

# Do NLP Models Know Numbers? Probing Numeracy in Embeddings

Eric Wallace<sup>\*1</sup>, Yizhong Wang<sup>\*2</sup>, Sujian Li<sup>2</sup>, Sameer Singh<sup>3</sup>, Matt Gardner<sup>1</sup>

<sup>1</sup>Allen Institute for Artificial Intelligence

<sup>2</sup>Peking University

<sup>3</sup>University of California, Irvine

{ericw, mattg}@allenai.org, {yizhong, lisujian}@pku.edu.cn, sameer@uci.edu

## Abstract

The ability to understand and work with numbers (numeracy) is critical for many complex reasoning tasks. Currently, most NLP models treat numbers in text in the same way as other tokens—they embed them as distributed vectors. Is this enough to capture numeracy? We begin by investigating the numerical reasoning capabilities of a state-of-the-art question answering model on the DROP dataset. We find this model excels on questions that require numerical reasoning, i.e., it already captures numeracy. To understand how this capability emerges, we probe token embedding methods (e.g., BERT, GloVe) on synthetic list maximum, number decoding, and addition tasks. A surprising degree of numeracy is naturally present in standard embeddings. For example, GloVe and word2vec accurately encode magnitude for numbers up to 1,000. Furthermore, character-level embeddings are even more precise—ELMo captures numeracy the best for all pre-trained methods—but BERT, which uses sub-word units, is less exact.

## 1 Introduction

Neural NLP models have become the de-facto standard tool across language understanding tasks, even solving basic reading comprehension and textual entailment datasets (Yu et al., 2018; Devlin et al., 2019). Despite this, existing models are incapable of complex forms of reasoning, in particular, we focus on the ability to reason *numerically*. Recent datasets such as DROP (Dua et al., 2019), EQUATE (Ravichander et al., 2019), or Mathematics Questions (Saxton et al., 2019) test numerical reasoning; they contain examples which require comparing, sorting, and adding numbers in natural language (e.g., Figure 2).

The first step in performing numerical reasoning over natural language is *numeracy*: the abil-

<sup>\*</sup>Equal contribution; work done while interning at AI2.

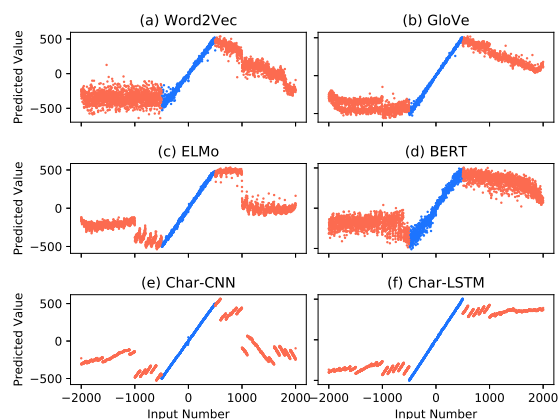


Figure 1: We train a probing model to decode a number from its word embedding over a random 80% of the integers from  $[-500, 500]$ , e.g., “71”  $\rightarrow$  71.0. We plot the model’s predictions for all numbers from  $[-2000, 2000]$ . The model accurately decodes numbers within the training range (in blue), i.e., pre-trained embeddings like GloVe and BERT capture numeracy. However, the probe fails to extrapolate to larger numbers (in red). The Char-CNN (e) and Char-LSTM (f) are trained jointly with the probing model.

ity to understand and work with numbers in either digit or word form (Spithourakis and Riedel, 2018). For example, one must understand that the string “23” represents a bigger value than “twenty-two”. Once a number’s value is (perhaps implicitly) represented, reasoning algorithms can then process the text, e.g., extracting the list of field goals and computing that list’s maximum (first question in Figure 2). Learning to reason numerically over paragraphs with only question-answer supervision appears daunting for end-to-end models; our work seeks to understand *if* and *how* “out-of-the-box” neural NLP models already learn this.

We begin by analyzing the state-of-the-art NAQANet model (Dua et al., 2019) for DROP—testing it on a subset of questions that evaluate numerical reasoning (Section 2). To our surprise,

the model exhibits excellent numerical reasoning abilities. Amidst reading and comprehending natural language, the model successfully computes list maximums/minimums, extracts superlative entities (argmax reasoning), and compares numerical quantities. For instance, despite NAQANet achieving only 49 F1 on the entire validation set, it scores 89 F1 on numerical comparison questions. We also stress test the model by perturbing the validation paragraphs and find one failure mode: the model struggles to extrapolate to numbers outside its training range.

