

# Linguistic generalization and compositionality in modern artificial neural networks

Marco Baroni\*<sup>1,2,3</sup>

<sup>1</sup>Catalan Institute for Advanced Studies and Research

<sup>2</sup>Department of Translation and Language Sciences,  
Universitat Pompeu Fabra

<sup>3</sup>Facebook Artificial Intelligence Research

June 21, 2019

## Abstract

In the last decade, deep artificial neural networks have achieved astounding performance in many natural language processing tasks. Given the high productivity of language, these models must possess effective generalization abilities. It is widely assumed that humans handle linguistic productivity by means of algebraic compositional rules: Are deep networks similarly compositional? After reviewing the main innovations characterizing current deep language processing networks, I discuss a set of studies suggesting that deep networks are capable of subtle grammar-dependent generalizations, but also that they do not rely on systematic compositional rules. I argue that the intriguing behaviour of these devices (still awaiting a full understanding) should be of interest to linguists and cognitive scientists, as it offers a new perspective on possible computational strategies to deal with linguistic productivity beyond rule-based compositionality, and it might lead to new insights into the less systematic generalization patterns that also appear in natural language.

**Keywords:** artificial neural networks, deep learning, linguistic productivity, compositionality

## 1 Introduction

Neural networks have been a prominent tool to model cognitive phenomena at the mechanistic level since at least the mid eighties [1]. In the last decade, under their “deep learning” re-branding, neural networks have also proven their

---

\*Contact information: Department of Translation and Language Sciences (UPF), Carrer Roc Boronat 138, 08018 Barcelona, Spain. Email: [mbaroni@gmail.com](mailto:mbaroni@gmail.com). ORCID ID: 0000-0001-5066-3580

worth as astonishingly successful general-purpose, large-scale machine-learning algorithms [2]. In the domain of natural language, today neural networks are core components of effective machine-translation engines such as Google Translate<sup>1</sup> and DeepL.<sup>2</sup> OpenAI recently caused controversy when it announced that it would not make its new language-modeling network publicly available, as it generates novel text about arbitrary topics that is so realistic and coherent that it could easily be deployed in malicious applications, such as bulk creation of faked or abusive content.<sup>3</sup> The debate on the linguistic abilities of neural networks of the 80s and 90s involved small experiments and theoretical speculation about whether neural networks would ever be able to process language at scale (e.g., [3, 4, 5, 6, 7], among many others). Given the impressive empirical results achieved by modern neural networks, the interesting question today is not *whether*, but *how* neural networks achieve their language skills, and what causes the surprising and sometimes dramatic failures that still affect them [8, 9].

Many early neural networks were developed with the specific purpose of understanding mental processes, and thus cognitive or biological plausibility was a central concern. Modern deep networks are instead optimized for practical goals, such as better translation quality or information extraction. It is thus unlikely that their behaviour will closely mimic human cognitive processes. I contend, however, that their high natural language processing performance makes them very worth studying from the perspective of cognitive science. Following an early proposal by McCloskey [10], we should treat this as comparative psychology. Just like the communication systems of primates and other species can shed light on the unique characteristics of human language (e.g., [11, 12]), studying how artificial neural networks accomplish (or fail to accomplish) sophisticated linguistic tasks can provide important insights on the nature of such tasks, and the possible ways in which a computational device can (or cannot) solve them. This is the perspective I adopt here in looking at linguistic productivity and compositionality in deep networks.

Natural languages are characterized by immense *productivity*, in the sense that they license a theoretically infinite set of possible expressions. Linguists almost universally agree that *compositionality*, the ability to construct larger linguistic expressions by combining simpler parts, subtends productivity. The focus is typically on *semantic* compositionality, the principle whereby the meaning of a linguistic expression is a function of the meaning of its components and the rules used to combine them [13, 14]. When studying the generalization properties of neural networks, I believe it is more useful to consider a broader notion of compositionality, also encompassing, for example, the syntactic derivation rules allowing us to judge the grammaticality of nonce sentences independently of their meaning [15, 16]. Indeed, compositionality is conjectured to be a landmark not only of language but of human thought in general [17, 5, 8], and the compositional abilities of neural networks have been tested on tasks that are not semantic [7] or even linguistic in nature [18]. If a system is not compositional in

---

<sup>1</sup><https://translate.google.com/>

<sup>2</sup><https://www.deepl.com/translator>

<sup>3</sup><https://openai.com/blog/better-language-models/>

this more general sense, it will not, *a fortiori*, be able to build complex semantic representations by parallel composition of syntactic and semantic constituents.

