

On the proper role of linguistically-oriented deep net analysis in linguistic theorizing

Marco Baroni^{1,2,3}

¹Facebook AI Research

²Universitat Pompeu Fabra

³Catalan Institution for Research and Advanced Studies

June 16, 2021

Abstract

A lively research field has recently emerged that uses experimental methods to probe the linguistic behavior of modern deep networks. While work in this tradition often reports intriguing results about the grammatical skills of deep nets, it is not clear what their implications for linguistic theorizing should be. As a consequence, linguistically-oriented deep net analysis has had very little impact on linguistics at large. In this chapter, I suggest that deep networks should be treated as theories making explicit predictions about the acceptability of linguistic utterances. I argue that, if we overcome some obstacles standing in the way of seriously pursuing this idea, we will gain a powerful new theoretical tool, complementary to mainstream algebraic approaches.

1 Introduction

During the last decade, deep neural networks have come to dominate the field of natural language processing (NLP) (Sutskever et al., 2014; Vaswani et al., 2017; Devlin et al., 2019). While earlier approaches to NLP relied on tools, such as part-of-speech taggers and parsers, that extracted linguistic knowledge from explicit manual annotation of text corpora (Jurafsky and Martin, 2008), deep-learning-based methods typically adopt an “end-to-end” approach: A deep net is directly trained to associate some form of natural linguistic input (e.g., text in a language) to a corresponding linguistic output (e.g., the same text in a different language), dispensing with the traditional pipeline of intermediate linguistic modules and the related annotation of latent linguistic structure (e.g., syntactic parses of source and target sentences) (Goldberg, 2017; Lappin, 2021).

This paradigm shift has important implications for the relation between theoretical and computational linguistics. The issue of which linguistic formalisms

might provide the best annotation schemes to develop effective NLP tools is no longer relevant. Instead, the last few years have seen the raise of a new field of investigation consisting in the experimental analysis of the grammatical skills of deep nets trained without the injection of any explicit linguistic knowledge. For the remainder of the chapter, I will refer to this research area as LODNA, for *linguistically-oriented deep net analysis*. LODNA takes the perspective of a psycholinguist (Futrell et al., 2019), or perhaps more accurately that of an ethologist (McCloskey, 1991; Scholte, 2016), designing sophisticated experiments to “probe” the knowledge implicit in a species’ behavior.

LODNA is currently a very active research field, with many papers focusing on whether neural networks have correctly induced specific kinds of grammatical generalization (e.g., Linzen et al., 2016; Chowdhury and Zamparelli, 2018; Futrell et al., 2019; Chaves, 2020), as well as benchmarks attempting to probe their linguistic competence at multiple levels (Conneau et al., 2018; Warstadt et al., 2019). LODNA papers account for a significant proportion of the work presented at annual events such as the Society for Computation in Linguistics conference and the BlackBox NLP workshop.

LODNA is well-motivated from a machine-learning perspective. Understanding how a system behaves is a prerequisite to improve it, and might be important in the perspective of AI safety and explainability (Belinkov and Glass, 2019; Xie et al., 2020). However, there is no doubt that the grammatical performance of deep nets is also extremely intriguing from a linguistic perspective, particularly because the architectural primitives of these models (such as distributed representations and structures that linearly propagate information across time) are profoundly different from those postulated in linguistics (such as categorical labels and tree structures). Still, as we will see below, for all the enthusiasm for LODNA within NLP, this line of work is hardly having any impact on the current debate in theoretical linguistics.

In this chapter, after introducing, as an example of LODNA, the by-now “classic” domain of long-distance agreement probing (Section 2), I will present evidence for the claim that this sort of research, despite the intriguing patterns it uncovers, is hardly affecting contemporary linguistics (Section 3), I will argue that this gap stems from lack of clarity about its theoretical significance (Section 4). In particular, I will show that modern deep networks cannot be treated as blank slates meant to falsify innateness claims. They should rather be seen as algorithmic linguistic theories making predictions about utterance acceptability. I will then outline several issues that are currently standing in the way of taking deep nets seriously as linguistic theories. I will conclude in Section 5 by briefly discussing why taking this stance might be beneficial to computational and theoretical linguistics, and by sketching two possible ways to pursue LODNA-based linguistic theorizing.

2 Linguistic-oriented analysis of deep nets: The case of long-distance agreement

Linguists identify sensitivity to syntactic structure that is not directly observable in the signal as one of the core properties of human grammatical competence (Everaert et al., 2015). A paradigmatic test for structure sensitivity comes from agreement phenomena. For example, subject-verb number agreement in an English clause depends on the c-command relation between the subject noun phrase and the corresponding verb, and it is not affected by nouns intervening between the NP head and the verb:

- (1) [The **kid** [near the *toys* in the *boxes*]] **is** tired.

In (1), an example of *long-distance agreement*, the fact that two plural nouns (*toys* and *boxes*) directly precede the main verb *is* does not affect its number, as the only noun that entertains the right relation with the verb is *kid*. As Everaert and colleagues' (2015) motto goes, it's all about *structures*, not strings!