We are especially intrigued by the model’s ability to learn numeracy, i.e., how does the model know the value of a number given its embedding? The model uses standard embeddings (GloVe and a Char-CNN) and receives no direct supervision for number magnitude/ordering. To understand how numeracy emerges, we probe token embedding methods (e.g., BERT, GloVe) using synthetic list maximum, number decoding, and addition tasks (Section 3).

We find that all widely-used pre-trained embeddings, e.g., ELMo (Peters et al., 2018), BERT (Devlin et al., 2019), and GloVe (Pennington et al., 2014), capture numeracy: number magnitude is present in the embeddings, even for numbers in the thousands. Among all embeddings, character-level methods exhibit stronger numeracy than word- and sub-word-level methods (e.g., ELMo excels while BERT struggles), and character-level models learned directly on the synthetic tasks are the strongest overall. Finally, we investigate why NAQANet had trouble extrapolating—was it a failure in the model or the embeddings? We repeat our probing tasks and test for model extrapolation, finding that neural models struggle to predict numbers outside the training range.

## 2 Numeracy Case Study: DROP QA

This section examines the state-of-the-art model for DROP by investigating its accuracy on questions that require numerical reasoning.

### 2.1 DROP Dataset

DROP is a reading comprehension dataset that tests numerical reasoning operations such as counting, sorting, and addition (Dua et al., 2019). The dataset’s input-output format is a superset of SQuAD (Rajpurkar et al., 2016): the answers are paragraph spans, as well as question

... JaMarcus Russell completed a 91-yard touchdown pass to rookie wide receiver **Chaz Schilens**. The Texans would respond with fullback Vonta Leach getting a 1-yard touchdown run, yet the Raiders would answer with kicker Sebastian Janikowski getting a 33-yard and a 21-yard field goal. Houston would tie the game in the second quarter with kicker Kris Brown getting a 53-yard and a 24-yard field goal. Oakland would take the lead in the third quarter with wide receiver Johnnie Lee Higgins catching a **29-yard** touchdown pass from Russell, followed up by an 80-yard punt return for a touchdown.

Q: How many yards was the longest field goal? A: 53

Q: How long was the shortest touchdown pass? A: 29-yard

Q: Who caught the longest touchdown? A: Chaz Schilens

Figure 2: Three DROP questions that require numerical reasoning; the state-of-the-art NAQANet answers every question correct. Plausible answer candidates to the questions are underlined and the model’s predictions are shown in bold.

spans, number answers (e.g., 35), and dates (e.g., 03/01/2014). The only supervision provided is the question-answer pairs, i.e., a model must learn to reason numerically while simultaneously learning to read and comprehend.

### 2.2 NAQANet Model

Modeling approaches for DROP include both semantic parsing (Krishnamurthy et al., 2017) and reading comprehension (Yu et al., 2018) models. We focus on the latter, specifically on Numerically-augmented QANet (NAQANet), the current state-of-the-art model (Dua et al., 2019).<sup>1</sup> The model’s core structure closely follows QANet (Yu et al., 2018) except that it contains four output branches, one for each of the four answer types (passage span, question span, count answer, or addition/subtraction of numbers.)

Words and numbers are represented as the concatenation of GloVe embeddings and the output of a character-level CNN. The model contains no auxiliary components for representing number magnitude or performing explicit comparisons. We refer readers to Yu et al. (2018) and Dua et al. (2019) for further details.

### 2.3 Comparative and Superlative Questions

We focus on questions that NAQANet requires numeracy to answer, namely *Comparative* and *Superlative* questions.<sup>2</sup> Comparative questions

<sup>1</sup>Result as of May 21st, 2019.

<sup>2</sup>DROP addition, subtraction, and count questions do not require numeracy for NAQANet, see Appendix A.

Question Type	Example	Reasoning Required
Comparative (Binary)	Which country is a bigger exporter, Brazil or Uruguay?	Binary Comparison
Comparative (Non-binary)	Which player had a touchdown longer than 20 yards?	Greater Than
Superlative (Number)	How many yards was the shortest field goal?	List Minimum
Superlative (Span)	Who kicked the longest field goal?	Argmax

Table 1: We focus on DROP *Comparative* and *Superlative* questions which test NAQANet’s numeracy.

probe a model’s understanding of quantities or events that are “larger”, “smaller”, or “longer” than others. Certain comparative questions ask about “either-or” relations (e.g., first row of Table 1), which test binary comparison. Other comparative questions require more diverse comparative reasoning, such as greater than relationships (e.g., second row of Table 1).

