

Probing Linguistic Systematicity

Emily Goodwin,^{*1,5} Koustuv Sinha,^{2,3,5} Timothy J. O’Donnell^{1,4,5}

¹Department of Linguistics, ²School of Computer Science, McGill University, Canada

³Facebook AI Research (FAIR), Montreal ⁴Canada CIFAR AI Chair, Mila

⁵Quebec Artificial Intelligence Institute (Mila)

{emily.goodwin, koustuv.sinha}@mail.mcgill.ca

timothy.odonnell@mcgill.ca

Abstract

Recently, there has been much interest in the question of whether deep natural language understanding models exhibit *systematicity*—generalizing such that units like words make consistent contributions to the meaning of the sentences in which they appear. There is accumulating evidence that neural models often generalize non-systematically. We examined the notion of systematicity from a linguistic perspective, defining a set of probes and a set of metrics to measure systematic behaviour. We also identified ways in which network architectures can generalize non-systematically, and discuss why such forms of generalization may be unsatisfying. As a case study, we performed a series of experiments in the setting of natural language inference (NLI), demonstrating that some NLU systems achieve high overall performance despite being non-systematic.

1 Introduction

Language allows us to express and comprehend a vast variety of novel thoughts and ideas. This creativity is made possible by *compositionality*—the linguistic system builds utterances by combining an inventory of primitive units such as morphemes, words, or idioms (the *lexicon*), using a small set of structure-building operations (the *grammar*; Camap, 1947; Fodor and Pylyshyn, 1988; Hodges, 2012; Janssen et al., 2012; Lake et al., 2017b; Szabó, 2012; Zdrozny, 1994; Lake et al., 2017a).

One property of compositional systems, widely studied in the cognitive sciences, is the phenomenon of *systematicity*. Systematicity refers to the fact that lexical units such as words make consistent contributions to the meaning of the sentences in which they appear. Fodor and Pylyshyn

(1988) provided a famous example: If a competent speaker of English knows the meaning of the sentence *John loves the girl*, they also know the meaning of *The girl loves John*. This is because for speakers of English knowing the meaning of the first sentence implies knowing the meaning of the individual words *the*, *loves*, *girl*, and *John* as well as grammatical principles such as how transitive verbs take their arguments. But knowing these words and principles of grammar implies knowing how to compose the meaning of the second sentence.

Deep learning systems now regularly exhibit very high performance on a large variety of natural language tasks, including machine translation (Wu et al., 2016; Vaswani et al., 2017), question answering (Wang et al., 2018; Henaff et al., 2016), visual question answering (Hudson and Manning, 2018), and natural language inference (Devlin et al., 2018; Storks et al., 2019). Recently, however, researchers have asked whether such systems generalize systematically (see §4).

Systematicity is the property whereby words have consistent contributions to composed meaning; the alternative is the situation where words have a high degree of *contextually conditioned meaning variation*. In such cases, generalization may be based on *local heuristics* (McCoy et al., 2019b; Niven and Kao, 2019), *variegated similarity* (Albright and Hayes, 2003), or *local approximations* (Veldhoen and Zuidema, 2017), where the contribution of individual units to the meaning of the sentence can vary greatly across sentences, interacting with other units in highly inconsistent and complex ways.

This paper introduces several novel probes for testing systematic generalization. We employ an artificial language to have control over systematicity and contextual meaning variation. Applying our probes to this language in an NLI setting reveals

*Corresponding author

that some deep learning systems which achieve very high accuracy on standard holdout evaluations do so in ways which are non-systematic: the networks do not consistently capture the basic notion that certain classes of words have meanings which are consistent across the contexts in which they appear.

The rest of the paper is organized as follows. §2 discusses degrees of systematicity and contextually conditioned variation; §3 introduces the distinction between open- and closed-class words, which we use in our probes. §5 introduces the NLI task and describes the artificial language we use; §6 discusses the models that we tested and the details of our training setup; §7 introduces our probes of systematicity and results are presented in §8.¹

2 Systematicity and Contextual Conditioning

Compositionality is often stated as the principle that the meaning of an utterance is determined by the meanings of its parts and the way those parts are combined (see, e.g., Heim and Kratzer, 2000).

Systematicity, the property that words mean the same thing in different contexts, is closely related to compositionality; nevertheless, compositional systems can vary in their degree of systematicity. At one end of the spectrum are systems in which primitive units contribute exactly one identical meaning across all contexts. This high degree of systematicity is approached by artificial formal systems including programming languages and logics, though even these systems don't fully achieve this ideal (Cantwell Smith, 1996; Dutilh Novaes, 2012).

The opposite of systematicity is the phenomenon of *contextually conditioned variation in meaning* where the contribution of individual words varies according to the sentential contexts in which they appear. Natural languages exhibit such context dependence in phenomena like homophony, polysemy, multi-word idioms, and co-compositionality. Nevertheless, there are many words in natural language—especially closed-class words like quantifiers (see below)—which exhibit very little variability in meaning across sentences.

At the other end of the spectrum from programming languages and logics are systems where many or most meanings are highly context dependent.

¹Code for datasets and models can be found here: <https://github.com/emilygoodwin/systematicity>

The logical extreme—a system where each word has a different and unrelated meaning every time it occurs—is clearly of limited usefulness since it would make generalization impossible. Nevertheless, learners with sufficient memory capacity and flexibility of representation, such as deep learning models, can learn systems with very high degrees of contextual conditioning—in particular, higher than human language learners. An important goal for building systems that learn and generalize like people is to engineer systems with inductive biases for the right degree of systematicity. In §8, we give evidence that some neural systems are likely too biased toward allowing contextually conditioned meaning variability for words, such as quantifiers, which do not vary greatly in natural language.