Current deep network architectures, such as long-short-term memory networks (LSTMs, Hochreiter and Schmidhuber, 1997), convolutional networks (CNNs, Kalchbrenner et al., 2014) or Transformers (Vaswani et al., 2017), do not encode any prior favoring a structural analysis of their input over a sequential one. It is natural then to ask whether they are able to correctly handle structure-dependent phenomena, such as long-distance agreement. Consequently, starting with the influential work of Linzen et al. (2016), long-distance agreement has become a standard test to probe their linguistic abilities.

Probably the most thorough analysis of long-distance number agreement in deep networks was the one we carried out in Gulordava et al. (2018). We focused on LSTMs trained as *language models*. That is, the networks were trained by exposing them to large amounts of textual data (samples from the Wikipedias of the relevant languages), with the task of predicting the next word given the context. No special tuning for the long-distance agreement challenge was applied. After this generic training, the networks were presented with sentence prefixes up to and excluding the second item in an agreement relation (e.g., *The kid near the toys in the boxes. . .*), and the probability they assigned to continuations with the right or wrong agreement (*is/are*) was measured. The experiment was conducted with a test set of genuine corpus examples in 4 languages (English, Hebrew, Italian and Russian), and considering various agreement relations (not only noun-verb but, also, for example, verb-verb and adjective-noun). The networks got the right agreement with high accuracy in all languages and for all constructions.

Even more impressively, the networks were also able to get agreement right when tested with nonsense sentences such as the one in (2), showing that they must extract syntactic rules at a rather abstract level.

- (2) The **carrot** around the *lions* for the *disasters. . . sings/*sing*.

Finally, we compared the agreement accuracy of the Italian network with that

of native speakers (both on corpus-extracted and nonsense sentences), finding that the network is only marginally lagging behind human performance.

Other studies tested different deep architectures, such as CNNs (Bernardy and Lappin, 2017) and Transformers (Goldberg, 2019b), confirming that they also largely succeed at long-distance agreement.

Deep nets have been tested for a number of other linguistic generalizations, such as those pertaining to filler-gap constructions, auxiliary fronting and case assignment. See Linzen and Baroni (2021) for a recent survey of LODNA specifically aimed at linguists.¹ In pretty much all cases, while they departed here and there from human intuition, deep nets captured at least the general gist of the phenomena being investigated.

3 The gap

Results such as the ones on long-distance agreement I briefly reviewed should be of interest to theoretical linguists, since, as already mentioned, deep nets possess very different priors from those postulated by linguists as part of the universal language faculty, such as a predisposition for hierarchical structures (Hauser et al., 2002; Berwick and Chomsky, 2016; Adger, 2019). In reality, however, the growing body of work on LODNA is almost completely ignored by the current theoretical linguistics debate.

To sustain this claim with quantitative evidence, I looked at the impact of Tal Linzen’s original paper on long-distance agreement in deep nets (Linzen et al., 2016). This is a highly-cited paper, having amassed 514 Google Scholar citations in less than 5 years.² I sifted through these citations, keeping track of how many came from theoretical linguistics (under a very broad notion of what counts as theoretical linguistics). I found that only 6 citations qualified. Of these, 3 were opinion pieces, one of them written by Linzen himself. Note that the article does not lack general interdisciplinary appeal, as shown by many citations from psycho- and neuro-linguistics, and even 4 citations from the field of computational agricultural studies!

Perhaps Google Scholar does a poor job at tracking theoretical linguistic work. Indeed, David Adger’s recent *Language Unlimited* volume (Adger, 2019) does extensively discuss Linzen’s article, but I did not find it among the studies citing it according to Scholar. Thus, as a supplementary source of evidence, I also downloaded all papers from the front page of LingBuzz, a popular linguistics preprint archive.³ I filtered out papers that do not qualify as theoretical linguistics. Again, I tried to be inclusive: I excluded, for example, one paper about the aftermath of the “Pinker LSA letter” controversy (Kastner et al., 2021), but I did include one about phonosymbolism in Pokémon character

¹Lappin (2021) also provides a review of some of the relevant work, as part of a book-length treatment of the linguistic and cognitive implications of deep learning models.

²Google Scholar queried on May 27th, 2021.

³Papers downloaded from <https://ling.auf.net/lingbuzz> on May 27th, 2021. I downloaded all *freshly changed* and *new* papers, as well as all the papers in the *Top Recent Downloads* and *Last 6 months* sections of the front page.

names (Kawahara et al., 2021). This left me with a corpus of 37 papers. I then went through their bibliographies, looking for references to deep learning work, and finding... none!⁴

It is not fair to impute this lack of references to a possible endogamous bent of theoretical linguistics. To the contrary, the papers in my mini-corpus reveal considerable interdisciplinary breadth, with frequent references to neuroscience, ethology, psycholinguistics and sociolinguistics; they include statistical treatments of experimental and corpus data; and they use sophisticated computational tools, such as graph-theoretical analyses. It is really NLP, and in particular deep-learning-based NLP, that is missing from the party.