Compositional operations in language (and thought) are argued to constitute a rule-based algebraic system, of the sort that can be formally captured by symbolic functions with variable slots. It follows that compositionality is “systematic”, in the sense that a function must apply in the same way to all variables of the right type. As famously put by Fodor, if you know the correct compositional rules to understand “*John loves Mary*”, you *must* also understand “*Mary loves John*” [5] (Fodor and colleagues distinguish systematicity and compositionality: simplifying somewhat, they see compositionality, in the stricter semantic sense presented above, as the natural consequence of applying systematic rules in the domain of natural language). Neural networks are not thought to be capable of acquiring systematic rules. Their linguistic generalization abilities have thus been the focus of much research in the past (see, e.g., [19, 20, 21, 22, 23] among many others). Note that productivity *per se* does not entail systematic compositionality. Some forms of generalization outside language are not rule-based (and not systematic). For example, similarity-driven reasoning about concept instances is probably too fuzzy and prototype-based to be accounted for by systematic rules [24]. One could also imagine a language that is productive but not (systematically) compositional. For example, Hockett [25], reflecting about the origins of language, conjectured a stage in which new expressions are formed not by systematic composition of smaller parts, but by blending unanalyzed wholes in inconsistent ways. Modern languages also exhibit many corners of non-systematic, partial productivity, a point I’ll return to in the conclusion. Still, systematic composition rules are an extremely powerful generalization mechanism. Once you know that *super-* attaches to adjectives to form other adjectives, you can in principle understand an infinite (if rather contrived) set of words: *super-good*, *super-super-good*, etc. In this context, it has been argued that lack of compositionality is one reason why modern neural networks, in striking contrast to humans, require huge amounts of data to induce correct generalizations [8].

In this article, I would like to introduce researchers interested in compositionality from a cognitive perspective to some relevant recent work about linguistic productivity in modern deep networks. After briefly reviewing the main novelties characterizing current deep language-processing architectures, I will present experimental evidence that these systems are at the same time able to capture subtle syntactic generalizations about novel forms (thus handling a sophisticated form of grammatical productivity), and failing to show convincing signs of rule-based compositionality. I will conclude with some considerations about the significance of these results for the general study of linguistic productivity.