3 Compositional Structure in Natural Language

Natural language distinguishes between *content* or *open-class* lexical units and *function* or *closed-class* lexical units. The former refers to categories, such as nouns and verbs, which carry the majority of contentful meaning in a sentence and which permit new coinages. Closed-class units, by contrast, carry most of the grammatical structure of the sentence and consist of things like inflectional morphemes (like pluralizing *-s* in English) and words like determiners, quantifiers, and negation (e.g., *all*, *some*, *the* in English). These are mostly fixed; adult speakers do not coin new quantifiers, for example, the way that they coin new nouns.

Leveraging this distinction gives rise to the possibility of constructing probes based on *jabberwocky-type sentences*. This term references the poem *Jabberwocky* by Lewis Carroll, which combines nonsense open-class words with familiar closed-class words in a way that allows speakers to recognize the expression as well formed. For example, English speakers identify a contradiction in the sentence *All Jabberwocks flug, but some Jabberwocks don't flug*, without a meaning for *jabberwock* and *flug*. This is possible because we expect the words *all*, *some*, *but*, and *don't* to contribute the same meaning as they do when combined with familiar words, like *All pigs sleep, but some pigs don't sleep*.

Using jabberwocky-type sentences, we tested the generalizability of certain closed-class word representations learned by neural networks. Giving the networks many examples of each construction

with a large variety of different content words—that is, large amounts of highly varied evidence about the meaning of the closed-class words—we asked during the test phase how fragile this knowledge is when transferred to new open-class words. That is, our probes combine novel open-class words with familiar closed-class words, to test whether the closed-class words are treated systematically by the network. For example, we might train the networks to identify contradictions in pairs like *All pigs sleep; some pigs don't sleep*, and test whether the network can identify the contradiction in a pair like *All Jabberwocks flug; some Jabberwocks don't flug*. A systematic learner would reliably identify the contradiction, whereas a non-systematic learner may allow the closed-class words (*all, some, don't*) to take on contextually conditioned meanings that depend on the novel context words.

4 Related Work

There has been much interest in the problem of systematic generalization in recent years (Bahdanau et al., 2019; Bentivogli et al., 2016; Lake et al., 2017a,b; Gershman and Tenenbaum, 2015; McCoy et al., 2019a; Veldhoen and Zuidema, 2017; Soulos et al., 2019; Prasad et al., 2019; Richardson et al., 2019; Johnson et al., 2017, *inter alia*).

In contrast to our approach (testing novel words in familiar combinations), many of these studies probe systematicity by testing familiar words in novel combinations. Lake and Baroni (2018) adopt this approach in semantic parsing with an artificial language known as SCAN. Dasgupta et al. (2018, 2019) introduce a naturalistic NLI dataset, with test items that shuffle the argument structure of natural language utterances. In the inductive logic programming domain, Sinha et al. (2019) introduced the CLUTTR relational-reasoning benchmark. The novel-combinations-of-familiar-words approach was formalized in the CFQ dataset and associated distribution metric of Keysers et al. (2019). Ettinger et al. (2018) introduced a semantic-role-labeling and negation-scope labeling dataset, which tests compositional generalization with novel combinations of familiar words and makes use of syntactic constructions like relative clauses. Finally, Kim et al. (2019) explore pre-training schemes' abilities to learn prepositions and wh-words with syntactic transformations (two kinds of closed-class words which our work does not address).

A different type of systematicity analysis directly

investigates learned representations, rather than developing probes of model behavior. This is done either through visualization (Veldhoen and Zuidema, 2017), training a second network to approximate learned representations using a symbolic structure (Soulos et al., 2019) or as a *diagnostic classifier* (Giulianelli et al., 2018), or reconstructing the semantic space through similarity measurements over representations (Prasad et al., 2019).

5 Study Setup

5.1 Natural Language Inference

We make use of the *Natural language inference* (NLI) task to study the question of systematicity. The NLI task is to infer the relation between two sentences (the *premise* and the *hypothesis*). Sentence pairs must be classified into one of a set of predefined logical relations such as *entailment* or *contradiction*. For example, the sentence *All mammals growl* entails the sentence *All pigs growl*. A rapidly growing number of studies have shown that deep learning models can achieve very high performance in this setting (Evans et al., 2018; Conneau et al., 2017; Bowman et al., 2014; Yoon et al., 2018; Kiela et al., 2018; Munkhdalai and Yu, 2017; Rocktäschel et al., 2015; Peters et al., 2018; Parikh et al., 2016; Zhang et al., 2018; Radford et al., 2018; Devlin et al., 2018; Storks et al., 2019).

5.2 Natural Logic

We adopt the formulation of NLI known as *natural logic* (MacCartney and Manning, 2014, 2009; Lakoff, 1970). Natural logic makes use of seven logical *relations* between pairs of sentences. These are shown in Table 1. These relations can be interpreted as the set theoretic relationship between the extensions of the two expressions. For instance, if the expressions are the simple nouns *warthog* and *pig*, then the entailment relation (\sqsubset) holds between these extensions (*warthog* \sqsubset *pig*) since every warthog is a kind of pig.