Superlative questions ask about the “shortest”, “largest”, or “biggest” quantity in a passage. When the answer type is a *number*, superlative questions require finding the maximum or minimum of a list (e.g., third row of Table 1). When the answer type is a *span*, superlative questions usually require an argmax operation, i.e., one must find the superlative action or quantity and then extract the associated entity (e.g., fourth row of Table 1). We filter the validation set to comparative and superlative questions by writing templates to match words in the question.

## 2.4 Emergent Numeracy in NAQANet

NAQANet’s accuracy on comparative and superlative questions is significantly higher than its average accuracy on the validation set (Table 2).<sup>3</sup>

NAQANet achieves 89.0 F1 on binary (either-or) comparative questions, approximately 40 F1 points higher than the average validation question and within 7 F1 points of human test performance. The model achieves a lower, but respectable, accuracy on non-binary comparisons. These questions require multiple reasoning steps, e.g., the second question in Table 1 requires (1) extracting all the touchdown distances, (2) finding the distance that is greater than twenty, and (3) selecting the player associated with the touchdown of that distance.

We divide the superlative questions into questions that have number answers and questions with span answers according to the dataset’s provided answer type. NAQANet achieves nearly 70 F1 on superlative questions with number answers, i.e., it can compute list maximum and minimums.

<sup>3</sup>We have a public NAQANet demo available <https://demo.allennlp.org/reading-comprehension>.

Question Type	Count	EM	F1
Human (Test Set)	9622	92.4	96.0
Full Validation	9536	46.2	49.2
Number Answers	5842	44.3	44.4
Comparative	704	73.6	76.4
Binary (either-or)	477	86.0	89.0
Non-binary	227	47.6	49.8
Superlative Questions	861	64.6	67.7
Number Answers	475	68.8	69.2
Span Answers	380	59.7	66.3

Table 2: NAQANet achieves higher accuracy on questions that require numerical reasoning (*Superlative* and *Comparative*) than on standard validation questions. Human performance is reported from Dua et al. (2019).

The model answers about two-thirds of superlative questions with span answers correctly (66.3 F1), i.e., it can perform argmax reasoning.

Figure 2 shows examples of superlative questions answered correctly by NAQANet. The first two questions require computing the maximum/minimum of a list: the model must recognize which digits correspond to field goals and touchdowns passes, and then extract the maximum/minimum of the correct list. The third question requires argmax reasoning: the model must first compute the longest touchdown pass and then find the corresponding receiver “Chaz Schilens”.

## 2.5 Stress Testing NAQANet’s Numeracy

Just how far does the numeracy of NAQANet go? Here, we stress test the model by automatically modifying DROP validation paragraphs.

We test two phenomena: larger numbers and word-form numbers. For larger numbers, we generate a random positive integer and multiply or add that value to the numbers in each paragraph. For word forms, we replace every digit in the paragraph with its word form (e.g., “75” → “seventy-five”). Since word-form numbers are usually small in magnitude when they occur in DROP, we perform word replacements for integers in the range [0, 100]. We guarantee the ground-truth answer is

Stress Test Dataset	All Questions		Superlative	
	F1	$\Delta$	F1	$\Delta$
Original Validation Set	49.2	-	67.7	-
Add [1, 20]	47.7	-1.5	64.1	-3.6
Add [21, 100]	41.4	-7.8	40.4	-27.3
Multiply [2, 10]	41.1	-8.1	39.3	-28.4
Multiply [11, 100]	38.8	-10.4	32.0	-35.7
Digits to Words [0, 20]	45.5	-3.7	63.8	-3.9
Digits to Words [21, 100]	41.9	-7.3	46.1	-21.6

Table 3: We stress test NAQANet’s numeracy by manipulating the numbers in the validation paragraphs. *Add* or *Multiply*  $[x, y]$  indicates adding or multiplying all of the numbers in the passage by a random integer in the range  $[x, y]$ . *Digits*  $\rightarrow$  *Words*  $[x, y]$  converts all integers in the passage within the range  $[x, y]$  to their corresponding word form (e.g., “75”  $\rightarrow$  “seventy-five”).

still valid by only modifying NAQANet’s internal representation (Appendix E).