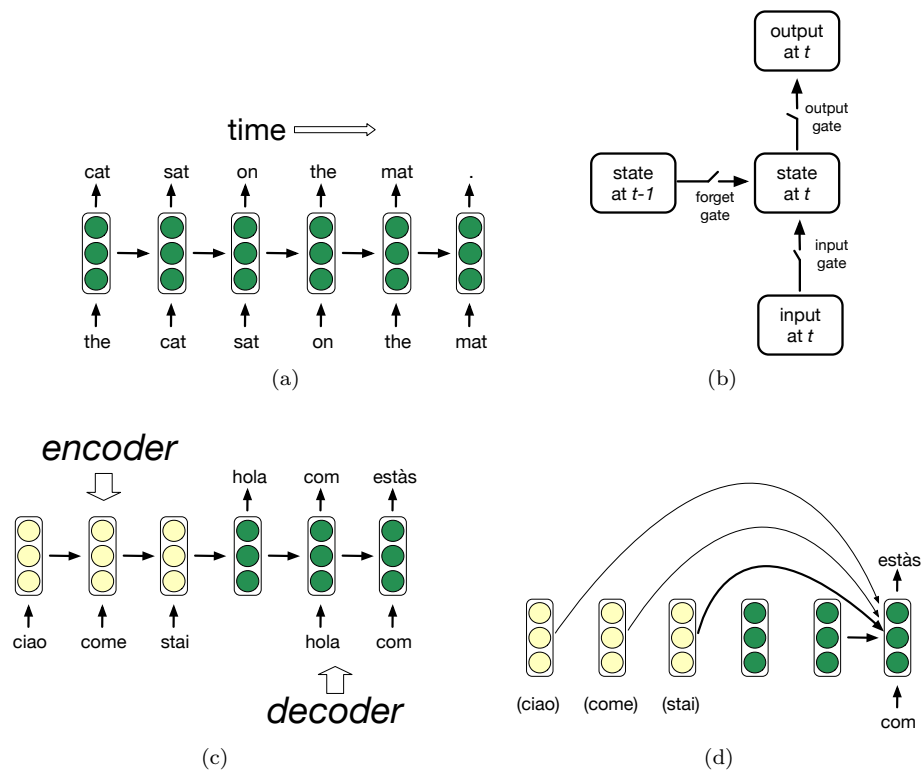


Figure 1: Architectural features of modern sequence-processing networks. (a) A traditional **recurrent network** processing a sequence of words in multiple time steps. The arrows represent sets of weighted connections (a single arrow stands for multiple unit-to-unit connections). The 3 green circles represent the same network at different time steps, which might be structured into multiple layers (not depicted in the figure). The output at time  $t$  is a (non-linear) function of the current input as well as the state of the network at time  $t-1$  (information is carried through time by the *recurrent* connections represented by vertical arrows in the figure). (b) In modern sequence-processing networks, a set of **gates** modulate the amount of information flowing through the connections of the network. (c) The **encoder-decoder** architecture is modular, with separate sub-networks (in yellow and green in the figure) trained to process input (encoder) and generate output (decoder). The decoder is typically initialized from the last state of the encoder. (d) In **attention**-enhanced architectures, the state of the encoder at each time step is separately stored in memory, and at each decoding step the network dynamically determines how much information to read from each of the memory slots. The figure schematically depicts the step in which the decoder makes a prediction based on the last word it produced (*com*), its previous state and attention-mediated memorized states from the decoder (with line thickness symbolizing the relative weight assigned to each memorized state). Unlike the previous diagrams, this one only shows the connections that are active in the last depicted processing step.

## 2 Modern deep networks for language processing: what has changed

Much of the last-decade improvements in neural network performance are due to the availability of larger training data sets that, together with computational power and better optimization methods, enabled large-scale data-driven training of complex, multi-layer architectures [2, 26]. In the domain of language, large corpora made it possible to train networks with the simple *language modeling* method [27, 28]. In this setup, the weights of a sequence-processing network are set by optimizing the objective of predicting the next word in a text, given the previous context. This is schematically illustrated for a recurrent network in Fig. 1(a). Nowadays, language modeling is used as a general-purpose way to pre-train networks to perform linguistic tasks [29]. It is also an interesting training regime from a cognitive point of view, since humans in many cultures are also exposed to large amounts of raw language data during acquisition (e.g., [30]), and predicting what comes next plays a central role in cognition [31, 32, 33, 34]. Of course, prediction is not the only task humans perform when learning a language, and how to design more varied and human-like training environments is an open research issue.

Important advances have also been made in architectural terms. The original sequence-processing recurrent network schematically illustrated in Fig. 1(a) reads some input (e.g., a word), and produces an output (e.g., a guess about the next word) at each time step. The output of the network at time  $t$  is a non-linear function of the input at time  $t$ , as well as of the state of the network itself at step  $t-1$  (weighted by *recurrent connections* that propagate activations across time) [35].

*Gated* recurrent networks, such as long short-term memory networks [36] and gated recurrent units [37], possess mechanisms regulating the dynamics of information processing across time, whose parameters are jointly induced with the rest of the network in the training phase. In particular, the network gates can “decide”, at each time step and for each unit, how much it should be updated with information about the current input (vs. preserving currently stored information), and how much it should contribute to (the hidden representation determining) the current output. Such gating mechanisms, schematically illustrated in Fig. 1(b), allow longer-term and more nuanced control of the information flow. Gates have proven empirically extremely effective, and are standard in modern language-processing networks [38].

Another important innovation consisted in decoupling input and output processing through *encoder-decoder* architectures [39]. As sketched in Fig. 1(c), separate sub-networks are trained to process the input and generate the output, with the last state of the first network (the encoder) used to initialize the second (the decoder). Input-output decoupling allows effective handling of *sequence-to-sequence* tasks, in which a sequence (e.g., a sentence in a language) has to be mapped onto another sequence (e.g., a sentence in another language), especially where input and output sequences are very different. This approach is by now

standard in machine translation (where, for example, it permits flexible mapping between languages with different word orders), but it is extremely general, and it has for example also been employed to convert linguistic instructions to actions and sentences to semantic representations [40, 41].

A (learned) *attention* mechanism, in its original form [42], automatically allows the decoder to read more or less information from different encoder states (on the basis of similarity computations between vectors representing current and past states). As schematically illustrated in Fig. 1(d), when it is about to translate the word following *com* (“*how*” in Catalan), an attention-augmented network might decide to read more from the encoder state corresponding to the word that immediately follows *come* (“*how*”) in the Italian source. Attention plays an increasingly important role in modern encoder-decoder architectures, to the point that the most successful contemporary models dispense with recurrent connections altogether, and rely instead on a rich attention mechanism to keep track of relevant past information [43].

Modern sequence-processing networks are complex systems, equipped with strong structural priors such as gates, encoding and decoding modules and attention. They should not be thought of as “*tabulae rasae*”, as they often were in early debates on connectionism. At the same time, the “innate” biases they encode are rather different from those assumed to shape human linguistic competence. Some researchers *are* trying to inject into modern networks priors closer to those traditionally postulated by linguists, such as a preference for hierarchical tree structures, (e.g., [44, 45]; see also Brennan’s contribution to this issue). Models of this latter kind have however not yet proven their worth as generic language processing devices, and I will not delve further into them. Intriguingly, [46] recently found that, when such models are not provided with explicit information about conventional compositional derivations, they come up with tree structures that do not resemble those posited by linguists at all. This is in line with the basic tenet of this paper, that neural networks might solve complex linguistic tasks, but not in the way we expect them to be solved.