For higher-order operators such as quantifiers, relations can be defined between sets of possible worlds. For instance, the set of possible worlds consistent with the expression *All blickets wug* is a subset of the set of possible worlds consistent with the logically weaker expression *All red blickets wug*. Critically, the relationship between composed expressions such as *All X Y* and *All P Q* is determined entirely by the relations between X/Y and P/Q, respectively. Thus, natural logic allows

us to compute the relation between the whole expressions using the relations between parts. We define an artificial language in which such alignments are easy to compute, and use this language to probe deep learning systems’ ability to generalize systematically.

Symbol	Name	Example	Set-theoretic definition
$x \equiv y$	equivalence	<i>pig</i> \equiv <i>pig</i>	$x = y$
$x \sqsubset y$	forward entailment	<i>pig</i> \sqsubset <i>mammal</i>	$x \subset y$
$x \supset y$	reverse entailment	<i>mammal</i> \supset <i>pig</i>	$x \supset y$
$x \wedge y$	negation	<i>pig</i> \wedge <i>not pig</i>	$x \cap y = \emptyset \wedge x \cup y = U$
$x \mid y$	alternation	<i>pig</i> \mid <i>cat</i>	$x \cap y = \emptyset \wedge x \cup y \neq U$
$x \sim y$	cover	<i>mammal</i> \sim <i>not pig</i>	$x \cap y \neq \emptyset \wedge x \cup y = U$
$x \# y$	independence	<i>hungry</i> $\#$ <i>warthog</i>	(all other cases)

Table 1: MacCartney and Manning (2009)’s implementation of natural logic relations

5.3 The Artificial Language

In our artificial language, sentences are generated according to the six-position template shown in Table 2, and include a quantifier (position 1), noun (position 3), and verb (position 6), with optional pre- and post-modifiers (position 2 and 4) and optional negation (position 5). For readability, all examples in this paper use real English words; however, simulations can use uniquely identified abstract symbols (i.e., generated by `gensym`).

We compute the relation between position-aligned pairs of sentences in our language using the natural logic system (described in §5.2). Quantifiers and negation have their usual natural-language semantics in our artificial language; pre- and post-modifiers are treated intersectively. Open-class items (nouns and verbs) are organized into linear hierarchical taxonomies, where each open-class word is the sub- or super-set of exactly one other open-class item in the same taxonomy. For example, since *dogs* are all *mammals*, and all mammals *animals*, they form the entailment hierarchy *dogs* \sqsubset *mammals* \sqsubset *animals*. We vary the number of distinct noun and verb taxonomies according to an approach we refer to as *block* structure, described in the next section.

5.4 Block Structure

In natural language, most open-class words do not appear with equal probability with every other word. Instead, their distribution is biased and clumpy, with words in similar topics occurring together. To mimic such topic structure, we group nouns and verbs into *blocks*. Each block consists of six nouns and six verbs, which form taxonomic hi-

erarchies (e.g., *lizards/animals*, *run/move*). Nouns and verbs from different blocks have no taxonomic relationship (e.g., *lizards* and *screwdrivers* or *run* and *read*) and do not co-occur in the same sentence pair. Because each block includes a six verbs and six nouns in a linear taxonomic hierarchy, no single block is intrinsically harder to learn than any other block.

The same set of closed-class words appear with all blocks of open-class words, and their meanings are systematic regardless of the open-class words (nouns and verbs) they are combined with. For example, the quantifier *some* has a consistent meaning when it is applied to *some screwdrivers* or *some animals*. Because closed-class words are shared across blocks, models are trained on extensive and varied evidence of their behaviour. We present closed-class words in a wide variety of sentential contexts, with a wide variety of different open-class words, to provide maximal pressure against overfitting and maximal evidence of their consistent meaning.

5.5 Test and Train Structure

We now describe the structure of our training blocks, holdout test set, and jabberwocky blocks. We also discuss our two test conditions, and several other issues that arise in the construction of our dataset.

Training set: For each training block, we sampled (without replacement) one sentence pair for every possible combination of open-class words, that is, every combination of nouns and verbs $\langle \text{noun}_1, \text{noun}_2, \text{verb}_1, \text{verb}_2 \rangle$. Closed-class words were sampled uniformly to fill each remaining positions in the sentence (see Table 2). A random subset of 20% of training items were reserved for validation (early stopping) and not used during training.

Holdout test set: For each training block, we sampled a holdout set of forms using the same nouns and verbs, but disjoint from the training set just described. The sampling procedure was identical to that for the training blocks. These holdout items allow us to test the generalization of the models with known words in novel configurations (see §8.1).

Jabberwocky test set: Each *jabberwocky block* consisted of novel open-class items (i.e., nouns and verbs) that did not appear in training blocks. For each jabberwocky block, we began by following a

Position	1	2	3	4	5	6
Category	quantifier	nominal premodifier	noun	nominal postmodifier	negation	verb
Status	Obligatory	Optional	Obligatory	Optional	Optional	Obligatory
Class	Closed	Closed	Open	Closed	Closed	Open
Example	<i>All</i>	<i>brown</i>	<i>dogs</i>	<i>that bark</i>	<i>don't</i>	<i>run</i>

Table 2: A template for sentences in the artificial language. Each sentence fills the obligatory positions 1, 3, and 6 with a word: a quantifier, noun, and verb. Optional positions (2, 4 and 5) are filled by either a word (adjective, postmodifier or negation) or by the empty string. Closed-class categories (Quantifiers, adjectives, post modifiers and negation) do not include novel words, while open-class categories (nouns and verbs) includes novel words that are only exposed in the test set.

sampling procedure identical to that for the training/holdout sets with these new words. Several of our systematicity probes are based on the behavior of neighboring pairs of test sentences (see §7). To ensure that all such necessary pairs were in the jabberwocky test set, we extended the initial sample with any missing test items.