To understand this gap, we need to ask: why should linguists care about the grammatical analysis of deep networks? What is it supposed to tell us about human linguistic competence? In other words, what is the theoretical significance of LODNA?

4 The theoretical significance of linguistically-oriented deep net analysis

When LODNA researchers situate their work within a broader theoretical context, it is invariably in terms of nature-or-nurture arguments resting on a view of deep nets as blank slates. For example, when asked about the significance of his work for theoretical linguistics, Tal Linzen told me that deep-net simulations “can help linguists focus on the aspects [...] that truly require explanation in terms of innate constraints. If the simulation shows that there is plenty of data for the learner to acquire a particular phenomenon, maybe there’s nothing to explain!” (Tal Linzen, p.c.).

Similar claims are sprinkled throughout LODNA papers. Here are just a few examples (from otherwise excellent papers): “Our results also contribute to the long-running nature-nurture debate in language acquisition: whether the success of neural models implies that unbiased learners can learn natural languages with enough data, or whether human abilities to acquire language given sparse stimulus implies a strong innate human learning bias” (Papadimitriou and Jurafsky, 2020). “The APS [(argument from the poverty of the stimulus)] predicts that any artificial learner trained with no prior knowledge of the principles of syntax [...] must fail to make acceptability judgments with human-level accuracy. [...] If linguistically uninformed neural network models achieve human-level performance on specific phenomena [...], this would be clear evidence limiting the scope of phenomena for which the APS can hold” (Warstadt et al., 2019). “[I]f such a device [(a neural network)] could manage to replicate fine-grained human intuitions inducing them from the raw training input this would be evidence that exposure to language structures [...] should in principle be sufficient

⁴I had performed a similar experiment in March 2021, by collecting papers from the latest issues of *Linguistic Inquiry*, *Natural Language and Linguistic Theory* and *Syntax*, with the very same outcome (no reference whatsoever to deep learning work).

to derive a syntactic competence, against the innatist hypothesis” (Chowdhury and Zamparelli, 2018).

Deep nets are linguistic theories, not blank slates

If blank slate arguments were (perhaps) valid when looking at the simple connectionist models of the eighties (Rumelhart et al., 1986; Churchland, 1989; Clark, 1989), all modern deep networks possess highly-structured innate architectures that considerably weaken such arguments. Consider, for example, the Transformer (Vaswani et al., 2017), the current darling of NLP. A Transformer network is structured into a number of layered modules, each involving a complex bank of linear and non-linear transformations. These, in turn, differ in profound ways from the innate structure of a LSTM (Hochreiter and Schmidhuber, 1997). For example, a LSTM will read a sentence one word at a time, and will use a recurrent function to preserve information across time, whereas the Transformer will read the whole sentence in parallel, and use an extended backward and forward attention system to incorporate contextual information.

Even more importantly, as demonstrated by the widespread interest of NLP and machine learning researchers in proposing new architectures, differences in the supposedly “weak” and “general” biases of different deep nets lead them to behave very differently, given the same input data.

A striking illustration of this was recently provided by Kharitonov and Chaabouni (2021), in a study of *seq2seq* deep nets, that is, networks trained to associate input and output sequences (as in, e.g., a translation task).

Kharitonov and Chaabouni trained such networks on really tiny corpora that severely underspecify the input-output relation. The test-time behavior of the network in cases where different generalizations lead to different outputs was then inspected, to reveal which innate preferences the networks brought to the task.

In one of their experiments, the whole training corpus consists of the following *input* \rightarrow *output* examples.

- (3) aabaa \rightarrow b
 bbabb \rightarrow a
 aaaaa \rightarrow a
 bbbbb \rightarrow b

The mini-corpus in (3) is compatible with (at least) two rules: a “hierarchical” one, stating that the output is generated by taking the character in the middle of the input; and a “linear” generalization, stating that the output is the third character in the input sequence.⁵

After training it with just the examples in (3), a network is exposed to a new input where the two rules lead to different predictions, e.g., *aaabaaa*, where the hierarchical generalization would pick *b* and the linear one *a*.

⁵This can be seen as a schematic reproduction of classic poverty-of-the-stimulus thought experiments, such as the one built around English auxiliary fronting by Chomsky (1965).

Of four widely-used seq2seq models tested by Kharitonov and Chaabouni, two (LSTMs with attention and Transformers) show a strong preference for the hierarchical generalization, and two (LSTMs without attention and CNNs) show a strong preference for the linear generalization.

Studies such as this invalidate any blank-slate claim about deep nets. It is more appropriate, instead, to look at deep nets as *linguistic theories*, encoding non-trivial structural priors facilitating language acquisition and processing.