Table 3 shows the results for different paragraph modifications. The model exhibits a tiny degradation in performance for small magnitude changes (e.g., NAQANet drops 1.5 F1 overall for Add [1,20]) but severely struggles on larger changes (e.g., NAQANet drops 35.7 F1 on superlative questions for Multiply [11,200]). Similar trends hold for word forms: the model exhibits small drops in accuracy when converting small numbers to words (3.9 degradation on Digits to Words [0,20]) but fails on larger magnitude word forms (21.6 F1 drop over [21,100]). These results show that NAQANet has a strong understanding of numeracy for numbers in the training range, but, the model can fail to extrapolate to other values.

## 2.6 Whence this behavior?

NAQANet exhibits numerical reasoning capabilities that exceed our expectations. What enables this behavior? Aside from reading and comprehending the passage/question, this kind of numerical reasoning requires two components: numeracy (i.e., representing numbers) and comparison algorithms (i.e., computing the maximum of a list).

Although the natural emergence of comparison algorithms is surprising, previous results show neural models are capable of learning to count and sort synthetic lists of scalar values when given explicit supervision (Weiss et al., 2018; Vinyals et al., 2016). NAQANet demonstrates that a model can learn comparison algorithms while simultane-

ously learning to read and comprehend, even with only question-answer supervision.

How, then, does NAQANet know numeracy? The source of numerical information eventually lies in the token embeddings themselves, i.e., the character-level convolutions and GloVe embeddings of the NAQANet model. Therefore, we can understand the source of numeracy by isolating and probing these embeddings.

## 3 Probing Numeracy of Embeddings

We use synthetic numerical tasks to probe the numeracy of token embeddings.

### 3.1 Probing Tasks

We consider three synthetic tasks to evaluate numeracy (Figure 3). Appendix C provides further details on training and evaluation.

**List Maximum** Given a list of the embeddings for five numbers, the task is to predict the index of the maximum number. Each list consists of values of similar magnitude in order to evaluate fine-grained comparisons (see Appendix C). As in typical span selection models (Seo et al., 2017), an LSTM reads the list of token embeddings, and a weight matrix and softmax function assign a probability to each index using the model’s hidden state. We use the negative log-likelihood of the maximum number as the loss function.

**Decoding** The decoding task probes whether number magnitude is captured (rather than the relative ordering of numbers as in list maximum). Given a number’s embedding, the task is to regress to its value, e.g., the embedding for the string “five” has a target of 5.0. We consider a linear regression model and a three-layer fully-connected network with ReLU activations. The models are trained using a mean squared error (MSE) loss.

**Addition** The addition task requires number manipulation—given the embeddings of two numbers, the task is to predict their sum. Our model concatenates the two token embeddings and feeds the result through a three-layer fully-connected network with ReLU activations, trained using MSE loss. Unlike the decoding task, the model needs to capture number magnitude internally without direct label supervision.

**Training and Evaluation** We focus on a numerical *interpolation* setting (we revisit extrapolation

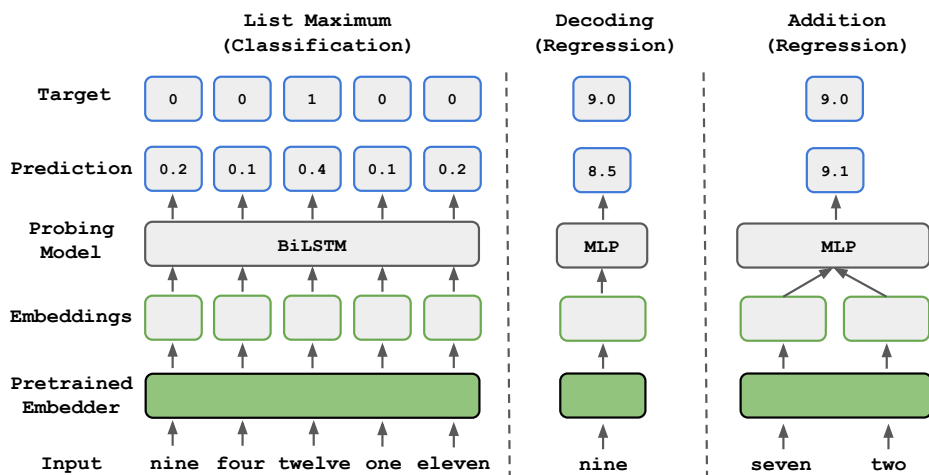


Figure 3: Our probing setup. We pass numbers through a pre-trained embedder (e.g., BERT, GloVe) and train a probing model to solve numerical tasks such as finding a list’s maximum, decoding a number, or adding two numbers. If the probing model generalizes to held-out numbers, the pre-trained embeddings must contain numerical information. We provide numbers as either words (shown here), digits (“9”), floats (“9.1”), or negatives (“-9”).

in Section 3.4): the model is tested on values that are within the training range. We first pick a range (we vary the range in our experiments) and randomly shuffle the integers over it. We then split 80% of the numbers into a training set and 20% into a test set. We report the mean and standard deviation across five different random shuffles for a particular range, using the exact same shuffles across all embedding methods.

