researchers to hypothesise that it was this process of generational transmission that drove the emergence of compositionality (see again Figure 1). Kirby argues that, since the only pressures the agents felt were those of learnability and those of production, the emergence of compositional structure can only be attributed to their influence. He describes the language used by the model's agents as "adapting to improve its own survival" (Kirby, 2001, p. 108), and the process of linguistic transmission imposes a form of natural selection, described as a "bottleneck", on the features of the language. Compositionality, for Kirby, represents an efficient and effective manner of learning to express oneself.

In more recent computational work, Ren et al. (2020) build on Kirby's iterated learning model by incorporating neural networks. The Neural Iterated Learning model (NIL) was similarly trained on an artificial language, with agents that were periodically reset and then re-trained on the language, to mimic the inter-generational transmission in nature. Similar to Kirby (2001), each successive generation of agents was only trained of a subset of the output of the previous generation, thereby recreating the "bottleneck" effect of the process of induction seen in the ILM. Ren et al. found that compositionality again emerged over successive generations, as a result of the iterative learning process. The process of acquiring the language each generation was made faster and more efficient by the emergence of compositionality: the researchers argue that the agents were able to use the compositional nature of the language to generalise their linguistic rules to new inputs, thereby speeding up the learning process.

The claim that compositionality aids generalisability has been contested by other computational studies. Chaabouni et al. (2020) developed a neural network similar to Ren et al. (2020), with a 'speaker' agent and a 'hearer' agent. The speaker was tasked with accurately communicating the label for a target object and the hearer with choosing the correct meaning based on the communicated signal. Importantly, the agents were incentivised to communicate successfully with rewards being given when the hearer chose the correct label for the speaker's intended object. The agents were then tested on unseen stimuli that contained features that had previously been seen (e.g. a green triangle, when the agents had only been trained on green circles and red triangles). This test was intended to determine whether the 'language' the model developed could be used to generalise the existing labels to new stimuli containing the same features.

The researchers found that generalisation did emerged as a result of prolonged exposure to an increasingly varied input set. However, they demonstrated that compositionality was in fact *not* necessary for this process of generalisation: it was possible for languages to emerge that were capable of generalisation and that did not demonstrate compositional features. Instead, the researchers hypothesised that, when compositionality *did* emerge in the neural agents' interaction, it meant the language was much easier to learn by successive 'generations'. This aligns with functionalist thought that compositionality may emerge "by a random mutation but, once present, it will survive and thrive, as it guarantees that languages possessing it will generalize and will be easier to learn" (Chaabouni et al., 2020, p. 4434). In short, while generalisation did not require compositionality, compositional languages were both capable of generalising and easier to learn quickly and thoroughly.

3.2 Artificial Language Learning studies

In later work, Kirby modelled the effect of these communicative pressures on the development of artificial toy languages in human participants. The pressure of learnability, referred to as *compressibility*, drives a language to have as simple a form as possible. Taken to its extreme, a language that is fully compressible is referred to as *degenerate* (Kirby et al., 2015), where all meanings map onto the same form. In such a case, the language is the easiest to acquire as possible, but especially poor for communication and expressing oneself. The second pressure, therefore, is for *expressivity*, the need to be able to refer to each and every possible concept. A language that is ultimately expressive can be thought of has having a one-to-one form-to-meaning mapping, referred to as a *holistic* language, and represents a significant challenge for learning.

In Kirby et al. (2015), participants were presented with a variety of visual stimuli, each containing everal features (shading, size, shape, orientation), building on the artificial language learning paradigm from Kirby et al. (2008). These participants were 'trained' on an artificial language by learning a series of inputs labels, each describing one of the visual stimuli. This initial artificial language therefore was holistic, as there was a one-to-one form-to-meaning mapping. The participants then were tasked with a 'playing phase', wherein the 'speaker' was presented with a stimulus and needed to communicate its name (the 'signal') to the 'hearer'.

In this particular study, the researchers evaluated the impact that generational transmission had on the form of the language. The participants were split into two groups: one group consisted of a "diffusion chain" structure, whereby the resulting set of labels was then taught to new, naive participants who repeated the experiment, while the other group only communicated within their pairs (see Figure 2). In the diffusion chain group, the language learned by the second partner was then used as the learning input for the next round, and the form of the language was followed throughout the iterations, or ‘generations’. Unlike the computational ILM (2001), participants were given immediate feedback after every round: the hearer is shown the correct label, and the speaker sees which label the hearer selected. If the selection is correct, a point is added to a running total, thereby creating a pressure for effective communication between the pairs.