Most linguists agree, explicitly or implicitly, that a linguistic theory is a computational system that, given an input utterance in a language, can predict whether the sequence would be acceptable to an idealized speaker of the language, typically under the view that such computational system is a high-level representation of the speaker’s mental knowledge of grammar (e.g., Chomsky, 1986; Sag et al., 2003; Müller, 2020). And a deep net trained on text with the standard language-modeling (that is, next-word prediction) objective will indeed do just that.⁶

It is undoubtedly easier to inspect the inner workings of a symbolic linguistic theory than those of a trained deep net. However, the classic objection that artificial neural networks can’t be taken seriously as theories of mental faculties because they are unopenable black boxes (e.g., McCloskey, 1991) is much weaker today, in light of extensive progress in the development of methods to analyze network states and behaviors (Belinkov and Glass, 2019), including the whole LODNA tradition.

Why don’t we see, then, many articles positioning deep nets as alternative or complementary theories to traditional grammatical formalism? I believe that two crucial ingredients are still missing, before deep nets can seriously contribute to contemporary linguistic theorizing.

The problem of low commitment to models

Differences between deep nets, as we have discussed above, are huge. *Mutatis mutandis*, the difference between an LSTM, reading a word at a time and building a joint representation through its recurrent state, and a Transformer, processing all words in parallel to create multiple context-weighted representations, might be as large as that between a derivational and a constraint-based theory in formal linguistics, if not bigger. Other differences, such as those between the memory structures used by LSTMs and GRU recurrent networks (Cho et al., 2014), might not be as large, but certainly making similarly discrepant structural assumptions about, say, possible tree configurations would be considered a big deal in formal linguistics.

And, yet, researchers investigating the linguistic behavior of these architectures almost never provide a theoretically grounded motivation for why they

⁶Just like in linguistics (e.g., Murphy, 2007; Lau et al., 2017; Sprouse and Schütze, 2019), there is considerable debate on the best way to elicit acceptability judgments from the models to compare them to human data, and on whether such judgments should be probabilistic or categorical (e.g., Linzen et al., 2016; Chowdhury and Zamparelli, 2018; Warstadt et al., 2019; Niu and Penn, 2020).

focused on one architecture or the other. Interest tends to shift with the state of the art in applied tasks such as machine translation or natural language inference. So, if nearly all early LODNA papers focused on LSTMs and GRUs, nowadays the field has nearly entirely shifted to analyzing Transformer networks, not because the latter were found to be more plausible models of human language processing (if anything, their ability to read and process massive windows of text in parallel makes them *less* plausible models than recurrent networks), but because they became the mainstream approach in applied NLP, thanks to their astounding performance in applied tasks.

As a concrete illustration of this phenomenon, we can compare the LODNA papers from the first (2018) and third (2020) editions of the BlackBox NLP workshop (one of the core events in the area).⁷ In the 2018 edition, I found 13 full papers that broadly qualify as LODNA. Of them, 12 focus on LSTM analysis, with the remaining one already looking at the Transformer. In two years, the balance has completely shifted. All 9 relevant papers in the 2020 editions analyze some variant of the Transformer, with two also including LSTM variants among the comparison models. Importantly, in none of these papers there is a linguistically-oriented (or even engineering-oriented) discussion of *why* the Transformer was picked over the LSTM or other architectures. Indeed, in a few cases, earlier work that was based on LSTMs is cited as corroborating evidence, only mentioning in passing that this earlier work was based on a profoundly different architecture.

The problem is mostly sociological: NLP puts a strong (and reasonable) emphasis on whichever models work best in applications, and consequently analytical work will also tend to concentrate on such models. However, if radical changes in the underlying architecture are not motivated by linguistic considerations, and indeed they tend to be completely glossed over, it is hard to take this work seriously from the perspective of linguistic theorizing.⁸

The problem of paucity of interesting predictions

A good linguistic theory should not only fit what is already known about a language, but also make predictions about previously unexplored patterns. This is the typical *modus operandi*, for example, in generative syntax, where, e.g., hypotheses about possible syntactic configurations lead to strong typological predictions about acceptable adverb and adjective orders (e.g., Cinque, 1999, 2010).

The standard approach in LODNA, instead, is to check whether models capture well-known patterns, such as vanilla English subject-verb number agreement. The occasional focus on cases outside the standard paradigm is typically

⁷<https://www.aclweb.org/anthology/volumes/W18-54/>;
<https://www.aclweb.org/anthology/volumes/2020.blackboxnlp-1/>

⁸There are important exceptions. Work that does put an emphasis on the linguistic motivation of architectural choices includes that of Chris Dyer and colleagues on recurrent neural network grammars (e.g., Kuncoro et al., 2018b), and that of Paul Smolensky and colleagues on tensor product decomposition networks (e.g., McCoy et al., 2019).

meant to highlight obviously *wrong* predictions made by the model (e.g., Kuncoro et al., 2018a, show how in some syntactic configurations LSTMs let the verb agree with the first noun in a sentence even if it is not its subject).

What we are doing, then, is an extensive (and important!) sanity check of our systems, rather than using them to widen the coverage of linguistic phenomena we are able to explain through computational modeling.

In order to move from sanity checking to employing our models as interesting-predictions generators, there are at least two issues we need to overcome. The first is that apparently insignificant changes in the very same model or in the way it is trained can lead to qualitatively different linguistic behavior.