Numbers are provided as integers (“75”), single-word form (“seventy-five”), floats (“75.1”), or negatives (“-75”). We consider positive numbers less than 100 for word-form numbers to avoid multiple tokens. We report the classification accuracy for the list maximum task (5 classes), and the Root Mean Squared Error (RMSE) for decoding and addition. Note that larger ranges will naturally amplify the RMSE error.

### 3.2 Embedding Methods

We evaluate various token embedding methods.

**Word Vectors** We use 300-dimensional GloVe (Pennington et al., 2014) and word2vec vectors (Mikolov et al., 2018). We ensure all values are in-vocabulary for word vectors.

**Contextualized Embeddings** We use ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019) embeddings.<sup>4</sup> ELMo uses character-level convo-

<sup>4</sup>Since our inputs are numbers, not natural sentences, language models may exhibit strange behavior. We experimented with extracting the context-independent feature vector immediately following the character convolutions for ELMo but found little difference in results.

lutions of size 1–7 with max pooling. BERT represents tokens via sub-word pieces; we use lower-cased BERT-base with 30k pieces.

**NAQANet Embeddings** We extract the GloVe embeddings and Char-CNN from the NAQANet model trained on DROP. We also consider an ablation that removes the GloVe embeddings.

**Learned Embeddings** We use a character-level CNN (Char-CNN) and a character-Level LSTM (Char-LSTM). We use left character padding, which greatly improves numeracy for character-level CNNs (details in Appendix B).

**Untrained Embeddings** We consider two untrained baselines. The first baseline is random token vectors, which trivially fail to generalize (there is no pattern between train and test numbers). These embeddings are useful for measuring the improvement of pre-trained embeddings. We also consider a randomly initialized and untrained Char-CNN and Char-LSTM.

**Number’s Value as Embedding** The final embedding method is simple: map a number’s embedding directly to its value (e.g., “seventy-five” embeds to [75]). We found this strategy performs poorly for large ranges; using a base-10 logarithmic scale improves performance. We report this as *Value Embedding* in our results.<sup>5</sup>

<sup>5</sup>We suspect the failures result from the raw values being too high in magnitude and/or variance for the model. We also experimented with normalizing the values to mean 0 and variance 1; a logarithmic scale performed better.

All pre-trained embeddings (all methods except the Char-CNN and Char-LSTM) are fixed during training. The probing models are trained on the synthetic tasks on top of these embeddings.

### 3.3 Results: Embeddings Capture Numeracy

We find that all pre-trained embeddings *contain* fine-grained information about number magnitude and order. We first focus on integers (Table 4).

**Word Vectors Succeed** Both word2vec and GloVe significantly outperform the random vector baseline and are among the strongest methods overall. This is particularly surprising given the training methodology for these embeddings, e.g., a continuous bag of words objective can teach fine-grained number magnitude.

**Character-level Methods Dominate** Models which use character-level information have a clear advantage over word-level models for encoding numbers. This is reflected in our probing results: character-level CNNs are the best architecture for capturing numeracy. For example, the NAQANet model without GloVe (only using its Char-CNN) and ELMo (uses a Char-CNN) are the strongest pre-trained methods, and a learned Char-CNN is the strongest method overall. The strength of the character-level convolutions seems to lie in the architectural prior—an *untrained* Char-CNN is surprisingly competitive. Similar results have been shown for images (Saxe et al., 2011): random CNNs are powerful feature extractors.

**Sub-word Models Struggle** BERT struggles for large ranges (e.g., 52% accuracy for list maximum for [0,9999]). We suspect this results from sub-word pieces being a poor method to encode digits: two numbers which are similar in value can have very different sub-word divisions.

**A Linear Subspace Exists** For small ranges on the decoding task (e.g., [0,99]), a linear model is competitive, i.e., a linear subspace captures number magnitude (Appendix D). For larger ranges (e.g., [0,999]), the linear model’s performance degrades, especially for BERT.