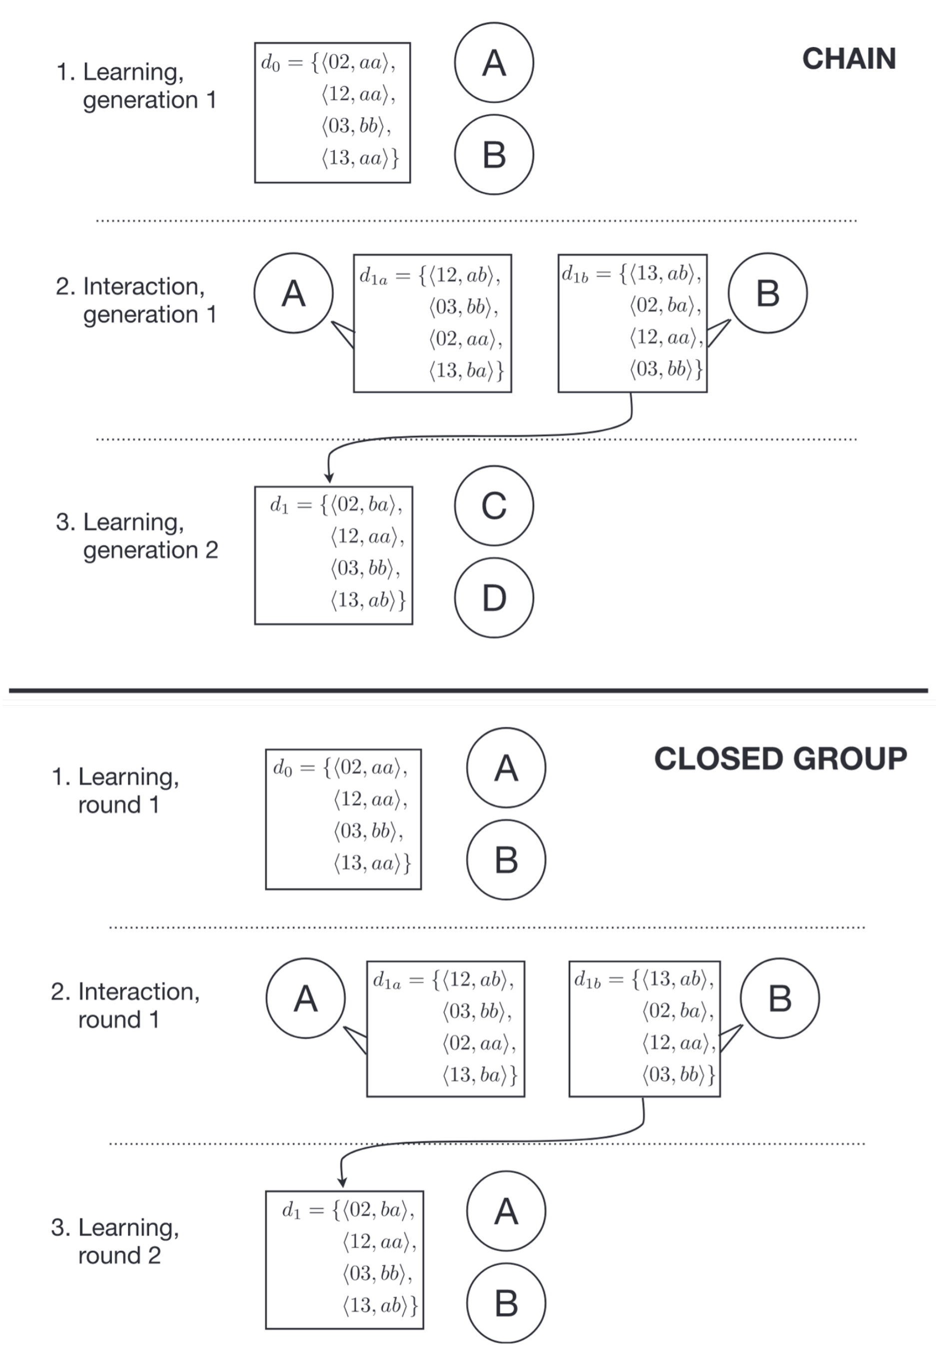

Figure 2. Simulation procedure in the chain (top) and closed group (bottom) conditions. During learning (1), agents A and B are exposed to some data and sample a hypothesis according to *P(d|h)*. During interaction (2), the agents take turns to produce meaning; forming pairs according to *P(d|h)*, updating *P(d|h)* according to their partner’s productions. The data produced by one randomly-selected agent during interaction is used (3) to train the next generation of agents. In chains (top) these are fresh, naïve learners; in closed groups (bottom) they are the same two agents. From Kirby et al. (2015).

The researchers found that only when a language was passed on to new participants in a chain-like format did the final language display significant levels of compositional structure, with morpheme-like elements representing features of the visual stimulus. The alternative group who only spoke with the same partner each round did not result in a compositional language, and instead continued to use a mostly holistic language. The researchers modelled only the pressures of expressivity and compressibility as acting on the participants. The researchers argued that in the closed group, there was a weaker pressure for compressibility, leading to participants being able to use the more expressive, holistic language throughout each iteration. In the chain format, participants needed to construct a language that was in some way compressible so that it could be learned by new participants, resulting in competing pressures of expressivity and compressibility. The researchers therefore argued that these findings suggest that compositional structure only emerges when there is a balance between the pressures for communication and for learning. Their resulting hypothesis is that, under certain conditions, communities that have less of either of these needs may tend more towards either holistic or degenerate languages.

However, throughout the experiment, a third pressure for *communicative success* was imposed, thereby differentiating the experiment from early computational work such as the ILM (Kirby, 2001). In addition to the explicit incentive to communicate the correct label, participants were told from the beginning that they were to learn a language, rather than, for example, to generate an image as a group. Since all participants were adults who spoke at least one language, it is likely that mentioning they will be building a new language already implied some sort of communicative aspect. It is difficult, therefore, to divorce the effect observed of compositionality emerging from a balance of expressivity and compressibility from the potential impact that a pressure for communicative success had. Moreover, those who subscribe to the view that language is primarily an instrument for thought may disagree with the inclusion of a pressure for communicative success. Language, it can be argued, is not always used purely for the purpose of successful communication (consider metaphorical or ironic language that often is purposefully chosen for its non-transparent meaning (Cormac, 1990)). It could be argued therefore that the ‘language’ that emerged in this experiment may in fact be more of a ‘code’ that enabled participants to succeed at the task at hand but cannot act as a proxy for the mental processing of language. Including a pressure for communicative success can be argued to result in a circular experiment design, where the conclusion that only communicative pressures drive the emergence of linguistic features is built-into the experiment design from the outset.

Moreover, it can be argued that the pressure for expressivity was engineered in the experiment by the inclusion of the communicative success metrics and not inherent in a language learner’s use of language. Languages with multiple meanings for a certain form are just as expressive as fully ‘holistic’ languages. Consider the English words *sheep* or *fish* that can refer to either the singular or plural concept. It is unlikely that one would consider English to be less expressive than, for example, Modern Greek, that uses different suffixes to express distinguish the plural from the singular meanings. It is therefore important to note that, in this experiment, the participants began initially with a holistic language, thereby taking the language to the extreme end of expressivity, a fact which Kirby acknowledges as a limitation to the study. It is possible that if the experiment were to begin with a degenerate language instead and if the aforementioned pressure for communicative success were removed, the participants might have no problem learning and passing on a degenerate or minimally ‘expressive’ language. It would be interesting to see if under such conditions, compositional structure could be observed to emerge in the closed group as well as the diffusion chain, as this would imply that Kirby et al.’s concept of ‘expressivity’ may not be as significant a pressure on language learners as that of learnability, as proposed by Chaabouni et al. (2020), among others.