For example, McCoy et al. (2020a) study how seq2seq recurrent networks generalize auxiliary fronting. Consider how English yes/no questions are formed by moving an auxiliary to the first position of the sentence, as in:

- (4) a. The zebra **does** chuckle.
- b. **Does** the zebra chuckle?

Chomsky (1965) observed that examples such as (4) are compatible with *two* generalizations: move the auxiliary of the *main* clause, as in (5-b); but also: move the *first* auxiliary in the sentence, as in the ungrammatical (5-c):

- (5) a. Your zebras that *don't* dance **do** chuckle.
- b. **Do** your zebras that *don't* dance chuckle?
- c. **Don't* your zebras that dance **do** chuckle?

Chomsky (1965) famously built a poverty-of-the-stimulus argument around these examples, noting that children rarely or never hear examples such as (5-b) when learning English, and yet they unfailingly go for the structure-sensitive generalization (move the *main* auxiliary).

McCoy and colleagues created a similar poverty-of-the-stimulus setup for LSTMs. Among other results, they found, strikingly, that the very same network, depending of different random initializations of its weights, could converge to a strong preference for the *move-main* or the *move-first* generalization.

McCoy et al. (2020b) showed that Transformer networks sharing a large proportion of weights, with random initialization only affecting their topmost linear layers and the order in which training examples are presented, can converge to accuracies between 0% and 66% in an experiment probing the ability to generalize across natural language inference data-sets.

Liska et al. (2018) reported that the very same recurrent network trained from the same initialization could discover or fail to discover compositionality, simply depending on different random orders in which the training data were presented.

In all these cases, our deep net “theories” make radically different predictions based on very trivial perturbations of the training process, which makes it difficult to trust such predictions. What’s worse, in none of these cases we understand *how* the relevant small, random differences in initialization or training schedule catastrophically affect the outcome. This leads to my second point:

In order to use our models as generators of interesting predictions, we need to achieve a good mechanistic understanding of which components lead them to adopt a certain behavior: something which is seldom done in LODNA.

As an example of the kind of study combining a granular understanding of the model inner working with a non-trivial prediction tested in humans, I will briefly summarize the detailed analysis of deep network behavior with respect to long-distance number agreement that we reported in Lakretz et al. (2019, 2021).

In the first of these studies, a cell-by-cell analysis of LSTMs performing the subject-verb agreement task revealed that they develop a sparse mechanism to store and propagate a single number feature between subject and verb. This sparse grammar-aware circuit is complemented by a distributed system that can fill in the number feature based on purely sequential heuristics.

This leads to an interesting prediction for sentences with two embedded long-distance dependencies, such as:

(6) The **kid**₁ that the **dogs**₂ near the *toy like*₂/***likes**₂ **is**₁/***are**₁ tired.

Here, the sparse grammar-aware mechanism will be activated when *kid* is encountered, and, due to its sparsity, it will not be able to also record the number of *dogs*. Consequently, once *like(s)* is encountered, the heuristic distributed system will take over, and it will wrongly predict the singular form, since the sequentially closer noun is *toy*. On the other hand, once *is/are* is reached, the feature stored in the sparse long-distance circuit can be released, correctly predicting a preference for singular *is*. This is an interesting prediction because, intuitively, the longer distance *kid-is* relation should be harder to track than the shorter-distance one connecting *dogs* and *like*.

In Lakretz et al. (2021), we proceeded to test the prediction both in LSTMs and with human subjects. We did indeed find the predicted inner/outer agreement asymmetry both in machines and (more weakly) in humans. This suggests that agreement might be implemented by means of sparse feature-carrying mechanisms in humans as well.

Arguably, this study focused on linguistic performance rather than competence. Nobody would claim that “*The kid that the dogs near the toy likes is tired*” is grammatical—it is simply an easier mistake to make during online language processing. Still, this work points the way towards how deep nets could profitably be used to lead linguistic theorizing through an in-depth process of model analysis, hypothesis formulation and human subject testing.

Lakretz’ study took about 4 years to run. By the time it was completed, it presented a detailed analysis of a model, the LSTM, that many in NLP would find obsolete. Its focus on a single grammatical construction might look quaint, now that the field has moved towards large-scale evaluation suites probing models on a variety of phenomena and tasks (e.g., Conneau and Kiela, 2018; Conneau et al., 2018; Marvin and Linzen, 2018; Wang et al., 2019; Warstadt et al., 2019). Yet, if we want to reach the sort of understanding of a deep model’s inner working that can be useful to gain new insights on human linguistic competence

and behavior, we should have more studies running at the same slow, thorough, narrow-focused pace of this project!

5 Conclusion

Language models based on deep network architectures such as the LSTM and the Transformer are computational devices that, by being exposed to large amounts of natural text, learn to assign probabilities to arbitrary word sequences. In the last five years or so, a rich tradition of studies has emerged that analyzes such models in order to understand what kind of grammatical competence they possess.

The results of these studies are often intriguing, revealing the sophisticated linguistic skills of deep nets, as well as interesting error patterns. However, such studies have had very little impact on theoretical linguistics.