Training conditions: Since a single set of closed-class words is used across all blocks, adding more blocks increases evidence of the meaning of these words without encouraging overfitting. To study the effect of increasing evidence in this manner, we use two training conditions: *small* with 20 training blocks and *large* with 185 training blocks. Both conditions contained 20 jabberwocky blocks. The small condition consisted of 51,743 training, 10,399 validation, and 3,694,005 test (holdout and jabberwocky) pairs. The large condition consisted of 478,649 training, 96,005 validation, and 3,694,455 test items.

Balancing: One consequence of the sampling method is that logical relations will not be equally represented in training. In fact, it is impossible to simultaneously balance the distributions of syntactic constructions, logical relations, and instances of words. In this trade-off, we chose to balance the distribution of open-class words in the vocabulary, as we are focused primarily on the ability of neural networks to generalize closed-class word meaning. Balancing instances of open-class words provided the greatest variety of learning contexts for the meanings of the closed-class items.

6 Simulations

6.1 Models

We analyze performance on four simple baseline models known to perform well on standard NLI tasks, such as the Stanford Natural Language Inference datasets, (Bowman et al., 2015). Following Conneau et al. (2017), the hypothesis u and

premise v are individually encoded by neural sequence encoders such as a long short-term memory (LSTM; Hochreiter and Schmidhuber, 1997) or gated recurrent unit (GRU; Cho et al., 2014). These vectors, together with their element-wise product $u * v$ and element-wise difference $u - v$ are fed into a fully connected multilayer perceptron layer to predict the relation. The encodings u and v are produced from an input sentence of M words, w_1, \dots, w_M , using a recurrent neural network, which produces a set of a set of M hidden representations h_1, \dots, h_t , where $h_t = f(w_1, \dots, w_M)$. The sequence encoding is represented by its last hidden vector h_T .

The simplest of four models sets f to be a bidirectional gated recurrent unit (BGRU). This model concatenates the last hidden state of a GRU run forwards over the sequence and the last hidden state of GRU run backwards over the sequence, for example, $u = [\overleftarrow{h}_M, \overrightarrow{h}_M]$.

Our second embedding system is the Infersent model reported by Conneau et al. (2017), a bidirectional LSTM with max pooling (INFS). This is a model where f is an LSTM. Each word is represented by the concatenation of a forward and backward representation: $h_t = [\overleftarrow{h}_t, \overrightarrow{h}_t]$. We constructed a fixed vector representation of the sequence h_t by selecting the maximum value over each dimension of the hidden units of the words in the sentence.

Our third model is a self-attentive sentence encoder (SATT) which uses an attention mechanism over the hidden states of a BiLSTM to generate the sentence representation (Lin et al., 2017). This attention mechanism is a weighted linear combination of the word representations, denoted by $u = \sum_M \alpha_i h_i$, where the weights are calculated as follows:

$$\begin{aligned} \bar{h}_i &= \tanh(W h_i + b_w) \\ \alpha_i &= \frac{e^{\bar{h}_i^\top u_w}}{\sum_i e^{\bar{h}_i^\top u_w}} \end{aligned}$$

where, u_w is a learned context query vector and (W, b_w) are the weights of an affine transformation. This self-attentive network also has multiple views of the sentence, so the model can attend to multiple parts of the given sentence at the same time.

Finally, we test the Hierarchical ConvolutionalNetwork (CONV) architecture from (Conneau et al., 2017) which is itself inspired from the model *AdaSent* (Zhao et al., 2015). This model has four convolution layers; at each layer the intermediate representation u_i is computed by a max-pooling operation over feature maps. The final representation is a concatenation $u = [u_1, \dots, u_l]$ where l is the number of layers.

7 Probing Systematicity

In this section, we study the systematicity of the models described in §6.1. Recall that systematicity refers to the degree to which words have consistent meaning across different contexts, and is contrasted with contextually conditioned variation in meaning. We describe three novel probes of systematicity which we call the *known word perturbation probe*, the *identical open-class words probe*, and the *consistency probe*.

All probes take advantage of the distinction between closed-class and open-class words reflected in the design of our artificial language, and are performed on sentence pairs with novel open-class words (jabberwocky-type sentences; see §5.5). We now describe the logic of each probe.

7.1 Known Word Perturbation Probe

We test whether the models treat the meaning of closed-class words systematically by perturbing correctly classified jabberwocky sentence pairs with a closed-class word. More precisely, for a pair of closed-class words w and w' , we consider test items which can be formed by substitution of w by w' in a correctly classified test item. We allow both w and w' to be any of the closed-class items, including quantifiers, negation, nominal post-modifiers, or the empty string ϵ (thus modeling insertions and deletions of these known, closed-class items). Suppose that Example 1 was correctly classified. Substituting *some* for *all* in the premise of yields Example 2, and changes the relation from entailment (\sqsubset) to reverse entailment (\sqsupset).

- (1) **All** blickets wug.
All blockets wug.

- (2) **Some** blickets wug.
All blockets wug.

There are two critical features of this probe. First, because we start from a correctly-classified jabberwocky pair, we can conclude that the novel words (e.g., *wug* and *blickets* above) were assigned appropriate meanings.

Second, since the perturbation only involves closed-class items which do not vary in meaning and have been highly trained, the perturbation should not affect the models ability to correctly classify the resulting sentence pair. If the model does misclassify the resulting pair, it can only be because a perturbed closed-class word (e.g., *some*) interacts with the open-class items (e.g., *wug*), in a way that is different from the pre-perturbation closed-class item (i.e., *all*). This is non-systematic behavior.