Finally, it can be argued that Kirby invokes that fallacy that compositionality must be visible within the word. There are many isolating languages in the world whose words tend to not be decomposable into constituent morphemes. These languages, nevertheless, still exhibit compositionality in the manner in which words are combined to create meaning in phrases, clauses and sentences. It is plausible that, in another hypothetical task, the participants in Kirby's studies could combine the labels of the 'holistic' starting language in a way that expressed some sort of meaning. It would be incorrect, therefore, to assume that a one-to-one form-meaning mapping implies a lack of compositionality.

3.3 Linguistic universals as tendencies for simplicity

Recently, further empirical evidence has emerged that language learners apply structure and regularisation patterns to a variety of inputs (linguistic and otherwise) as a result of pattern-learning biases. Culbertson et al. (2024) conducted several pattern-learning experiments across both linguistic and non-linguistic domains. The aim of these studies was to identify which patterns are easiest to learn in order to determine whether certain learning biases may play a role in the emergence of linguistic 'universals'.

The various experiments were similar in their design. The one focussing on the linguistic domain dealt with an artificial language consisting of 'heads' and 'dependents'. Heads were longer (two syllables) while dependents were shorter (one syllable). Participants were placed into one of three groups: harmonic, non-harmonic across-heads, and non-harmonic within heads. The first had a consistent ordering of dependents and heads, be it head-first or head-last, imitating languages such as English or Turkish that align the order of the head and dependent across phrases of different types. The non-harmonic across-heads group differed in the head-dependent order across different categories of heads: some heads came before the dependent, some after. This imitated languages such as Farsi or Tiringa that have, for example, verbs before the object but postpositions after the dependent. Finally, the third group had all potential orders, with dependents coming both before and after heads of each category, as is possible in Mandarin, for example. The subjects were tasked with learning to recognise two new sequences consisting of shorter sequences (dependents) and longer sequences (heads) and to understand how the shorter ones attach to the longer ones. For the visual, auditory and tactile experiments, corresponding sequences of stimuli were presented in their respective modalities, with smaller or shorter stimuli representing dependents and larger or longer stimuli representing heads.

The participants were tested with novel stimuli and asked to determine whether these fit into the patterns they had learned. For the linguistic experiment, this equates to determining whether the test stimuli were grammatical. The results demonstrated that the artificial languages with harmonic word order (a consistent order of heads and dependents across categories) were easiest to learn. Participants in this group correctly judged the test stimuli to be grammatical far more than the non-harmonic groups. The non-harmonic across-heads group answered correctly the second most, with the non-harmonic within-heads group (the most varied input) answered correctly the least. This trend was also observed in the other modalities, although it was most pronounced in the linguistic and visual stimuli compared to the aural and tactile modalities.

The findings from these experiments and the resulting cross-domain analysis were taken to support the view that linguistic patterns emerge due to a tendency for simplicity, rather than an innate part of the linguistic processing system. The fact that similar patterns were detected in the visual, aural and tactile experiments led the researchers to conclude that this tendency for a harmonic pattern "lies outside the language faculty in a domain-general bias for simplicity. This bias would operate on the linearization of language-specific categories and representations" (Culbertson et al., 2024, p. 14). According to this view, therefore, it is possible that domain-general biases may also apply to the language-specific representation of compositional structure. Compositionality may therefore emerge as a result of domain-general biases for simplicity. While not explicitly studying the emergence of compositionality, the arguments presented in this study appear to support the view that linguistic 'universals' such as compositionality may not be innate to the human mind, but rather are a result of a preference for simple representations and therefore happen to be favoured by cognitive biases.

3.4 Neural networks without cognitive biases

Interestingly, further support for this view can be found in other computational experiments with neural agents. Lian et al. (2023, 2025) proposed the Neural-agent Language Learning and Communication framework (NeLLCom) which consists of neural agents acting as 'speakers' and 'listeners', similar to Kirby's (2001) and Ren et al.'s (2020) iterated models. The agents were trained on a miniature, artificial language, and then were tasked with playing a communication game. This was a means of 'optimising' the language for communicative efficacy by punishing the mismatch between the information conveyed in the meaning of the speaker's 'utterance' and the listener's prediction of its meaning.

Rather than compositionality, in Lian et al. (2023) the object of study was the trade-off between case marking and freedom of word order, a pattern that is generally observed in languages worldwide (Comrie, 1989; Greenberg, 1963). This trade-off refers to the tendency of languages with more flexible constituent order to have more complex morphological case systems. Previous work on this phenomenon has hypothesised that this trade-off emerges as a strategy to balance the effort of production with the reduction of ambiguity and information loss Fedzechkina et al. (2017). A freer constituent order reduces production effort, but risks raising ambiguity; this is therefore mitigated by a more complex morphological system that reduces ambiguity but may increase production effort.

Despite not training NeLLCom with any explicit 'cognitive' biases other than the desire to be understood by the conversation partner, Lian et al. found that the agents' language eventually developed the case marking – word order trade-off observed in natural languages. Crucially, in the control group, wherein the agents were not tasked with communicating, simply learning the artificial language, the trade-off did not emerge. The researchers propose that this particular linguistic 'universal' therefore emerges purely as a result of a pressure communicative success, as that was the only pressure present in the experiment.

A similar result was observed in the later study (2025) which focussed on differential case marking (DCM). DCM refers to phenomenon in which arguments may be marked differently based on certain semantic or pragmatic properties (e.g. animacy). Similar to the case marking-word order trade-off, it has been proposed that DCM represents a strategy to minimise ambiguity but requires more effort in production (Fedzechkina et al. (2012); Smith and Culbertson (2020)). NeLLCom was again taught a language that had the option of employing DCM and was tasked with communicating a variety of events and settings. The researchers found that DCM emerged only when the agents were tasked with communicating, not in the learning process alone. Similar to the earlier study, this was taken as further support that DCM is employed specifically in cases where communicative efficiency would be increased.