### 3 Colorless green grammatical generalization in deep networks

There is no doubt that modern language-processing neural networks can generalize beyond their training data. Without such ability, their astounding performance in machine translation [43, 47] would remain unexplained, as most sentences, or even long word sequences, in any text to be translated are extremely unlikely to have ever been produced before [48]. Recently, there has been widespread interest in understanding whether this performance depends on shallow heuristics, or whether the networks are indeed capturing grammar-based generalizations, of the sort that would be supported by symbolic compositional rules (e.g., [49, 50, 51]).

Gulordava and colleagues [52] test the networks’ grammatical “intuitions” in

a setup that is strictly controlled to insure they are tapping into their productive competence (as opposed to memorized patterns). They train a gated recurrent network on large Wikipedia-derived corpora, using the language modeling objective. They then feed it minimal pairs of sentences respecting/violating long-distance number agreement. Crucially, the test items are semi-randomly generated nonsense sentences. For example, one minimal pair is: “*I realize the **wars** on which I should revise your hunt **understand/understands**” (here, the plural verb variant with *understand* is the grammatical one). The model, without further task-specific tuning, is tasked with computing the probability of the two variants, and it is said to have produced the correct judgment if it assigns a higher probability to the grammatical one. Gulordava’s nonsensical twist strips off possible semantic, lexical and collocational confounds, focusing on the abstract grammatical generalization.*

Gulordava and colleagues run the experiment in English, Hebrew, Italian and Russian. In all cases, neural networks display a preference for the grammatical sentences that is well above chance level and competitive baselines (the lowest performance occurs in English, where the network still guesses correctly 74% of the cases, where chance level is at 50%). Moreover, for Italian, they compare the network performance to human subjects taking the same test. Human accuracy turns out to be only marginally above that of the network (88.4% vs. 85.5%).

These results suggest that neural networks capture abstract, structure-based grammatical generalizations. However, the evidence is indirect, and others [53, 54] have suggested that the networks are really capitalizing on shallow heuristics (such as: “percolate the number of the first noun in a sentence to all verbs”). Lakretz and colleagues [55] conducted extensive ablation and connectivity studies of the Gulordava network. They found that the network specialized very few units to the task of carrying long-distance number information. For example, when the activation of these units is fixed to 0, the network performance on agreement tasks slide towards chance level. Importantly, these units are strongly connected to a sub-network of nodes that can be independently shown to be sensitive to hierarchical syntactic constituency. Unveiling this circuit in the network suggests that the latter has indeed developed genuine grammatical processing mechanisms, and it is not simply relying on surface heuristics when computing agreement.

The kind of productivity that was probed in these studies is grammatical in nature. Just like in Chomsky’s famous “*colorless green ideas*” example [15], Gulordava’s network can tell apart subtly different grammatical and ungrammatical nonsensical sentences, that are certainly very far from anything it was exposed to during training. Lakretz’ analysis of the network further suggests that its behaviour relies on genuine sensitivity to grammatical structure. The traditional linguistic story about a cognitive device displaying this behaviour would be that it possesses a compositional rule-based system, allowing it to reliably process novel linguistic input (such as Gulordava’s stimuli). We have no evidence, yet, about whether Gulordava’s network possesses something akin to such system, but we will now turn to experiments with similar networks directly probing their compositionality through a miniature language designed for this

purpose, that suggest that they don't.

## 4 Compositional generalization: can deep networks dax twice?

Lake and Baroni [56] introduced SCAN, a benchmark to test the compositional abilities of sequence processing networks, later extended by [57]. The SCAN miniature language is characterized by a grammar generating a large but finite number of linguistic navigation commands, and an interpretation function associating a semantic representation (a sequence of action symbols) to each possible command. The primitives of the language are verbs such as *jump* and *run*, mapped to the corresponding actions (e.g., JUMP, RUN). Primitives are combined with a set of adverb-like modifiers and conjunctions, resulting in composite expressions denoting action sequences. For example, if  $[[x]]$  is the action associated to expression “x” by the interpretation function, then “x and y” maps to  $[[x]] [[y]]$  and “x twice” maps to  $[[x]] [[x]]$ . Consequently, “*jump twice and run*” is compositionally mapped to the action sequence: JUMP JUMP RUN.

The general evaluation paradigm is as follows. An encoder-decoder network is trained on a set of SCAN commands for long enough that the network learns to accurately execute them (that is, to map them to the corresponding action sequences). The network performance is then evaluated on executing a set of test commands that were not encountered during training. Note that, unlike in the experiments reviewed in the previous section, where it had to assign a probability to pre-determined sentences, here the network has to actively produce an output action sequence, thus its generative abilities are more directly probed.