I attributed this gap to the fact that similar studies typically lack a clear theoretical standing and, when they do, it is one based on the wrong idea that we should treat modern deep nets as *tabulae rasae* lacking strong innate priors. Deep nets do possess such innate priors, as shown by the fact that different models trained on the same data can extract dramatically different generalizations. I proposed that a more solid theoretical standing for the linguistic analysis of deep nets can be achieved by treating them as *algorithmic linguistic theories*.

I discussed above some concrete roadblocks we must overcome if we want to seriously adopt this stance. I will conclude by briefly explaining why I think that such stance is beneficial for both computational and theoretical linguists, and by providing quick sketches of how deep-net-based linguistic theorizing could look like.

Why should computational linguists care?

The incredible progress in deep learning for NLP we've seen in the last few years must be entirely credited to NLP and machine-learning practitioners interested in solving concrete challenges such as machine translation. Ideas from theoretical linguistics have played no role in the area (Lappin, 2021), and there is no clear reason, in turn, why computational linguists interested in practical NLP technologies should care about the implications of their work for linguistics.

However, the success of events such as the already mentioned Society for Computation in Linguistics conference and BlackBox NLP workshop, as well as the fact that all major NLP conferences now feature special tracks on linguistic analysis of computational models, suggest that there is a significant sub-community of computational linguists who *are* interested in the linguistic implications of deep learning models.

These researchers should be bothered by the fact that their work is not having an impact on mainstream theoretical linguistics. Clarifying the theoretical status of deep net simulations, and in particular boldly presenting them as alternative linguistic theories, might finally attract due attention from the linguistics

community.

Why should theoretical linguists care?

Deep nets attained incredible empirical results in tasks that heavily depend on linguistic knowledge, such as machine translation (Edunov et al., 2018), well beyond what was ever achieved by symbolic or hybrid systems. While it is possible that deep nets are relying on a completely different approach to language processing than the one encoded in human linguistic competence, theoretical linguists should investigate what are the building blocks making these systems so effective: if not for other reasons, at least in order to explain why a model that is supposedly encoding completely different priors than those programmed into the human brain should be so good at handling human language.

I conjecture however that deep nets and traditional symbolic theories are both valid algorithmic approaches to modeling human linguistic competence, and that they are complementary in the aspects they best explain. The more algebraic features of language, such as recursive structures, are elegantly handled by traditional linguistic formalisms such as generative syntax (Müller, 2020) and formal semantics (Heim and Kratzer, 1998). However, language has other facets, in particular those where the fuzzy, large-scale knowledge that characterizes the lexicon is involved, where such theories struggle. Neural language models, by inducing a large set of context-dependent and fuzzy patterns from natural input, and by being inherently able to probabilistically generate and process text, should be better equipped to handle phenomena such as polysemy, the partial productivity of morphological derivation, non-fully-compositional phrase formation and diachronic shift (e.g., Marelli and Baroni, 2015; Vecchi et al., 2017; Lenci, 2018; Boleda, 2020).⁹

From this angle, the current emphasis of LODNA on exactly those phenomena (such as long-distance agreement) that are already satisfactorily captured by traditional algebraic models might be misguided. Curiously, even staying within the syntax domain, there is no work I am aware of focusing instead on those patterns, such as partially lexicalized constructions (e.g., Goldberg, 2005, 2019a), where the fuzzier rules typically learned by neural networks might give us novel insights into human generalization.

Do neural network theories require a switch from algebraic to distributed models of linguistic competence?

The main topic of this volume is the role of algebraic systems in the representation of linguistic knowledge. By proposing a trained Transformer, with its billions of weights and its continuous activation vectors, as a linguistic theory, I am *de facto* implying that the appropriate level to represent linguistic knowledge is not algebraic, but massively distributed. This requires a radical

⁹These references mostly discuss a precursor of neural language models known as *distributional semantics*, but the same accounts could be replicated and extended using latest-generation neural language models.

methodological shift in the way linguistic models are studied. Standard rule- or constraint-based systems can easily be probed by direct inspection. With deep networks, model probing requires sophisticated experiments: indeed, the whole area of LODNA can be seen as an example of how to gain insights from a linguistic theory based on distributed representations.

However, I would like to leave the issue of the right level for deep-net-based linguistic theorizing open. Optimality Theory (Prince and Smolensky, 2004) was the most fruitful outcome of early attempts to bring together linguistics and connectionism. Optimality Theory is an algebraic approach whose principles are inspired by how linguistic constraints might be implemented by a neural network. Could the way in which LSTMs or Transformers process linguistic information similarly inspire a symbolic theory of language? Perhaps, one that is not based on tree structures but on storage and retrieval mechanisms akin to gating and attention?

To conclude, despite the criticism I vented to some aspects of the field, I think that LODNA is one of the most exciting things that happened to cognitive science in the last five years. I hope that, once we clarify its theoretical standing, the body of evidence assembled in this area will finally have the impact it deserves on linguistics at large.