While not explicitly studying the emergence of compositionality, these studies can be taken as support for the view that linguistic universals may only emerge because they represent efficient strategies in communication. As such, it is not difficult to imagine that compositionality, with the advantage it lends to the acquisition of language in newer generations and in expressive capabilities, may simply represent a tendency emerging from cultural evolution, rather than an innate linguistic necessity.

3.5 Conclusion

It is important to keep in mind that, to nativists, compositionality is a mental tool by which meaning and form are mapped onto each other. They would argue that even when words or signals do not display overtly combinatorial properties, language users still process the possible meaning of a signal through compositional means. With the example stimuli of Kirby et al.'s 2015 study, even though the initial word set was holistic, nativists could argue that the participants would mentally decompose the image into its constituent features (shape, shading, orientation, etc) in order to associate the features with the correct label. Whether or not this decomposition appears in the label that is chosen in order to achieve a certain communicative goal could be seen as arbitrary to nativists who would be more concerned with the mental processing of such concepts.

4 COMPOSITIONALITY AS INNATE

4.1 Compositionality emerging within a generation

If compositionality is an innate mental process, what form of empirical evidence could be found in artificial language learning experiments to support this viewpoint? One potential avenue of investigation is whether individual language learners produce compositional languages within their own time-frame of participation in the artificial language learning experiments. Where Kirby and others have shown linguistic universals such as compositionality emerging over successive generations as part of the process of transmitting language, other studies show that it can emerge within a single generation.

Expanding on Kirby's experimental work, Raviv (2019a; 2019b) has demonstrated that, while the pressures of compressibility and expressivity may indeed be what drives the emergence of compositionality, they are not a direct result of the 'diffusion chain' generational format of learning. Raviv (2019a) was able to procure the emergence of compositionality in a similar experimental design to Kirby et al.'s (2015), but was able to do so even in the closed group. Like in Kirby's experiment, participants communicated in pairs, with a speaker and a hearer, but crucially these pairs were alternated with others, resulting in a fully connected group of four participants communicating in every round. Additionally, a growing semantic space was used, with new inputs being presented to participants throughout the process of the experiment. This and the subsequent study (Raviv et al. (2019b)) produced structurally compositional languages by the final round, without the need for new, naïve participants to pass on the language to.

Raviv's work raises questions as to whether the compressibility pressure only arises through generational transmission, as proposed by Kirby. Theoretically, within a single 'generation' there is less of a pressure for compressibility, as all conversation partners have already learned the language. Instead, it appears that the larger variety of inputs resulting from multiple interaction partners can also exert this same pressure for compressibility within the same generation. This implies that there may be a type of instinct within the language learner's mind to impose a structure on linguistic inputs as they grow in quantity. Raviv hypothesises that this may be caused by domain-general working memory limitations, whereby the language learner aims to optimise their communicative performance with limited memory capacity.

This view could be seen as support for the modern nativist's conception of the language learner's mind: the linguistic universal of compositional structure can develop within an individual as a response to complex stimuli, rather than purely as a result of interaction. Counter-examples, such as ABSL or Jiwarli, would therefore represent exceptional cases where the pressure from multiple conversation partners or an expanding meaning space is not sufficiently strong to require holistic words or signs to be decomposable.

For example, communication in ABSL has been described as being relatively highly redundant and contextually dependent (Sandler et al., 2011) meaning that the tight-knit nature of the community results in signers sharing a significant proportion of daily experiences. This produces a relatively concentrated semantic space, with few new concepts or inputs being introduced, thereby enabling signers to retain holistic or iconic signs throughout the generations. Additionally, the close-knight community structure allows for relatively difficult-to-learn conventions and idiosyncrasies to be passed down generationally. With Raviv's hypotheses in mind, it is proposed that with a larger pool of concepts to refer to, or a growing signing community, ABSL would eventually exhibit more instances of compositionality. Again, it is important to note that, even if ABSL does indeed exhibit less or no duality of patterning, nativists maintain that the meaning of the concepts communicated by the signs in the language is still processed compositionally in the mind. This view challenges Evans and Levinson who reiterate the relativist claim that compositionality does not exist in the *minds* of speakers of languages such as Jiwarli (Evans and Levinson, 2009, p. 441), since it is (arguably) not expressed in the language.

4.2 Compositionality in other domains

Additional support for the view of compositionality as an innate instinct can be seen when one removes the pressure from semantics (and therefore expressivity) entirely. Verhoef (2012; 2014) conducted artificial 'language' learning experiments similar to Kirby and Raviv's paradigm, but using a slide-whistle toy instead of words. Participants were asked to listen to and replicate a series of twelve signals produced by a toy sliding whistle. In the testing phase, the subjects were required to produce twelve unique signals, thereby requiring multiple forms and preventing underspecification that would normally come about in natural language through a need to describe different concepts. Similar to Kirby's (2008, 2014, 2015) experiment design, the reproduced output was then used as the learning input for the next set of participants, creating a diffusion chain.

The aims of Verhoef's studies were to understand what processes the mind applies to acoustic signals through iterated exposure in an effort to understand what might be driving the formation of combinatorial structure in spoken language. Verhoef's aim was to see whether duality of patterning emerges as a result of cognitive processes being applied to a non-linguistic input or if it is specific to language. She found that, by the end of the experiment, participants tended to produce a form of combinatorial 'language' whereby the overall signal was made up of smaller elements (staccato, elongated sounds, down-up slides) that were repeated in different combinations. Quantitatively, Verhoef identified that an increase in combinatorial structure in the sounds correlated with a decrease in recall error in the signals produced by participants. This implied that compositionality was being employed as a strategy to reduce information loss in the learning process by reducing a potential infinite input to a finite set of underlying components, thereby making the acoustic signal more easily learnable.