**Value Embedding Fails** The Value Embedding method fails for large ranges. This is surprising as the embedding directly provides a number’s value, thus, the synthetic tasks *should* be easy to solve. However, we had difficulty training models for

large ranges, even when using numerous architecture variants (e.g., tiny networks with 10 hidden units and tanh activations) and hyperparameters. Trask et al. (2018) discuss similar problems and ameliorate them using new neural architectures.

### Words, Floats, and Negatives are Captured

Finally, we probe the embeddings on word-form numbers, floats, and negatives. We observe similar trends for these inputs as integers: pre-trained models exhibit natural numeracy and learned embeddings are strong (Tables 5, 6, and 10). The ordering of the different embedding methods according to performance is also relatively consistent across the different input types. One notable exception is that BERT struggles on floats, which is likely a result of its sub-word pieces. We do not test word2vec and GloVe on floats/negatives because they are out-of-vocabulary.

### 3.4 Probing Models Struggle to Extrapolate

Thus far, our synthetic experiments evaluate on held-out values *within* the same range as the training data (i.e., numerical interpolation). In Section 2.5, we found that NAQANet struggles to extrapolate to values *outside* the training range. Is this an idiosyncrasy of NAQANet or is it a more general problem? We investigate this using a numerical extrapolation setting: we train models on a specific integer range and test them on values greater than the largest training number and smaller than the smallest training number.

**Extrapolation for Decoding and Addition** For decoding and addition, models struggle to extrapolate. Figure 1 shows the predictions for models trained on 80% of the values from [-500,500] and tested on held-out numbers in the range [-2000, 2000] for six embedding types. The embedding methods fail to extrapolate in different ways, e.g., predictions using word2vec decrease almost monotonically as the input increases, while predictions using BERT are usually near the highest training value. Trask et al. (2018) also observe that models struggle outside the training range; they attribute this to failures in neural models themselves.

**Extrapolation for List Maximum** For the list maximum task, accuracies are closer to those in the interpolation setting, however, they still fall short. Table 7 shows the accuracy for models trained on the integer range [0,150] and tested on

Interpolation <i>Integer Range</i>	List Maximum (5-classes)			Decoding (RMSE)			Addition (RMSE)		
	[0,99]	[0,999]	[0,9999]	[0,99]	[0,999]	[0,9999]	[0,99]	[0,999]	[0,9999]
Random Vectors	0.16	0.23	0.21	29.86	292.88	2882.62	42.03	410.33	4389.39
Untrained CNN	0.97	0.87	0.84	2.64	9.67	44.40	1.41	14.43	69.14
Untrained LSTM	0.70	0.66	0.55	7.61	46.5	210.34	5.11	45.69	510.19
Value Embedding	<b>0.99</b>	0.88	0.68	<b>1.20</b>	11.23	275.50	<b>0.30</b>	15.98	654.33
<i>Pre-trained</i>									
Word2Vec	0.90	0.78	0.71	2.34	18.77	333.47	0.75	21.23	210.07
GloVe	0.90	0.78	0.72	2.23	13.77	174.21	0.80	16.51	180.31
ELMo	0.98	0.88	0.76	2.35	13.48	62.20	0.94	15.50	45.71
BERT	0.95	0.62	0.52	3.21	29.00	431.78	4.56	67.81	454.78
<i>Learned</i>									
Char-CNN	0.97	<b>0.93</b>	<b>0.88</b>	2.50	<b>4.92</b>	<b>11.57</b>	1.19	<b>7.75</b>	<b>15.09</b>
Char-LSTM	0.98	0.92	0.76	2.55	8.65	18.33	1.21	15.11	25.37
<i>DROP-trained</i>									
NAQANet	0.91	0.81	0.72	2.99	14.19	62.17	1.11	11.33	90.01
- GloVe	0.88	0.90	0.82	2.87	5.34	35.39	1.45	9.91	60.70

Table 4: *Interpolation with integers (e.g., “18”).* All pre-trained embedding methods (e.g., GloVe and ELMo) surprisingly capture numeracy. The probing model is trained on a randomly shuffled 80% of the *Integer Range* and tested on the remaining 20%. The probing model architecture and train/test splits are equivalent across all embeddings. We show the mean over 5 random shuffles (standard deviation in Appendix D).