Acknowledgments

I would like to thank Jelke Bloem, Grzegorz Chrupała, Ido Dagan, Roberto Dessí, Emmanuel Dupoux, Dieuwke Hupkes, Shalom Lappin, Yair Lakretz, Paola Merlo, the members of the UPF Computational Linguistics and Linguistic Theory group, the participants in the EACL 2021 Birds-of-a-Feather Meetup on Linguistic Theories, the audience at EACL 2021 and, especially, David Adger, Gemma Boleda, Roberta D’Alessandro, Tal Linzen, Louise McNally and Adina Williams for a mixture of advice, stimulating discussion and constructive feedback.

References

- David Adger. *Language Unlimited: The Science Behind Our Most Creative Power*. Oxford University Press, Oxford, UK, 2019.
- Yonatan Belinkov and James Glass. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 7:49–72, 2019.
- Jean-Philippe Bernardy and Shalom Lappin. Using deep neural networks to learn syntactic agreement. *Linguistic Issues in Language Technology*, 15(2): 1–15, 2017.

- Robert Berwick and Noam Chomsky. *Why Only Us: Language and Evolution*. MIT Press, Cambridge, MA, 2016.
- Gemma Boleda. Distributional semantics and linguistic theory. *Annual Review of Linguistics*, 6:213–234, 2020.
- Rui Chaves. What don’t RNN language models learn about filler-gap dependencies? In *Proceedings of SCiL*, pages 20–30, New Orleans, LA, 2020.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of EMNLP*, pages 1724–1734, Doha, Qatar, 2014.
- Noam Chomsky. *Aspects of the theory of syntax*. MIT Press, Cambridge, MA, 1965.
- Noam Chomsky. *Knowledge of Language: Its Nature, Origin, and Use*. Praeger, Westport, CT, 1986.
- 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.
- Paul Churchland. *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. MIT Press, Cambridge, MA, 1989.
- Guglielmo Cinque. *Title Adverbs and Functional Heads: A Cross-Linguistic Perspective*. Oxford University Press, Oxford, UK, 1999.
- Guglielmo Cinque. *The Syntax of Adjectives*. MIT Press, Cambridge, MA, 2010.
- Andy Clark. *Microcognition: Philosophy, Cognitive Science, and Parallel Distributed Processing*. MIT Press, Cambridge, MA, 1989.
- Alexis Conneau and Douwe Kiela. SentEval: An evaluation toolkit for universal sentence representations. In *Proceedings of LREC*, pages 1699–1704, Miyazaki, Japan, 2018.
- Alexis Conneau, Germán Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. What you can cram into a single $\&\!#*$ vector: Probing sentence embeddings for linguistic properties. In *Proceedings ACL*, pages 2126–2136, Melbourne, Australia, 2018.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL*, pages 4171–4186, Minneapolis, MN, 2019.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. Understanding back-translation at scale. In *Proceedings of EMNLP*, pages 489–500, Brussels, Belgium, 2018.

- Martin Everaert, Marinus Huybregts, Noam Chomsky, Robert Berwick, and Johan Bolhuis. Structures, not strings: Linguistics as part of the cognitive sciences. *Trends in Cognitive Sciences*, 19(12):729–743, 2015.
- Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. Neural language models as psycholinguistic subjects: Representations of syntactic state. In *Proceedings of NAACL*, pages 32–42, Minneapolis, MN, 2019.
- Adele Goldberg. *Constructions at Work: The Nature of Generalization in Language*. Oxford University Press, Oxford, UK, 2005.
- Adele Goldberg. *Explain Me This: Creativity, Competition, and the Partial Productivity of Constructions*. Princeton University Press, Princeton, NJ, 2019a.
- Yoav Goldberg. *Neural Network Methods for Natural Language Processing*. Morgan & Claypool, San Francisco, CA, 2017.
- Yoav Goldberg. Assessing BERT’s syntactic abilities. <https://arxiv.org/abs/1901.05287>, 2019b.
- 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.
- Marc Hauser, Noam Chomsky, and Tecumseh Fitch. The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298(5598):1569–1579, 11 2002.
- Irene Heim and Angelika Kratzer. *Semantics in Generative Grammar*. Blackwell, Malden, MA, 1998.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- Dan Jurafsky and James Martin. *Speech and Language Processing, 2nd ed.* Prentice Hall, Upper Saddle River, 2008.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. A convolutional neural network for modelling sentences. In *Proceedings of ACL*, pages 655–665, Baltimore, MD, 2014.
- Itamar Kastner, Hadas Kotek, Anonymous Anonymous, Rikker Dockum, Michael Dow, Maria Esipova, Caitlin Green, and Todd Snider. Who speaks for us? Lessons from the Pinker letter. <https://ling.auf.net/lingbuzz/005381>, 2021.
- Shigeto Kawahara, Gakuji Kumagai, and Mahayana Godoy. English speakers can infer Pokémon types using sound symbolism. <https://ling.auf.net/lingbuzz/005129>, 2021.