A crucial question arises in light of Verhoef's work: does the non-linguistic nature of the signal studied in her experiments (and the corresponding lack of reference to a 'language' in the participants' instructions) imply that the potentially innate instinct to organise information combinatorically applies not only to language, but to other cognitive subsystems as well? Such a suggestion would be acceptable to some nativists, but also to emergentists and bio-linguists who subscribe to the view that there are many components to the faculty of language more broadly, including possibly some shared with other subsystems of the brain.

5 COMPOSITIONAL STRUCTURE AS EMERGING FROM DOMAIN-GENERAL PROCESSES

5.1 Emergentism and biolinguistics

Many of the views described so far propose that, if biology plays a significant role in the emergence of linguistic universals such as compositionality in language, it is likely not the only factor at play and it is possible that it is not a language-specific faculty of the brain that is driving it. Kirby's, Raviv's and Verhoef's work acknowledge cognitive biases present in the process of learning language (most often from a limited input), but also acknowledge the need for interaction between language learners for it to emerge – it will not do so in a vacuum.

The idea that compositionality emerges from certain domain-general cognitive processes, and possibly also through interaction among language learners, has been described as *emergentism*. Here, scholars model language as a complex, dynamic system that emerges from the interaction of a range of domain-general subsystems (Smith and Thelen, 1993). This view appreciates that an individual's capacity for language acquisition is shaped by factors such as joint attention, analogy and conceptualisation which are not specific to language, but also acknowledges the resulting existence of linguistic categories such as verbs, verb phrases and tense. These categories are perceived as emerging from the interaction of various nodes on a cognitive 'network' that specifies associations "between different aspects of a speaker's linguistic knowledge" (Diessel, 2020). As such, linguistic concepts are a result of domain-general cognitive processes, and are seen as especially transient, rather than representing a fixed, in-built state. Linguistic universals, such as compositional structure and recursion, therefore, represent a culmination of an independent, complex, dynamic system responding to environmental pressures for communicative efficiency (Verhoef, 2012).

As evolutionary biology has developed over recent years, linguistic frameworks have incorporated biological findings into their viewpoints. Nativists refined the faculty of language into a 'broad sense' (FLB) and a 'narrow sense' (FLN) (Fitch, 2010; Hauser et al., 2002). This view of 'biolinguistics' holds that the language faculty as it has traditionally been termed is in fact only a narrow sense of the broad range of cognitive components that contribute to language use. Subsystems such as working memory, combinatorial phonology and theory of mind have been shown to not be unique to language, nor to human communication. Hauser et al. (2002) proposed that the only component unique to humans and language is that of recursion, which may have evolved from a domain-specific system related to navigation, for example, to one that can be used for other domains, such as language (Hauser et al., 2002, p. 1578).

Emergentism does not preclude the significant role that innate biological factors have on the development of language. As mentioned, nativists do acknowledge that there are multiple components to the processing and production of language, many of which are not uniquely associated with this function. Beyond that, it is also acknowledged that an individual learner's experience and interaction plays a role in the development of the language faculty over time, but the extent to which it does may be debated. Both Hauser et al. (2002) and Kirby (2001) accept the role of biology in the development of language. While Kirby models the compositionality emerging purely through transmission, it is domain-general cognitive constraints that human biology imposes on the processes of induction and production that cause the emergence of compositionality. These domain-general cognitive constraints can also be found in Hauser et al.'s concept of the FLB.

5.2 Entropy model of rule induction

Researchers have also investigated and modelled these domain-general systems that have been hypothesised to drive the emergence of language using artificial language learning paradigms. Radulescu et al. (2019) conducted pattern-recognition studies where participants were presented with sequences of three syllables from a 'forgotten language'. Participants were exposed to the sequences and were told that this forgotten language "had its own rules for the forms of words, and that those words were not known to them from any other language they might be familiar with." After the familiarisation phase, the participants were presented with new inputs and were asked to determine whether these inputs would be possible in this 'language' they learned. Crucially, the quantity of information contained in the familiarisation set of inputs was measured and adjusted from 'low entropy' to 'high entropy'. In each familiarisation phase, participants were exposed to the same number of stimuli, but with different quantities of unique combinations. In the low entropy group, six variants of the first syllable and six of the second were presented each time; in the medium entropy group, 12 variants of the first and 12 of the second were presented twice each, and in the high entropy group 24 variants of each syllable were each presented once. This enabled the researchers to measure the quantity of information conveyed in each stimulus to see if it correlated with the rate of success.

The significant finding of this experiment was that participants' success rate of identifying possible syllable combinations in the language correlated to the degree of 'entropy' (essentially, the complexity of the input) contained in the familiarisation set. Participants were significantly more successful at identifying rule patterns in the high entropy set compared to the lower entropy sets. A second experiment expanded the difference in entropy values between the low entropy and high entropy groups in order to confirm the relationship between entropy levels and rule induction.

The results of this study and subsequent studies (Radulescu et al., 2021) are taken as support of the entropy model of for rule induction. Under this framework, rule induction is a strategy employed by the brain to maximise its encoding power of complex inputs (entropy) with a finite processing capacity (channel capacity). With a much larger channel capacity, language learners can rely on rote memorisation rather than rule induction. In the experiments' low entropy cases, input complexity did not exceed channel capacity, evident in the relatively high rate of correct rejections of familiar syllables, but low rate of correct rejections of new syllables. This meant that participants tended to memorise the familiarisation set rather than learning the rules to form new words. As entropy increased, the researchers found that the participants tended to use category-based generalisations, leading to a higher rate of correct acceptance of new syllables.

Radulescu et al.'s model is similar to Ren et al.'s (2020) hypothesis that generalisation allows for more efficient processing of a limited cognitive resource. In Radulescu et al's case, the generalisation is the result of rule induction, while in Ren et al it is compositionality that allows for generalisation. Compositionality therefore can be thought of as playing a similar role to rule induction in the learning process, as it allows the learner to break down a possibly infinite input set into a finite set of recognisable, learnable parts and to generalise these smaller parts to similar ones that the language learner is exposed to. Radulescu and other researchers in this field hypothesise that children therefore tend to acquire complex rules impressively fast due to their limited channel capacity requiring them to rely on rule induction more than memorisation. Under the entropy model of rule induction, therefore, compositionality would most likely emerge as a result of domain-general cognitive limitations (channel capacity) and would do so in settings where the complexity of the input is greater than the channel capacity of the language learner.