Interpolation <i>Float Range</i>	List Maximum (5-classes)	
	[0.0,99.9]	[0.0,999.9]
Rand. Vectors	0.18 ± 0.03	0.21 ± 0.04
ELMo	0.91 ± 0.03	0.59 ± 0.01
BERT	0.82 ± 0.05	0.51 ± 0.04
Char-CNN	0.87 ± 0.04	0.75 ± 0.03
Char-LSTM	0.81 ± 0.05	0.69 ± 0.02

Table 5: *Interpolation with floats (e.g., “18.1”) for list maximum.* Pre-trained embeddings capture numeracy for float values. The probing model is trained on a randomly shuffled 80% of the *Float Range* and tested on the remaining 20%. See the text for details on selecting decimal values. We show the mean alongside the standard deviation over 5 different random shuffles.

the ranges [151,160], [151,180], and [151,200]; all methods struggle, especially token vectors.

**Augmenting Data to Aid Extrapolation** Of course, in many real-world tasks it is possible to ameliorate these extrapolation failures by augmenting the training data (i.e., turn extrapolation into interpolation). Here, we apply this idea to aid in training NAQANet for DROP. For each superlative and comparative example, we duplicate the example and modify the numbers in its paragraph using the *Add* and *Multiply* techniques de-

Interpolation <i>Integer Range</i>	List Maximum (5-classes)
	[-50,50]
Rand. Vectors	0.23 ± 0.12
Word2Vec	0.89 ± 0.02
GloVe	0.89 ± 0.03
ELMo	0.96 ± 0.01
BERT	0.94 ± 0.02
Char-CNN	0.95 ± 0.07
Char-LSTM	0.97 ± 0.02

Table 6: *Interpolation with negatives (e.g., “-18”) on list maximum.* Pre-trained embeddings capture numeracy for negative values.

scribed in Section 2.5. Table 11 shows that this data augmentation can improve both interpolation and extrapolation, e.g., the accuracy on superlative questions with large numbers can *double*.

Extrapolation <i>Test Range</i>	List Maximum (5-classes)		
	[151,160]	[151,180]	[151,200]
Rand. Vectors	0.17	0.22	0.15
Untrained CNN	0.80	0.47	0.41
<i>Pre-trained</i>			
Word2Vec	0.14	0.16	0.11
GloVe	0.19	0.17	0.21
ELMo	0.65	0.57	0.38
BERT	0.35	0.11	0.14
<i>Learned</i>			
Char-CNN	0.81	0.75	0.73
Char-LSTM	<b>0.88</b>	<b>0.84</b>	<b>0.82</b>
<i>DROP</i>			
NAQANet	0.31	0.29	0.25
- GloVe	0.58	0.53	0.48

Table 7: *Extrapolation on list maximum.* The probing model is trained on the integer range [0,150] and evaluated on integers from the *Test Range*. The probing model struggles to extrapolate when trained on the pre-trained embeddings.

## 4 Discussion and Related Work

An open question is how the training process elicits numeracy for word vectors and contextualized embeddings. Understanding this, perhaps by tracing numeracy back to the training data, is a fruitful direction to explore further (c.f., influence functions (Koh and Liang, 2017; Brunet et al., 2019)).

More generally, numeracy is one type of emergent knowledge. For instance, embeddings may capture the size of objects (Forbes and Choi, 2017), speed of vehicles, and many other “commonsense” phenomena (Yang et al., 2018). Vendrov et al. (2016) introduce methods to encode the order of such phenomena into embeddings for concepts such as hypernymy; our work and Yang et al. (2018) show that a relative ordering naturally emerges for certain concepts.

In concurrent work, Naik et al. (2019) also explore numeracy in word vectors. Their methodology is based on variants of nearest neighbors and cosine distance; we use neural network probing classifiers which can capture highly non-linear dependencies between embeddings. We also explore more powerful embedding methods such as ELMo, BERT, and learned embedding methods.

**Probing Models** Our probes of *numeracy* parallel work in understanding the linguistic capabilities (*literacy*) of neural models (Conneau et al., 2018; Liu et al., 2019). LSTMs can remember sentence length, word order, and which words were present in a sentence (Adi et al., 2017). Khandel-

wal et al. (2018) show how language models leverage context, while Linzen et al. (2016) demonstrate that language models understand subject-verb agreement.

**Numerical Value Prediction** Spithourakis and Riedel (2018) improve the ability of language models to *predict* numbers, i.e., they go beyond categorical predictions over a fixed-size vocabulary. They focus on improving models; our focus is probing embeddings. Kotnis and García-Durán (2019) predict numerical attributes in knowledge bases, e.g., they develop models that try to predict the population of Paris.