By splitting the possible SCAN commands into different training and testing partitions, we can gain insights into what the network is actually learning. I will focus here on the results obtained with 3 of the proposed splits, and their implications. In the *random* split, 80% of the commands are used for training, the remaining 20% for testing. This split checks the network ability to handle generic productivity (of the sort that might occur in standard machine-translation benchmarks), since all test expressions are new for the network. However, there is no *controlled* difference between the two sets, and the network will in general have seen a number of examples quite similar to those it has to execute. For example, the test set contains the command “*look around left twice and jump right twice*”. This is new, but the training data contain examples of “*look around left twice*” and “*jump right twice*”, both on their own and in conjoined expressions (e.g., “***look around left twice and turn left***”, “***run twice and jump right twice***”).

In the *jump* split, the training set contains all possible commands with all primitives except *jump*. During training, *jump* is presented multiple times, but only in isolation. The test set is then made of all composite commands containing *jump*. For example, at training the network is exposed to “*run twice*”



and “*walk and look*”, and at test time it must execute “*jump twice*” and “*walk and jump*”. The split is straightforward for a system possessing composition rules such as: “*x twice*” maps to  $[[x]] [[x]]$ . This is akin to a human subject learning a new verb *daxing*, and being immediately able to understand what “*dax twice*” means.

However, since *jump* only occurs in isolation at training time, the tested specimen could also reasonably conclude that the latter has a different distribution from the other verbs, and refuse to generalize it to novel composite contexts [58]. Loula and colleagues [57] introduce different partitions that control for this factor. In particular, in the *around-right* split, the training data contain all possible commands, except those where *around* is combined with *right*. Still, the network is given plenty of distributional evidence that *right* and *left* function identically otherwise, and it is exposed to many examples illustrating the behaviour of *around* in combination with *left*. The training data contain, a.o., the commands “*run left*”, “*run right*”, “*jump opposite left*”, “*jump opposite right*”, “*look around left*”. The test data include “*look around right*”.

The results in the original papers, and the further experiments of [59] with more carefully-tuned models, tell a simple story. Modern recurrent sequence-processing networks, just like conjectured by Jerry Fodor and Gary Marcus, are able to generalize in a fuzzy, similarity-based way that allows them to succeed in the *random* split (100% average test accuracy and s.d.  $\approx 0.0\%$  across multiple runs with different initializations). However, they utterly fail at the *jump* and *around-right* splits, that require inducing systematic compositional rules (12.5% accuracy with 6.6% s.d. and 2.5% accuracy with 2.7% s.d., respectively).

In very recent work, Dessì and Baroni [60] found a somewhat more intriguing pattern. They replaced the gated recurrent network architectures used in earlier SCAN work with an out-of-the-box convolutional network that dispenses with recurrence by heavily relying on attention, and that has independently been shown to achieve competitive results in machine translation ([61]; Dessì and Baroni were not able, for the time being, to trace the difference in performance between this network and the previously tried models back to their architectural differences). The new model is still able to perfectly generalize in the random split (100% test accuracy, s.d.  $\approx 0.0\%$ ), but now it reaches a surprising middle ground in the “compositional” splits: 69.2% accuracy (8.2% s.d.) with *jump* and 56.7% (10.2% s.d.) with *around-right*. As chance level accuracy in these tasks is practically 0%, the results show that the network does get some important generalizations right. Still, if it extracted the correct rules, we would expect it to be perfectly accurate, which is not the case. Dessì and Baroni initially conjectured that the network learned a subset of compositional rules, explaining its partial success. For example, it could be that the network learned how to map “*x twice*” to  $[[x]] [[x]]$ , but failed to learn the corresponding “*x thrice*” rule. However, a follow-up error analysis showed this not to be the case. The network makes errors relatively uniformly across composition frames, and, qualitatively, it does not display any trace of systematicity. For example, in the *jump* split the network executes “*jump left after walk*” correctly, but fails “*jump left after run*”. In the *around-right* split, the network can execute “*run around right*”,

but not “*walk around right*”.

Similar conclusions were drawn by Andreas [62] from a very different experiment. Andreas trained sender and receiver sequence-processing networks to play a “communication game”: The sender must describe the properties of a set of objects to the receiver through a discrete communication channel, and the receiver must reconstruct the correct target objects. The agents were trained by rewarding successful communication. Across random initializations, the sender agent developed more or less compositional codes for its messages to the receiver. Interestingly, at least some codes with low degrees of compositionality were as good at generalizing to new object descriptions as more compositional ones. Again, neural networks can be productive without being compositional.