5.3 Re-examining artificial language learning experiments

Language acquisition, therefore, can be thought of as following the same pattern of gradual reliance of rule induction as input entropy increases, according to this model. If compositionality emerges as a rule that is learned and applied by learners to new input, then it stands to reason that higher rates of compositionality will be seen in input sets with higher rates of entropy or information. Raviv's (2019a; 2019b) studies appear to bear out this hypothesis: the expanding meaning space and larger number of conversation partners added more variants to the input pool for participants, thereby increasing the entropy of the input. At some point, the channel capacity of the language users was exceeded by the entropy of the linguistic input, and the participants resorted to decomposing the language into smaller, combinatorial parts. The resulting compositional structure of the language can be seen as a strategy by which to process the inputs with the finite channel capacity of the participants' brains.

Interestingly, Raviv's meta-analysis of the results from the initial study (2019a) identified a larger number of conversation partners as introducing a stronger pressure for compositionality than the expanding meaning space did. Both contributed to a higher rate of compositionality, but the pressure from the expanding meaning space seemed "to be less crucial when another strong pressure for compressibility (i.e., interacting with multiple partners) already [existed]." (Raviv et al., 2019b, pp. 162-163) For the entropy model of rule induction to hold, it would be important to investigate exactly how much more entropy an expanding meaning space adds, compared to that added by the increased population size. If both factors contribute the same quantity of increased information density, it poses a problem for the hypothesis that this is the key driver of rule induction, since one was shown to be much more powerful as a predictor than the other.

5.4 Non-Adjacent Dependency learning

Work has also been conducted into more general cognitive mechanisms hypothesised to be involved in the language acquisition process. Similar to the entropy of a linguistic input, Grama and Wijnen (2018) investigated the role that variability in the linguistic inputs plays in the process of acquiring and understanding non-adjacent dependencies (NADs). These refer to morpho-syntactic dependencies between morphemes that are not contiguous (e.g. "The princess <u>is</u> reluctantly kiss<u>ing</u> the frog"). Children and adults use frequent patterns or 'frames' in the linguistic input to infer lexical categories as part of the language acquisition process, such as *he* __ *it* leading to *he* VERB *it*. Adjacent dependencies have been show to inform an understanding of phrase structure, and non-adjacent ones lead to an understanding of more distant, morphosyntactic relationships (Gómez, 2002).

Grama and Wijnen studied the responses of 18-month-old infants to novel NADs that follow a similar *a-b* pattern across an intermediary element X. The aim was to determine whether they are able to generalise the syntactic relationship to new items and whether they can track dependencies between morpho-syntactic items whose dependency is not immediately salient. Across two experiments, the infants were presented with a range of inputs from an artificial language that followed the *aXb* format, with *a* and *b* representing the NAD such as 'is X _ing'. After a familiarisation phase where the infants were exposed to the 'correct' stimuli a_i-b_i pairs, their acceptance of test combinations in the *aXb* frame was studied, i.e. they were asked to deem whether a 'correct' pairing (a_1-X-b_1) or 'incorrect' pairing (a_1-X-b_2) was grammatical or not. This tested whether the infants were able to identify the correct a_i-b_i dependencies across an intervening variable *X*, or whether the intervening *X* 'blocked' the recognition of the corresponding part of the dependency. The second experiment tested whether the children accepted new intervening elements, essentially an a_iYb_i frame, to determine whether they are able to generalise the NADs to new variables. The researchers found that the infants displayed significant levels of learning of the correct dependency and were able to reject combinations that violated the pattern. They were also able to generalise the rule they had identified in order to recognise familiar *a-b* pairs across a variety of *X* intermediary elements.

This study, along with previous work by Gómez (2002), has highlighted the importance of variability in the input in identifying rules. In the second experiment, the infants were exposed to higher rates of variability in the input, which the researchers argue supported the generalisation of learned rules. If the input had been less varied, rote memorisation would have been an equally efficient strategy, but with a highly varied input, focus was directed to structural regularities rather than surface-level similarities. Variability in the input can be equated to Radulescu's concept of entropy, thereby aligning the findings of Gómez's and Grama and Wijnen's studies with the entropy model of rule induction.

The importance of dependency relations to the hierarchical representation of linguistic units means that these findings can readily be applied to the compositionality discussion. The results show that children appear to decompose linguistic elements as a strategy to understanding the relationship between constituent parts and to facilitate the language learning process by allowing for generalisation to new input. This strategy is not infallible, however, as some instances of novel intermediary items made the NADs more difficult to detect, which the researchers propose may be due to the novel item directing "attention and processing resources away from the *a_b* dependencies" (Grama and Wijnen, 2018, p. 19). A compositional representation of language, therefore, appears to have emerged in infants as young as 18 months (and even as young as 15 months: see Gómez and Maye (2005)). Whether this is caused by domain-general systems such as joint attention or distributional learning or whether it is an processing mechanism specific to the language faculty can still be debated. It is nevertheless convincing to conclude that children apply structural rules such as compositionality to an input set with a higher level of information density or variability in order to a) faster learn the input and b) detect structural similarities that can be re-applied in the future.

6 GENERAL DISCUSSION AND CONCLUSION

This paper examined findings from recent computational work and empirical artificial language learning experiments related to compositional structure and factors that may contribute to its emergence and use. At first glance, examining empirical results and computational models to glean insights into the nature of a potentially innate mental processing system may appear misguided. However, the results from these studies provide valuable evidence of the complex phenomenon of language by modelling the processes of acquisition, use and transmission, from which linguists can distil individual driving factors. Aside from implications for the theoretical debate regarding the universality of compositional structure and its innateness, the results from the studies in this paper broadly attribute the need for compositionality to two purposes. The first refers to the various advantages that compositional language lends to the process of acquiring language. The second frames compositional structure as a natural consequence of domain-general cognitive processes and biases. These two arguments certainly overlap in the work that has been discussed in this paper, and many of the researchers cited might agree with both.