- Eugene Kharitonov and Rahma Chaabouni. What they do when in doubt: A study of inductive biases in seq2seq learners. In *Proceedings of ICLR*, Online event, 2021. Proceedings at: <https://openreview.net/group?id=ICLR.cc/2021/Conference>.
- 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/>, 2018a.
- Adhiguna Kuncoro, Chris Dyer, John Hale, Dani Yogatama, Stephen Clark, and Phil Blunsom. LSTMs can learn syntax-sensitive dependencies well, but modeling structure makes them better. In *Proceedings of ACL*, pages 1426–1436, Melbourne, Australia, 2018b.
- 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.
- Yair Lakretz, Dieuwke Hupkes, Alessandra Vergallito, Marco Marelli, Marco Baroni, and Stanislas Dehaene. Mechanisms for handling nested dependencies in neural-network language models and humans. *Cognition*, 2021. In press.
- Shalom Lappin. *Deep learning and linguistic representation*. CRC Press, Boca Raton, FL, 2021.
- Jey Han Lau, Alexander Clark, and Shalom Lappin. Grammaticality, acceptability, and probability: A probabilistic view of linguistic knowledge. *Cognitive Science*, 41(5):1202–1241, 2017.
- Alessandro Lenci. Distributional models of word meaning. *Annual Review of Linguistics*, 4(14):1–21, 2018.
- Tal Linzen and Marco Baroni. Syntactic structure from Deep Learning. *Annual Review of Linguistics*, 7:195–212, 2021.
- 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.
- Adam Liska, Germán Kruszewski, and Marco Baroni. Memorize or generalize? searching for a compositional RNN in a haystack. In *Proceedings of the ICML AEGAP Workshop*, Stockholm, Sweden, 2018.
- Marco Marelli and Marco Baroni. Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, 122(3):485–515, 2015.

- Rebecca Marvin and Tal Linzen. Targeted syntactic evaluation of language models. In *Proceedings of EMNLP*, pages 1192–1202, Brussels, Belgium, 2018.
- Michael McCloskey. Networks and theories: The place of connectionism in cognitive science. *Psychological Science*, 2(6):387–395, 1991.
- Thomas McCoy, Tal Linzen, Ewan Dunbar, and Paul Smolensky. RNNs implicitly implement tensor-product representations. In *Proceedings of ICLR*, New Orleans, LA, 2019. Published online: <https://openreview.net/group?id=ICLR.cc/2019/conference>.
- Thomas McCoy, Robert Frank, and Tal Linzen. Does syntax need to grow on trees? sources of hierarchical inductive bias in sequence-to-sequence networks. *Transactions of the Association for Computational Linguistics*, 8:125–140, 2020a.
- Thomas McCoy, Junghyun Min, and Tal Linzen. BERTs of a feather do not generalize together: Large variability in generalization across models with similar test set performance. In *Proceedings of the EMNLP BlackboxNLP Workshop*, pages 217–227, Online event, 2020b.
- Stefan Müller. *Grammatical Theory: From Transformational Grammar to Constraint-Based Approaches, 4th ed.* Language Science Press, Berlin, Germany, 2020.
- Brian Murphy. *A Study of Notions of Participation and Discourse in Argument Structure Realisation*. Dissertation, Trinity College Dublin, 2007.
- Jingcheng Niu and Gerald Penn. Grammaticality and language modelling. In *Proceedings of the EMNLP Eval4NLP Workshop*, pages 110–119, Online event, 2020.
- Isabel Papadimitriou and Dan Jurafsky. Learning music helps you read: Using transfer to study linguistic structure in language models. In *Proceedings EMNLP*, pages 6829–6839, Online event, 2020.
- Alan Prince and Paul Smolensky. *Optimality Theory*. Blackwell, Malden, MA, 2004.
- 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.
- Ivan Sag, Thomas Wasow, and Emily Bender. *Syntactic Theory: A Formal Introduction*. CSLI, Stanford, CA, 2003.
- Steven Scholte. Fantastic DNimals and where to find them. *NeuroImage*, 180: 112–113, 2016.

- Jon Sprouse and Carson Schütze. Grammar and the use of data. In Bas Aarts, Jill Bowie, and Gergana Popova, editors, *The Oxford Handbook of English Grammar*, pages 40–58. Oxford University Press, Oxford, UK, 2019.
- Ilya Sutskever, Oriol Vinyals, and Quoc Le. Sequence to sequence learning with neural networks. In *Proceedings of NIPS*, pages 3104–3112, Montreal, Canada, 2014.
- 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.
- Eva Maria Vecchi, Marco Marelli, Roberto Zamparelli, and Marco Baroni. Spicy adjectives and nominal donkeys: Capturing semantic deviance using compositionality in distributional spaces. *Cognitive Science*, 41(1):102–136, 2017.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of ICLR*, New Orleans, LA, 2019. Published online: <https://openreview.net/group?id=ICLR.cc/2019/conference>.
- Alex Warstadt, Amanpreet Singh, and Samuel Bowman. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641, 2019.
- Ning Xie, Gabrielle Ras, Marcel van Gerven, and Derek Doran. Explainable deep learning: A field guide for the uninitiated. <https://arxiv.org/abs/2004.14545>, 2020.