## 5 Conclusion

When critics of classic connectionism argued that neural networks are intrinsically incapable to induce symbolic composition rules, they probably assumed that this would severely limit their practical ability to handle natural language. The empirical evidence concerning modern deep networks is surprising, as it suggests that they are extremely proficient at language, while indeed not being compositional. Note in particular that the experiments by Gulordava, Lakretz and others I reviewed in Section 3 above suggest that the linguistic proficiency of neural networks extends beyond shallow pattern recognition, to competence about structure-dependent generalizations of the sort traditionally attributed to the command of systematic compositional rules. Our current understanding of the strategies learned by these networks is very limited, and our highest priority should be to develop better analytical tools to uncover the mechanisms that lead to the detected dissociation of productive grammatical competence and systematic compositionality.

From the perspective of AI research, one central question is whether making neural networks more compositional, for example by means of more structured modular architectures [63], will also make them more adaptive and faster at learning, while not costing in terms of generality. Current deep networks are also brittle in surprising ways. For example they are easily fooled by “adversarial” examples (e.g., words with a few characters shifted) that would be trivially handled by humans [e.g., 64, 65, 66]. Explicitly compositional architectures might provide added robustness to similar attacks, or at least afford better insights into the often mysterious failings of the networks. In turn, this might lead to progress in ambitious natural language processing tasks where the success of modern deep networks are less clear-cut, such as machine reading and natural language inference [e.g., 67, 68].

Still, the way in which current models generalize without possessing anything resembling compositional rules also offer an intriguing opportunity for comparative studies to linguists and cognitive scientists. Classic and modern criticism of neural networks emphasizes the aspects of human language that are best characterized by clear-cut, algebraic rules. Language, however, is also host to

plenty of productive phenomena that obey less systematic, fuzzier laws, ranging from phonologically-driven generalizations of irregular inflections [69], to partial semantic transparency in derivational morphology [70], to semi-lexicalized constraints in syntax [71], to the early stages of grammaticalization in language change [72]. Progress in understanding the linguistic capabilities of neural networks might help us to make precise predictions about the origin, scope and mechanics of these phenomena, and ultimately to develop a more encompassing account of the amazing productivity and malleability of human language.

## Acknowledgments

I would like to thank Gemma Boleda, Roberto Dessì, Diane Bouchacourt, Emmanuel Dupoux, Kristina Gulordava, Dieuwke Hupkes, Douwe Kiela, Jean-Rémi King, Germán Kruszewski, Tal Linzen, Tomas Mikolov, Paul Smolensky, Matthijs Westera, the NTNU “*Towards mechanistic models of meaning composition*” workshop participants, my colleagues at FAIR, my co-authors, the reviewers for the Philosophical Transactions of the Royal Society and especially, Brenden Lake and Yair Lakretz for many enlightening discussions about compositionality and neural networks.

## Competing interests

I have no competing interests.

## References

- [1] David Rumelhart, James McClelland, and PDP Research Group, editors. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations*. MIT Press, Cambridge, MA, 1986.
- [2] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436–444, 2015.
- [3] Paul Smolensky. Tensor product variable binding and the representation of symbolic structures in connectionist networks. *Artificial Intelligence*, 46: 159–216, 1990.
- [4] Jeffrey Elman. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7:195–225, 1991.
- [5] Jerry Fodor and Ernest Lepore. *The Compositionality Papers*. Oxford University Press, Oxford, UK, 2002.
- [6] Steven Pinker and Micheal Ullman. The past and future of the past tense. *Trends in Cognitive Sciences*, 6:456–463, 2002.