In the process of acquisition, language learners must develop a fully functional mental grammar based on a limited input. In this respect, compositionality has been argued to grant a significant advantage. The 'bottlenecks' with each successive generation of learners modelled by Kirby, Kirby and Hurford and Ren et al. demonstrate that a limited input set that can be combined repeatedly to form new meaningful elements allows for an expressive, compositional language to emerge. Similarly, the work by Lian et al. lends credence to the notion that, even without any cognitive biases, the process of interaction towards a specific communicative goal over several iterations leads to the emergence of some of the linguistic universals observed cross-linguistically. The arguments made by these researchers is that therefore compositional structure may be a *consequence* of cultural evolution (iterated, generational learning), rather than existing of its own accord in the language learner's mental toolbox. Moreover, it is not solely attributed to a language-specific component of the mind, but rather may have emerged as a result of the limited nature of the input and a variety of domain-general cognitive biases.

Although not directly examining compositional structure, additional arguments were discussed concerning potential domain-general biases that may lead to its adoption across languages. The work by Culbertson et al. concludes that cognitive biases for simplicity impact linguistic tendencies such as word order harmony. If such cognitive biases that also apply to pattern recognition in the visual, aural and tactile modalities affect the process of language acquisition, is it possible that similar cognitive biases may also affect the tendency to structure language hierarchically? It is possible, but the crucial role compositionality plays in forming meaning out of linear sequences is difficult to view as merely a popular tendency among the world's languages. A recreation of the experiments, perhaps, with a specific focus on combinatorial properties of the stimuli may provide more evidence for this claim.

Specifically with regards to learning, the results from the work by Verhoef et al. demonstrating that compositionality reduces the recall error associated with learning a signal further supports the notion that it poses a significant advantage to the language learning process. Whether this implies that compositionality is built-in to the notion of language from the outset, or whether it has developed as a result of its advantage for learning, is unclear from these particular studies. However, the fact that the participants in these experiments still employed compositionality as a strategy despite the non-linguistic acoustic signal used seems convincing that there is some innate impetus to parse acoustic signals to form some sort of meaning. This would make the strategy of decomposition and recombination a potentially domain-general one, even if it is innate.

It is also possible, therefore, to acknowledge the effect of domain-general processes on the emergence of linguistic features while also holding the view that languages learners innately process languages combinatorially. The work by Radulescu et al., Grama and Wijnen and Gómez supports these views by suggesting that an awareness of combinatorial properties of language develops as a natural response to increasing complexity of information. Especially interestingly, these results appear to further explain results from artificial language learning experiments that showed compositionality emerging within a single generation (Raviv et al., 2019a) when the variability of the linguistic parameters significantly increased. It is convincing therefore to surmise that compositionality may have emerged gradually throughout the process of the evolution of language as the complexity of linguistic expression increased.

Whether or not one agrees with the notion that compositionality emerges as result of domain-general processes, the evidence that a compositional structure may in fact *not* grant an advantage to the process of generalisation to new input is particularly interesting. In the work by Chaabouni et al., the conclusion that language may have evolved with compositional structure not because it is *required* in order to generalise rules to new items, but rather because it lends a significant evolutionary advantage to the language by making it easier to learn is very convincing. It is plausible to imagine that languages evolved over time with compositionality built-in due to this significant advantage to learning. It is also plausible, therefore, to imagine situations at some point in time or in some place in the world in which compositionality was not required in order to learn sequences of inputs and generalise them to new ones. This may potentially explain the exceptional cases of young, emergent languages such as ABSL (Sandler et al., 2005), relatively isolated languages such as Pirahã (Everett, 2005) or the artificial languages in Kirby et al. (2015) that did not display combinatorial properties.

In closing, the strong evidence that hierarchical structure is likely a result of some sort of biologically defined cognitive processes (be they specific to language or not) makes it difficult to accept the strong form of functionalism that holds that compositionality in language only emerges in response to communicative pressures. It may at first appear difficult to link the empirical evidence presented in this paper to a mental process of decomposing language into constituent parts, as proposed by nativists. However, the significant evidence that compositionality lends advantages to the processes of learning underlying rules by induction (and potentially to applying them to new inputs) makes it much more convincing that language is indeed processed compositionally as a means of connecting limited forms to potentially infinite meanings.

REFERENCES

Chaabouni, R., Kharitonov, E., Bouchacourt, D., Dupoux, E., and Baroni, M. (2020). Compositionality and Generalization In Emergent Languages. In Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4427–4442, Online. Association for Computational Linguistics.

Chomsky, N. (1980). *Rules and representations*. Woodbridge lectures ; no. 11, 1978. Columbia University Press, New York, N.Y.

Chomsky, N. (1986). *Knowledge of Language: Its nature, origin, and use*. Praeger, Westport, CT.

Chomsky, N. (1993). *Lectures on Government and Binding: The Pisa Lectures*.

Chomsky, N. (1998). Comments: Galen Strawson, Mental Reality. *Philosophy and phenomenological research*, 58(2):437–441.

Chomsky, N. (2007). Of Minds and Language. *Biolinguistics*, 1:009–027.

Comrie, B. (1989). *Language Universals and Linguistic Typology: Syntax and Morphology*. University of Chicago Press.

Cormac, E. R. M. (1990). *A Cognitive Theory of Metaphor*. MIT Press.

Culbertson, J., Compostella, A., and Kirby, S. (2024). Language structure reflects biases in pattern learning across domains and modalities. *Quarterly Journal of Experimental Psychology*, page 17470218241282404. Publisher: SAGE Publications.

Dench, A., , and Evans, N. (1988). Multiple case-marking in Australian languages. *Australian Journal of Linguistics*, 8(1):1–47. Publisher: Routledge _eprint: https://doi.org/10.1080/07268608808599390.