**Synthetic Numerical Tasks** Similar to our synthetic numerical reasoning tasks, other work considers sorting (Graves et al., 2014), counting (Weiss et al., 2018), or decoding tasks (Trask et al., 2018). They use synthetic tasks as a testbed to prove or design better models, whereas we use synthetic tasks as a probe to understand token embeddings. In developing the Neural Arithmetic Logic Unit, Trask et al. (2018) arrive at similar conclusions regarding extrapolation: neural models have difficulty outputting numerical values outside the training range.

## 5 Conclusion

How much do NLP models know about numbers? By digging into a surprisingly successful model on a numerical reasoning dataset (DROP), we discover that pre-trained token representations naturally encode numeracy.

We analyze the limits of this numeracy, finding that CNNs are a particularly good prior (and likely the cause of ELMo’s superior numeracy compared to BERT) and that it is difficult for neural models to extrapolate beyond the values seen during training. There are still many fruitful areas for future research, including discovering *why* numeracy naturally emerges in embeddings, and what other properties are similarly emergent.

## Acknowledgements

We thank Mark Neumann, Suchin Gururangan, Pranav Goel, Shi Feng, Nikhil Kandpal, Dheeru Dua, the members of AllenNLP and UCI NLP, and the reviewers for their valuable feedback.



## References

- Yossi Adi, Einat Kermary, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. Finegrained analysis of sentence embeddings using auxiliary prediction tasks. In *ICLR*.
- Marc-Etienne Brunet, Colleen Alkalay-Houlihan, Ashton Anderson, and Richard Zemel. 2019. Understanding the origins of bias in word embeddings. In *ICML*.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. In *ACL*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *NAACL*.
- Maxwell Forbes and Yejin Choi. 2017. Verb physics: Relative physical knowledge of actions and objects. In *ACL*.
- Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. Neural Turing machines. *arXiv preprint arXiv:1410.5401*.
- Urvashi Khandelwal, He He, Peng Qi, and Dan Jurafsky. 2018. Sharp nearby, fuzzy far away: How neural language models use context. In *ACL*.
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *ICML*.
- Bhushan Kotnis and Alberto García-Durán. 2019. Learning numerical attributes in knowledge bases. In *AKBC*.
- Jayant Krishnamurthy, Pradeep Dasigi, and Matt Gardner. 2017. Neural semantic parsing with type constraints for semi-structured tables. In *EMNLP*.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of LSTMs to learn syntax-sensitive dependencies. In *TACL*.
- Nelson F Liu, Matt Gardner, Yonatan Belinkov, Matthew Peters, and Noah A Smith. 2019. Linguistic knowledge and transferability of contextual representations. In *NAACL*.
- Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhres, and Armand Joulin. 2018. Advances in pre-training distributed word representations. In *LREC*.
- Aakanksha Naik, Abhilasha Ravichander, Carolyn Rose, and Eduard Hovy. 2019. Exploring numeracy in word embeddings. In *ACL*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *EMNLP*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *North American Association for Computational Linguistics*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *EMNLP*.
- Abhilasha Ravichander, Aakanksha Naik, Carolyn Rose, and Eduard Hovy. 2019. EQUATE: A benchmark evaluation framework for quantitative reasoning in natural language inference. *arXiv preprint arXiv:1901.03735*.
- Andrew M Saxe, Pang Wei Koh, Zhenghao Chen, Maneesh Bhand, Bipin Suresh, and Andrew Y Ng. 2011. On random weights and unsupervised feature learning. In *ICML*.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. Analysing mathematical reasoning abilities of neural models.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *ICLR*.
- Georgios Spithourakis and Sebastian Riedel. 2018. Numeracy for language models: Evaluating and improving their ability to predict numbers. In *ACL*.
- Andrew Trask, Felix Hill, Scott Reed, Jack W. Rae, Chris Dyer, and Phil Blunsom. 2018. Neural arithmetic logic units. In *NeurIPS*.
- Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2016. Order-embeddings of images and language. In *ICLR*.
- Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. 2016. Order matters: Sequence to sequence for sets. In *ICLR*.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018. On the practical computational power of finite precision rnns for language recognition. In *ACL*.
- Yiben Yang, Larry Birnbaum, Ji-Ping Wang, and Doug Downey. 2018. Extracting commonsense properties from embeddings with limited human guidance. In *ACL*.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V. Le. 2018. QANet: Combining local convolution with global self-attention for reading comprehension. In *ICLR*.