- [7] Gary Marcus. *The Algebraic Mind*. MIT Press, Cambridge, MA, 2003.
- [8] Brenden Lake, Tomer Ullman, Joshua Tenenbaum, and Samuel Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:1–72, 2017.
- [9] Gary Marcus. Deep learning: A critical appraisal. <https://arxiv.org/abs/1801.00631>, 2018.
- [10] Michael McCloskey. Networks and theories: The place of connectionism in cognitive science. *Psychological Science*, 2(6):387–395, 1991.
- [11] Philippe Schlenker, Emmanuel Chemla, Anne Schel, James Fuller, Jean-Pierre Gautier, Jeremy Kuhn, Dunja Veselinović, Kate Arnold, Cristiane Cäsar, Sumir Keenan, Alban Lemasson, Karim Ouattara, Robin Ryder, and Klaus Zuberbühler. Formal monkey linguistics. *Theoretical Linguistics*, 42(1-2):1–90, 2016.
- [12] Simon Townsend, Sabrina Engesser, Sabine Stoll, Klaus Zuberbühler, and Balthasar Bickel. Compositionality in animals and humans. *PLOS Biology*, 16(8):1–7, 2018.
- [13] Gottlob Frege. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, 100:25–50, 1892.
- [14] Richard Montague. Universal Grammar. *Theoria*, 36:373–398, 1970.
- [15] Noam Chomsky. *Syntactic Structures*. Mouton, Berlin, Germany, 1957.
- [16] Noam Chomsky. *Aspects of the theory of syntax*. MIT Press, Cambridge, MA, 1965.
- [17] Jerry Fodor and Zenon Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28:3–71, 1988.
- [18] Brenden Lake, Ruslan Salakhutdinov, and Joshua Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [19] Morten Christiansen and Nick Chater. Generalization and connectionist language learning. *Mind & Language*, 9(3):273–287, 1994.
- [20] Gary Marcus. Rethinking eliminative connectionism. *Cognitive Psychology*, 282(37):243–282, 1998.
- [21] Steven Phillips. Are feedforward and recurrent networks systematic? analysis and implications for a connectionist cognitive architecture. *Connection Science*, 10:137–160, 1998.
- [22] Frank van der Velde, Gwendid van der Voort van der Kleij, and Marc de Kamps. Lack of combinatorial productivity in language processing with simple recurrent networks. *Connection Science*, 16(1):21–46, 2004.

- [23] Philémon Brakel and Stefan Frank. Strong systematicity in sentence processing by simple recurrent networks. In *Proceedings of CogSci*, pages 1599–1604, Amsterdam, the Netherlands, 2009.
- [24] Gregory Murphy. *The Big Book of Concepts*. MIT Press, Cambridge, MA, 2002.
- [25] Charles Hockett. The origin of speech. *Scientific American*, 203:88–111, 1960.
- [26] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, Cambridge, MA, 2016.
- [27] Alex Graves. *Supervised Sequence Labelling with Recurrent Neural Networks*. Springer, Berlin, 2012.
- [28] Tomas Mikolov. *Statistical language models based on neural networks*. Dissertation, Brno University of Technology, 2012.
- [29] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. <https://d4mucfpksywv.cloudfront.net/better-language-models/language-models.pdf>, 2019.
- [30] Thomas Landauer and Susan Dumais. A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2):211–240, 1997.
- [31] Moshe Bar. The proactive brain: using analogies and associations to generate predictions. *Trends in Cognitive Science*, 11(7):280–289, 2007.
- [32] Roger Levy. Expectation-based syntactic comprehension. *Cognition*, 106(3):1126–1177, 2008.
- [33] Martin Pickering and Simon Garrod. An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, 36(4):329–347, 2013.
- [34] Andy Clark. *Surfing Uncertainty*. Oxford University Press, Oxford, UK, 2016.
- [35] Jeffrey Elman. Finding structure in time. *Cognitive Science*, 14:179–211, 1990.
- [36] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [37] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Proceedings of the NIPS Deep Learning and Representation Learning Workshop*, Montreal, Canada, 2014. Published online: <http://www.dlworkshop.org/accepted-papers>.

- [38] Yoav Goldberg. *Neural Network Methods for Natural Language Processing*. Morgan & Claypool, San Francisco, CA, 2017.
- [39] Ilya Sutskever, Oriol Vinyals, and Quoc Le. Sequence to sequence learning with neural networks. In *Proceedings of NIPS*, pages 3104–3112, Montreal, Canada, 2014.
- [40] Li Dong and Mirella Lapata. Language to logical form with neural attention. In *Proceedings of ACL*, pages 33–43, Berlin, Germany, 2016.
- [41] Hongyuan Mei, Mohit Bansal, and Matthew Walter. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *Proceedings of AAAI*, pages 2772–2778, Phoenix, AZ, 2016.
- [42] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR Conference Track*, San Diego, CA, 2015. Published online: <http://www.iclr.cc/doku.php?id=iclr2015:main>.
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of NIPS*, pages 5998–6008, Long Beach, CA, 2017.
- [44] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of EMNLP*, pages 1631–1642, Seattle, WA, 2013.
- [45] Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah Smith. Recurrent neural network grammars. In *Proceedings of NAACL*, pages 199–209, San Diego, California, 2016.
- [46] Adina Williams, Andrew Drozdov, and Samuel Bowman. Do latent tree learning models identify meaningful structure in sentences? *Transactions of the Association for Computational Linguistics*, 6:253–267, 2018.
- [47] Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. Understanding back-translation at scale. In *Proceedings of EMNLP*, pages 489–500, Brussels, Belgium, 2018.
- [48] Marco Baroni. Distributions in text. In Anke Lüdeling and Merja Kytö, editors, *Corpus Linguistics: An International Handbook*, volume 2, pages 803–821. Mouton de Gruyter, Berlin, Germany, 2009.
- [49] Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. Assessing the ability of LSTMs to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics*, 4:521–535, 2016.
- [50] Shammur Chowdhury and Roberto Zamparelli. RNN simulations of grammaticality judgments on long-distance dependencies. In *Proceedings of COLING*, pages 133–144, Santa Fe, NM, 2018.