In order to rule out trivially correct behavior where the model simply ignores the perturbation, we consider only perturbations which result in a change of class (e.g., $\square \mapsto \square$) for the sentence pair. In addition to accuracy on these perturbed items, we also examine the variance of model accuracy on probes across different blocks. If a model’s accuracy varies depending only on the novel open-class items in a particular block, this provides further evidence that it does not treat word meaning systematically.

7.2 Identical Open-class Words Probe

Some sentence pairs are classifiable without any knowledge of the novel words’ meaning; for example, pairs where premise and hypothesis have identical open-class words. An instance is shown in Example 3: the two sentences must stand in contradiction, regardless of the meaning of *blicket* or *wug*.

- (3) All blickets wug.
Some blickets don’t wug.

The closed-class items and compositional structure of the language is sufficient for a learner to deduce the relationships between such sentences, even with unfamiliar nouns and verbs. Our second probe, the *identical open-class words probe*, tests the models’ ability to correctly classify such pairs.

7.3 Consistency Probe

Consider Examples 4 and 5, which present the same two sentences in opposite orders.

- (4) All blickets wug.
All red blickets wug.
- (5) All red blickets wug.
All blickets wug.

In Example 4, the two sentences stand in an entailment (\sqsubset) relation. In Example 5, by contrast, the two sentences stand in a reverse entailment (\sqsupset) relation. This is a logically necessary consequence of the way the relations are defined. Reversing the order of sentences has predictable effects for all seven natural logic relations: in particular, such reversals map $\sqsubset \mapsto \sqsupset$ and $\sqsupset \mapsto \sqsubset$, leaving all other relations intact. Based on this observation, we develop a consistency probe of systematicity. We ask for each correctly classified jabberwocky block test item, whether the corresponding reversed item is also correctly classified. The intuition behind this probe is that whatever meaning a model assumes for the novel open-class words, it should assume the same meaning when the sentence order is reversed. If the reverse is not correctly classified, then this is strong evidence of contextual dependence in meaning.

8 Results

In this section, we report the results of two control analyses, and that of our three systematicity probes described above.

8.1 Analysis I: Holdout Evaluations

We first establish that the models perform well on novel configurations of known words. Table 3 reports accuracy on heldout sentence pairs, described in §5.5. The table reports average accuracies across training blocks together with the standard deviations of these statistics. As can be seen in the table, all models perform quite well on holdout forms across training blocks, with very little variance. Because these items use the same sampling scheme and vocabulary as the trained blocks, these simulations serve as a kind of upper bound on the performance and a lower bound on the variance that we can expect from the more challenging jabberwocky-block-based evaluations below.

8.2 Analysis II: Distribution of Novel Words

Our three systematicity probes employ jabberwocky-type sentences—novel open-class words in sentential frames built from known closed-class words. Since models are not

Condition	BGRU	CONV	SATT	INFS
	mean (sd)	mean (sd)	mean (sd)	mean (sd)
small	95.1 ±0.21	95.43 ±0.12	93.14 ±0.94	96.02 ±0.51
large	95.09 ±1.03	95.22 ±0.55	94.89 ±1.09	96.17 ±0.74

Table 3: Accuracy on holdout evaluations (training conditions and holdout evaluation are explained in §5.5)




Figure 1: Visualization of trained and novel open-class word embeddings.

trained on these novel words, it is important to establish that they are from the same distribution as the trained words and, thus, that the models’ performance is not driven by some pathological feature of the novel word embeddings.

Trained word embeddings were initialized randomly from $\mathcal{N}(0, 1)$ and then updated during training. Novel word embeddings were simply drawn from $\mathcal{N}(0, 1)$ and never updated. Figure 1 plots visualizations of the trained and novel open-class word embeddings in two dimensions, using t-SNE parameters computed over all open-class words (Maaten and Hinton, 2008). Trained words are plotted as +, novel words as •. Color indicates the proportion of test items containing that word that were classified correctly. As the plot shows, the two sets of embeddings overlap considerably. Moreover, there does not appear to be a systematic relationship between rates of correct classification for items containing novel words and their proximity to trained words. We also performed a resampling analysis, determining that novel vectors did not differ significantly in length from trained vectors ($p = 0.85$). Finally, we observed mean and standard deviation of the pairwise cosine similarity between trained and novel words to be 0.999 and 0.058 respectively, confirming that there is little evidence the distributions are different.

8.3 Analysis III: Known Word Perturbation Probe

Recall from §7.1 that the known word perturbation probe involves insertion, deletion, or substitution

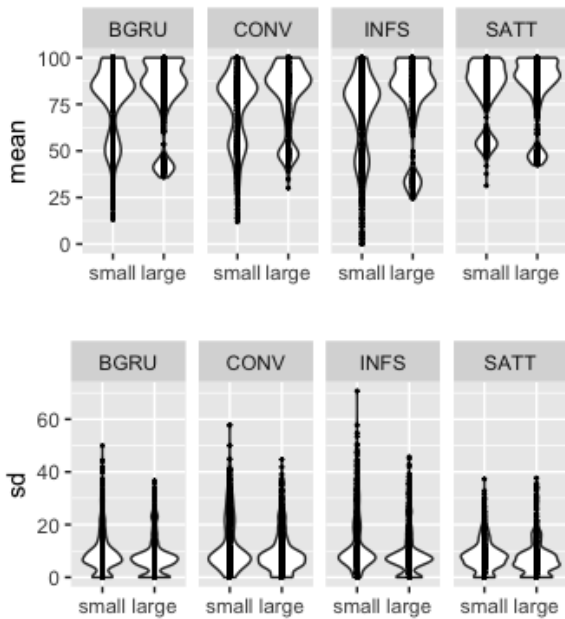


Figure 2: Performance on the known word perturbation probe, small and large training conditions (see §5.5).

of a trained closed-class word in a correctly classified jabberwocky-type sentence pair. Figure 2 plots the results of this probe. Each point represents a *perturbation type*—a group of perturbed test items that share their before/after target perturbed closed-class words and before/after relation pairs. The upper plot displays the mean accuracy of all perturbations, averaged across blocks, and the lower plot displays the standard deviations across blocks.