Diessel, H. (2020). A Dynamic Network Approach to the Study of Syntax. *Frontiers in Psychology*, 11. Publisher: Frontiers.

Evans, N. and Levinson, S. C. (2009). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, 32(5):429–448.

Evans, N. D. (1995a). *A Grammar of Kayardild: With Historical-Comparative Notes on Tangkic*. De Gruyter Mouton.

Evans, N. D. (1995b). Multiple Case in Kayardild: Anti-iconic Suffix Ordering and the Diachronic Filter. In *Double Case : Agreement by Suffixaufnahme*, OUP E-Books, pages 396–430. Oxford University Press, New York.

Everett, D. L. (2005). Cultural constraints on grammar and cognition in Pirahã. Another look at the design features of human language. *Current Anthropology*, 46:621–46.

Fedzechkina, M., Jaeger, T. F., and Newport, E. L. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences*, 109(44):17897–17902.

Fedzechkina, M., Newport, E. L., and Jaeger, T. F. (2017). Balancing Effort and Information Transmission During Language Acquisition: Evidence From Word Order and Case Marking. *Cognitive Science*, 41(2):416–446. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cogs.12346.

Fitch, W. T. (2010). The Lay of the Land. In *The Evolution of Language*, Approaches to the Evolution of Language, pages 11–202. Cambridge University Press.

Grama, I. and Wijnen, F. (2018). Learning and generalizing non-adjacent dependencies in 18-month-olds: A mechanism for language acquisition? *PLOS ONE*, 13(10):e0204481. Publisher: Public Library of Science.

Greenberg, J. (1963). Some Universals of Grammar with Particular Reference to the Order of Meaningful Elements. In Greenberg, J., editor, *Universals of Language*, pages 73–113. MIT Press, London.

Gómez, R. and Maye, J. (2005). The developmental trajectory of nonadjacent dependency learning. *Infancy*, 7(2):183–206.

Gómez, R. L. (2002). Variability and Detection of Invariant Structure. *Psychological Science*, 13(5):431–436. Publisher: SAGE Publications Inc.

Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). The Faculty of Language: What Is It, Who Has It, and How Did It Evolve? *Science*, 298(5598):1569–1579. _eprint: https://www.science.org/doi/pdf/10.1126/science.298.5598.1569.

Hockett, C. F. (1960). The Origin of Speech. *Scientific American*, 203(3):88–97. Publisher: Scientific American, a division of Nature America, Inc.

Keizer, E. (2015). *A functional discourse grammar for English*. Oxford textbooks in linguistics. Oxford University Press, Oxford, United Kingdom, first edition. edition. Publication Title: A functional discourse grammar for English.

Kirby, S. (2001). Spontaneous evolution of linguistic structure-an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110.

Kirby, S., Cornish, H., and 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*, 105(31):10681–10686. Publisher: Proceedings of the National Academy of Sciences.

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

Kirby, S. and Hurford, J. R. (2002). The Emergence of Linguistic Structure: An Overview of the Iterated Learning Model. In Cangelosi, A. and Parisi, D., editors, *Simulating the Evolution of Language*, pages 121–147. Springer, London.

Kirby, S., Tamariz, M., Cornish, H., and Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141:87–102.

Lian, Y., Bisazza, A., and Verhoef, T. (2023). Communication Drives the Emergence of Language Universals in Neural Agents: Evidence from the Word-order/Case-marking Trade-off. *Transactions of the Association for Computational Linguistics*, 11:1033–1047.

Lian, Y., Bisazza, A., and Verhoef, T. (2025). Simulating the Emergence of Differential Case Marking with Communicating Neural-Network Agents. arXiv:2502.04038 [cs].

Nordlinger, R. (1998). *Constructive case: evidence from Australian languages*. Dissertations in linguistics. Center for the Study of Language and Information, CSLI Publications, Stanford, Calif.

Radulescu, S., , Frank, W., , and Avrutin, S. (2019). Patterns Bit by Bit. An Entropy Model for Rule Induction. *Language Learning and Development*, 16(2):109–140. Publisher: Routledge _eprint: https://doi.org/10.1080/15475441.2019.1695620.

Radulescu, S., Kotsolakou, A., Wijnen, F., Avrutin, S., and Grama, I. (2021). Fast but Not Furious. When Sped Up Bit Rate of Information Drives Rule Induction. *Frontiers in Psychology*, 12. Publisher: Frontiers Media S.A.

Raviv, L., Meyer, A., and Lev-Ari, S. (2019a). Compositional structure can emerge without generational transmission. *Cognition*, 182:151–164.

Raviv, L., Meyer, A., and Lev-Ari, S. (2019b). Larger communities create more systematic languages. *Proceedings of the Royal Society B: Biological Sciences*, 286(1907):20191262. Publisher: Royal Society.

Ren, Y., Guo, S., Labeau, M., Cohen, S. B., and Kirby, S. (2020). Compositional Languages Emerge in a Neural Iterated Learning Model. arXiv:2002.01365 [cs].

Sandler, W., Aronoff, M., Meir, I., and Padden, C. (2011). The gradual emergence of phonological form in a new language. *Natural language & linguistic theory*, 29(2):503–543.

Sandler, W., Meir, I., Padden, C., and Aronoff, M. (2005). The emergence of grammar: Systematic structure in a new language. *Proceedings of the National Academy of Sciences*, 102(7):2661–2665. Publisher: Proceedings of the National Academy of Sciences.

Smith, K. and Culbertson, J. (2020). Communicative pressures shape language during communication (not learning): Evidence from case-marking in artificial languages.

Smith, L. B. and Thelen, E. (1993). *A Dynamic Systems Approach to the Development of Cognition and Action.* MIT Press, Cambridge, MA.

Verhoef, T. (2012). The origins of duality of patterning in artificial whistled languages. *Language and cognition*, 4(4):357–380.

Verhoef, T., Kirby, S., and de Boer, B. (2014). Emergence of combinatorial structure and economy through iterated learning with continuous acoustic signals. *Journal of Phonetics*, 43:57–68.