- [51] Ethan Wilcox, Roger Levy, Takashi Morita, and Richard Futrell. What do RNN language models learn about filler-gap dependencies? In *Proceedings of the EMNLP BlackboxNLP Workshop*, pages 211–221, Brussels, Belgium, 2018.
- [52] Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. Colorless green recurrent networks dream hierarchically. In *Proceedings of NAACL*, pages 1195–1205, New Orleans, LA, 2018.
- [53] Adhiguna Kuncoro, Chris Dyer, John Hale, and Phil Blunsom. The perils of natural behavioral tests for unnatural models: The case of number agreement. Poster presented at the Learning Language in Humans and in Machines conference, online at: <https://osf.io/view/L2HM/>, 2018.
- [54] Tal Linzen and Brian Leonard. Distinct patterns of syntactic agreement errors in recurrent networks and humans. In *Proceedings of CogSci*, pages 692–697, Austin, TX, 2018.
- [55] Yair Lakretz, Germán Kruszewski, Theo Desbordes, Dieuwke Hupkes, Stanislas Dehaene, and Marco Baroni. The emergence of number and syntax units in LSTM language models. In *Proceedings of NAACL*, pages 11–20, Minneapolis, MN, 2019.
- [56] Brenden Lake and Marco Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *Proceedings of ICML*, pages 2879–2888, Stockholm, Sweden, 2018.
- [57] Joao Loula, Marco Baroni, and Brenden Lake. Rearranging the familiar: Testing compositional generalization in recurrent networks. In *Proceedings of the EMNLP BlackboxNLP Workshop*, pages 108–114, Brussels, Belgium, 2018.
- [58] Elizabeth Wonnacott, Elissa Newport, and Michael Tanenhaus. Acquiring and processing verb argument structure: Distributional learning in a miniature language. *Cognitive Psychology*, 56(3):165–209, 2008.
- [59] Joost Bastings, Marco Baroni, Jason Weston, Kyunghyun Cho, and Douwe Kiela. Jump to better conclusions: SCAN both left and right. In *Proceedings of the EMNLP BlackboxNLP Workshop*, pages 47–55, Brussels, Belgium, 2018.
- [60] Roberto Dessì and Marco Baroni. CNNs found to jump around more skillfully than RNNs: Compositional generalization in seq2seq convolutional networks. In *Proceedings of ACL*, Firenze, Italy, 2019. In press.
- [61] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann Dauphin. Convolutional sequence to sequence learning. In *Proceedings of ICML*, pages 1243–1252, Sydney, Australia, 2017.

- [62] Jacob Andreas. Measuring compositionality in representation learning. In *Proceedings of ICLR*, New Orleans, LA, 2019. Published online: <https://openreview.net/group?id=ICLR.cc/2019/conference>.
- [63] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Learning to compose neural networks for question answering. In *Proceedings NAACL*, pages 1545–1554, San Diego, California, 2016.
- [64] Yonatan Belinkov and Yonatan Bisk. Synthetic and natural noise both break neural machine translation. In *Proceedings of ICLR Conference Track*, Vancouver, Canada, 2018. Published online: <https://openreview.net/group?id=ICLR.cc/2018/Conference>.
- [65] Javid Ebrahimi, Daniel Lowd, and Dejing Dou. On adversarial examples for character-level neural machine translation. In *Proceedings of COLING*, pages 653–663, Santa Fe, New Mexico, USA, 2018.
- [66] Zhengli Zhao, Dheeru Dua, and Sameer Singh. Generating natural adversarial examples. In *Proceedings of ICLR Conference Track*, Vancouver, Canada, 2018. Published online: <https://openreview.net/group?id=ICLR.cc/2018/Conference>.
- [67] Robin Jia and Percy Liang. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of EMNLP*, pages 2021–2031, Copenhagen, Denmark, 2017.
- [68] Thomas McCoy, Ellie Pavlick, and Tal Linzen. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of ACL*, Firenze, Italy, 2019. In press.
- [69] Adam Albright and Bruce Hayes. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition*, 90:119–161, 2003.
- [70] Marco Marelli and Marco Baroni. Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, 122(3):485–515, 2015.
- [71] Adele Goldberg and Ray Jackendoff. The English resultative as a family of constructions. *Language*, 80(3):532–568, 2004.
- [72] Paul Hopper and Elizabeth Traugott. *Grammaticalization*. Cambridge University Press, Cambridge, UK, 1990.