All models perform substantially worse than the holdout-evaluation on at least some of the perturbations. In addition, the standard deviation of accuracy between blocks is higher than the holdout tests. As discussed in §7.1, low accuracy on this probe indicates that closed-class words do not maintain a consistent interpretation when paired with different open-class words. Variance across blocks shows that under all models the behavior of closed-class words is highly sensitive to the novel words they appear with.

Performance is also susceptible to interference from sentence-level features. For example, consider the perturbation which deletes a post-modifier from a sentence pair in negation, yielding a pair in cover relation. The self-attentive encoder performs perfectly when this perturbation is applied to a premise ($100\% \pm 0.00\%$), but not when applied to a hypothesis ($86.60\% \pm 18.08\%$). Similarly, deleting the adjective *red* from the hypothesis of a forward-entailing pair results in an unrelated sen-

tence pair ($84.79\% \pm 7.50\%$) or another forward-entailing pair ($92.32\%, \pm 3.60\%$) or an equality pair ($100\% \pm 0.00\%$). All the possible perturbations we studied exhibit similarly inconsistent performance.

8.4 Analysis IV: Identical Open-Class Words Probe

Recall that the identical open-class words probe consist of sentence pairs where all open-class lexical items were identical. Table 4 shows the accuracies for these probes, trained on the small language. Average accuracies across jabberwocky blocks are reported together with standard deviations.

Relation	BGRU mean (sd)	CONV mean (sd)	SATT mean (sd)	INFS mean (sd)
#	100 ± 0	100 ± 0	99.94 ± 0.26	99.67 ± 0.98
^	55.68 ± 20.29	73.29 ± 10.8	23.71 ± 11.45	90.67 ± 10.98
⊂	90.78 ± 4.99	82.84 ± 6.51	75.22 ± 5.98	95.53 ± 2.64
≡	90.43 ± 17.1	38.12 ± 15.56	71.94 ± 24.1	95.93 ± 6.5
⊃	90.34 ± 4.18	77.11 ± 5.9	81.4 ± 6.67	93.81 ± 2.96
	93.08 ± 3.58	85.34 ± 5.47	74.05 ± 8.03	92.23 ± 4.6
⌋	88.01 ± 3.55	71.5 ± 7.32	78.4 ± 7.91	95.22 ± 3.58

Table 4: Identical open-class words probe performance, trained on the small language condition (trained on 51, 743 sentence pairs, see §5.5)

Accuracy on the probe pairs fails to reach the holdout test levels for most models and most relations besides #, and variance between blocks is much higher than in the holdout evaluation. Of special interest is negation (^), for which accuracy is dramatically lower and variance dramatically higher than the holdout evaluation.

The results are similar for the large language condition, shown in Table 5. Although model accuracies improve somewhat, variance remains higher than the heldout level and accuracy lower. Recall that these probe-items can be classified while ignoring the specific identity of their open-class words. Thus, the models inability to leverage this fact, and high variance across different sets novel open-class words, illustrates their sensitivity to context.

8.5 Analysis V: Consistency Probe

The consistency probe tests abstract knowledge of relationships between logical relations, such as the fact that two sentences that stand in a contradiction still stand in a contradiction after reversing their order. Results of this probe in the small-language condition are in Table 6: For each type of relation, we show the average percentage of correctly-labeled

Relation	BGRU mean (sd)	CONV mean (sd)	SATT mean (sd)	INFS mean (sd)
#	99.82 ±0.45	99.57 ±0.73	98.67 ±1.81	100 ±0
^	84.18 ±12.29	73.73 ±18.31	79.97 ±16.58	85.54 ±14.11
⊂	96.13 ±2.59	93.88 ±2.67	97.3 ±2.36	97.02 ±2.39
≡	89.33 ±12.5	77.84 ±12.08	94.44 ±11.23	94.59 ±7.02
⊃	95.4 ±2.48	94.55 ±2.04	98.05 ±1.51	97.6 ±2.08
	89.97 ±6.73	92.36 ±6.29	84.52 ±7.07	98.72 ±2.08
⊆	90.78 ±6.33	93.18 ±2.95	87.85 ±6.46	97.48 ±2.56

Table 5: Identical open-class words probe performance when trained on the large language training condition (trained on 478, 649 sentence pairs, see §5.5)

sentence pairs that, when presented in reverse order, were also correctly labeled.

The best-performing model on negation reversal is SATT, which correctly labeled reversed items 66.92% of the time. Although performance on negation is notably more difficult than the other relations, every model, on every relation, exhibited inter-block variance higher than that of the hold-out evaluations.

Relation	BGRU mean (sd)	CONV mean (sd)	SATT mean (sd)	INFS mean (sd)
#	97.4 ±0.86	97.8 ±0.93	98.58 ±0.74	97.03 ±0.87
^	63.03 ±36.19	63.42 ±35.91	66.92 ±31.45	57.16 ±38.24
⊂	92.45 ±6.26	88.1 ±8.16	93.16 ±5.42	90.64 ±6.76
≡	100 ±0	100 ±0	100 ±0	100 ±0
⊃	91.37 ±6.23	94.73 ±6.51	96.42 ±3.22	87.02 ±9.61
	96.02 ±2.6	96.29 ±2.51	96.95 ±2.14	94.2 ±3.48
⊆	93.57 ±3.56	95 ±2.97	96.4 ±2.83	93.1 ±3.77

Table 6: Consistency probe performance, trained on the small language condition (51, 743 sentence pairs, see §5.5).

Furthermore, as can be seen in Table 7, the large language condition yields little improvement. Negation pairs are still well below the hold-out test threshold, still with a high degree of variation. Variation remains high for many relations, which is surprising because the means report accuracy on test items that were chosen specifically because the same item, in a reverse order, was already correctly labeled. Reversing the order of sentences causes the model to misclassify the resulting pair, more often for some blocks than others.

9 Discussion and Conclusion

Systematicity refers to the property of natural language representations whereby words (and other units or grammatical operations) have consistent meanings across different contexts. Our probes test whether deep learning systems learn to represent linguistic units systematically in the natural lan-

Relation	BGRU mean (sd)	CONV mean (sd)	SATT mean (sd)	INFS mean (sd)
#	98.45 ±0.65	98.69 ±0.54	98.83 ±0.6	98.38 ±0.74
^	70.46 ±33.72	77.82 ±26	84.27 ±23.89	65.64 ±35.13
⊂	96.02 ±2.96	96.6 ±3.26	96.78 ±4.23	95.01 ±5.38
=	100 ±0	100 ±0	100 ±0	100 ±0
⊃	93.5 ±4.51	95.76 ±4.23	94.23 ±5.86	90.11 ±8.5
	96.31 ±2.73	97.25 ±2.05	97.17 ±2.23	94.46 ±4.24
⊆	96.25 ±2.49	96.98 ±2.66	97.18 ±2.17	93.88 ±4.78

Table 7: Consistency probe performance, trained on the large language condition (478, 649 sentence pairs).

guage inference task. Our results indicate that despite their high overall performance, these models tend to generalize in ways that allow the meanings of individual words to vary in different contexts, even in an artificial language where a totally systematic solution is available. This suggests the networks lack a sufficient inductive bias to learn systematic representations of words like quantifiers, which even in natural language exhibit very little meaning variation.

Our analyses contain two ideas that may be useful for future studies of systematicity. First, two of our probes (known word perturbation and consistency) are based on the idea of starting from a test item that is classified correctly, and applying a transformation that should result in a classifiable item (for a model that represents word meaning systematically). Second, our analyses made critical use of differential sensitivity (i.e., variance) of the models across test blocks with different novel words but otherwise identical information content. We believe these are a novel ideas that can be employed in future studies.

Acknowledgements

We thank Brendan Lake, Marco Baroni, Adina Williams, Dima Bahdanau, Sam Gershman, Ishita Dasgupta, Alessandro Sordani, Will Hamilton, Leon Bergen, the Montreal Computational and Quantitative Linguistics, and Reasoning and Learning Labs at McGill University for feedback on the manuscript. We are grateful to Facebook AI Research for providing extensive compute and other support. We also gratefully acknowledge the support of the Natural Sciences and Engineering Research Council of Canada, the Fonds de Recherche du Québec, Société et Culture, and the Canada CIFAR AI Chairs Program.

References

- Adam Albright and Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition*, 90(2):119–161.
- Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron Courville. 2019. Systematic generalization: What is required and can it be learned? *arXiv*, abs/1811.12889(arXiv:1811.12889v2 [cs.CL]).
- Luisa Bentivogli, Raffaella Bernardi, Marco Marelli, Stefano Menini, Marco Baroni, and Roberto Zamparelli. 2016. Sick through the semeval glasses. lesson learned from the evaluation of compositional distributional semantic models on full sentences through semantic relatedness and textual entailment. *Language Resources and Evaluation*, 50(1):95–124.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Samuel R Bowman, Christopher Potts, and Christopher D Manning. 2014. Recursive neural networks can learn logical semantics. *arXiv preprint arXiv:1406.1827*.
- Rudolf Camap. 1947. Meaning and necessity: A study in semantics and modal logic.
- Brian Cantwell Smith. 1996. *On the Origins of Objects*. MIT Press, Cambridge, Massachusetts.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680.
- Ishita Dasgupta, Demi Guo, Samuel J. Gershman, and Noah D. Goodman. 2019. [Analyzing machine-learned representations: A natural language case study](#).
- Ishita Dasgupta, Demi Guo, Andreas Stuhlmüller, Samuel J Gershman, and Noah D Goodman. 2018. Evaluating compositionality in sentence embeddings. *arXiv preprint arXiv:1802.04302*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Catarina Dutilh Novaes. 2012. *Formal Languages in Logic: A Philosophical and Cognitive Analysis*. Cambridge University Press, Cambridge, England.
- Allyson Ettinger, Ahmed Elgohary, Colin Phillips, and Philip Resnik. 2018. Assessing composition in sentence vector representations. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1790–1801, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Richard Evans, David Saxton, David Amos, Pushmeet Kohli, and Edward Grefenstette. 2018. Can neural networks understand logical entailment? *arXiv preprint arXiv:1802.08535*.
- Jerry A Fodor and Zenon W Pylyshyn. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1):3–71.
- Samuel Gershman and Joshua B Tenenbaum. 2015. Phrase similarity in humans and machines. In *CogSci*.
- Mario Giulianelli, Jack Harding, Florian Mohnert, Dieuwke Hupkes, and Willem Zuidema. 2018. Under the hood: Using diagnostic classifiers to investigate and improve how language models track agreement information. *arXiv preprint arXiv:1808.08079*.
- Irene Heim and Angelika Kratzer. 2000. *Semantics in Generative Grammar*. Blackwell Publishing, Malden, MA.
- Mikael Henaff, Jason Weston, Arthur Szlam, Antoine Bordes, and Yann LeCun. 2016. [Tracking the world state with recurrent entity networks](#).
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Wilfrid Hodges. 2012. Formalizing the relationship between meaning and syntax. In *The Oxford handbook of compositionality*. Oxford Handbooks in Linguistic.
- Drew A. Hudson and Christopher D. Manning. 2018. [Compositional attention networks for machine reasoning](#).
- Theo MV Janssen et al. 2012. Compositionality: its historic context. *The Oxford handbook of compositionality*, pages 19–46.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910.

- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2019. Measuring compositional generalization: A comprehensive method on realistic data.
- Douwe Kiela, Changan Wang, and Kyunghyun Cho. 2018. Dynamic meta-embeddings for improved sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1466–1477.
- Najoung Kim, Roma Patel, Adam Poliak, Patrick Xia, Alex Wang, Tom McCoy, Ian Tenney, Alexis Ross, Tal Linzen, Benjamin Van Durme, Samuel R. Bowman, and Ellie Pavlick. 2019. Probing what different NLP tasks teach machines about function word comprehension. In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, pages 235–249, Minneapolis, Minnesota. Association for Computational Linguistics.
- Brenden Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International Conference on Machine Learning*, pages 2879–2888.
- Brenden Lake, Tal Linzen, and Marco Baroni. 2017a. Human few-shot learning of compositional instructions. In Ashok Goel, Colleen Seifert, and Christian Freksa, editors, *Proceedings of the 41st Annual Conference of the Cognitive Science Society*, pages 611–616. Cognitive Science Society, Montreal, Canada.
- Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. 2017b. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.
- George Lakoff. 1970. Linguistics and natural logic. *Synthese*, 22(1-2):151–271.
- Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- Bill MacCartney and Christopher D. Manning. 2009. An extended model of natural logic. In *Proceedings of the Eighth International Conference on Computational Semantics, IWCS-8 '09*, pages 140–156, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Bill MacCartney and Christopher D. Manning. 2014. Natural logic and natural language inference. In *Computing Meaning: Volume 4*, pages 129–147, Dordrecht. Springer Netherlands.
- R Thomas McCoy, Junghyun Min, and Tal Linzen. 2019a. Berts of a feather do not generalize together: Large variability in generalization across models with similar test set performance. *arXiv preprint arXiv:1911.02969*.
- R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019b. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy.
- Tsendsuren Munkhdalai and Hong Yu. 2017. Neural semantic encoders. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 1, page 397. NIH Public Access.
- Timothy Niven and Hung-Yu Kao. 2019. Probing neural network comprehension of natural language arguments.
- Ankur P Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. *arXiv preprint arXiv:1606.01933*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- Grusha Prasad, Marten van Schijndel, and Tal Linzen. 2019. Using priming to uncover the organization of syntactic representations in neural language models.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/languageunsupervised/language_understanding_paper.pdf.
- Kyle Richardson, Hai Hu, Lawrence S. Moss, and Ashish Sabharwal. 2019. Probing natural language inference models through semantic fragments.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. 2015. Reasoning about entailment with neural attention. *arXiv preprint arXiv:1509.06664*.
- Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. Clutrr: A diagnostic benchmark for inductive reasoning from text.
- Paul Soulos, Tom McCoy, Tal Linzen, and Paul Smolensky. 2019. Discovering the compositional structure of vector representations with role learning networks.

- Shane Storks, Qiaozi Gao, and Joyce Y Chai. 2019. Recent advances in natural language inference: A survey of benchmarks, resources, and approaches. *arXiv preprint arXiv:1904.01172*.
- Zoltan Szabó. 2012. The case for compositionality. *The Oxford Handbook of Compositionality*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).
- Sara Veldhoen and Willem Zuidema. 2017. Can neural networks learn logical reasoning? *Proceedings of the Conference on Logic and Machine Learning in Natural Language*.
- Wei Wang, Ming Yan, and Chen Wu. 2018. Multi-granularity hierarchical attention fusion networks for reading comprehension and question answering. *arXiv preprint arXiv:1811.11934*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#).
- Deunsol Yoon, Dongbok Lee, and SangKeun Lee. 2018. Dynamic self-attention: Computing attention over words dynamically for sentence embedding. *arXiv preprint arXiv:1808.07383*.
- Wlodek Zadrozny. 1994. From compositional to systematic semantics. *Linguistics and philosophy*, 17(4):329–342.
- Zhuosheng Zhang, Yuwei Wu, Zuchao Li, Shexia He, Hai Zhao, Xi Zhou, and Xiang Zhou. 2018. I know what you want: Semantic learning for text comprehension. *arXiv preprint arXiv:1809.02794*.
- Han Zhao, Zhengdong Lu, and Pascal Poupart. 2015. Self-adaptive hierarchical sentence model. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.