

# *Probing with Noise:* Unpicking the Warp and Weft of Embeddings

Filip Klubička and John D. Kelleher

ADAPT Centre, Technological University Dublin, Ireland  
{filip.klubicka, john.kelleher}@adaptcentre.ie

## Abstract

Improving our understanding of how information is encoded in vector space can yield valuable interpretability insights. Alongside vector dimensions, we argue that it is possible for the vector norm to also carry linguistic information. We develop a method to test this: an extension of the probing framework which allows for relative intrinsic interpretations of probing results. It relies on introducing noise that ablates information encoded in embeddings, grounded in random baselines and confidence intervals. We apply the method to well-established probing tasks and find evidence that confirms the existence of separate information containers in English GloVe and BERT embeddings. Our correlation analysis aligns with the experimental findings that different encoders use the norm to encode different kinds of information: GloVe stores syntactic and sentence length information in the vector norm, while BERT uses it to encode contextual incongruity.

## 1 Introduction

Probing in NLP, as defined by [Conneau et al. \(2018\)](#), is a classification problem that predicts linguistic properties using dense embeddings as training data. The framework rests on the assumption that the probe’s success at a given task indicates that the encoder is storing information on the pertinent linguistic properties. Probing has quickly become an essential tool for encoder interpretability, by providing interesting insights into embeddings.

In essence, embeddings are vectors positioned in a shared multidimensional vector space, and vectors are geometrically defined by two aspects: having both a **direction** and **magnitude** ([Hefferon, 2018](#), page 36). Direction is the position in the space that the vector points towards (expressed by its dimension values), while magnitude is a vector’s length, defined as its distance from the origin (expressed by the vector norm) ([Anton and Rorres, 2013](#), page 131). It is understood that information

contained in a vector is encoded in the dimension values, which are most often studied in NLP research (see §6). However, information can be encoded in a representational vector space in more implicit ways, and relations can be inferred from more than just vector dimension values.

We hypothesise that it is possible for the vector magnitude—the norm—to carry information as well. Though it is a distributed property of a vector’s dimensions, the norm not only relates the distance of a vector from the origin, but indirectly also its distance from other vectors. Two vectors could be pointing in the exact same direction, but their distance from the origin might differ dramatically.<sup>1</sup> A similar effect has been observed in the literature: for many word embedding algorithms, the norm of the word vector correlates with the word’s frequency ([Schakel and Wilson, 2015](#)). E.g. in fastText embeddings the vectors of stop words (the most frequent words in English) are positioned closer to the origin than content words ([Balodis and Deksne, 2018](#)); and [Goldberg \(2017\)](#) notes that for many embeddings normalising the vectors removes word frequency information. Additionally, the norm plays an integral part in BERT’s attention layer, controlling the levels of contribution from frequent, less informative words by controlling the norms of their vectors ([Kobayashi et al., 2020](#)). It stands to reason that the norm could be leveraged by embedding models to encode other linguistic information as well. Hence, we argue that a vector representation has two **information containers**: vector *dimensions* and the vector *norm* (the titular *warp* and *weft*). In this paper, we test the assumption that these two components can be used to encode different types of information.

To this end, we need a probing method that provides an intrinsic evaluation of any given embed-

<sup>1</sup>Mathematically, two vectors can only be considered equal if both their direction and magnitude are equal ([Anton and Rorres, 2013](#), page 137).

ding representation, for which the typical probing pipeline is not suited. We thus extend the existing probing framework by introducing random noise into the embeddings. This enables us to do an intrinsic evaluation of a single encoder by testing whether the noise disrupted the information in the embedding being tested. The right application of noise enables us to determine which embedding component the relevant information is encoded in, by ablating that component’s information. In turn, this can inform our understanding of how certain linguistic properties are encoded in vector space. We call the method *probing with noise* and demonstrate its generalisability to both contextual and static encoders by using it to intrinsically evaluate English GloVe and BERT embeddings on a number of established probing tasks.

This paper’s main contributions are: (a) a methodological extension of the probing framework: *probing with noise*; (b) an array of experiments demonstrating the method on a range of probing tasks; and (c) an exploration of the importance of the vector norm in encoding linguistic phenomena in different embedding models.

## 2 Method: Probing With Noise

Our method is an extension of the typical probing pipeline (steps 1-6), incorporated as steps 7 and 8:

1. Choose a probing task
2. Choose or design an appropriate dataset
3. Choose a word/sentence representation
4. Choose a probing classifier (the probe)
5. Train the probe on the embeddings as input
6. Evaluate the probe’s performance on the task
- 7. Introduce systematic noise in the embedding**
- 8. Repeat training, evaluate and compare**

Usually, the evaluation score from step 6 is used as a basis to make inferences regarding the presence of the probed information in embeddings. Different encoders are compared based on their evaluation score and the probe’s relative performance can inform which model stores the information more saliently. Though ours may seem like a minor addition, it changes the approach conceptually. Now, rather than providing the final score, the output of step 6 establishes an intrinsic, *vanilla baseline*. Embeddings with noise injections can then be compared against it in steps 7 and 8, offering a relative intrinsic interpretation of the evaluation. In other

words, using relative information between a vector representation and targeted ablations of itself allows for inferences to be made on where information is encoded in embeddings.

The method relies on three supporting pillars: (a) random baselines, which in tandem with the vanilla baseline provide the basis for a relative evaluation; (b) statistical significance derived from confidence intervals, which informs the inferences we make based on the relative evaluation; and (c) targeted noise, which enables us to examine where the information is encoded. We describe them in the following subsections, starting with the noise.

### 2.1 Choosing the Noise

The nature of the noise is crucial for our method, as the goal is to systematically disrupt the content of the information containers in order to identify whether a container encodes the information. We use an ablation method to do this: by introducing noise into either container we “sabotage” the representation, in turn identifying whether the information we are probing for has been removed. Though we introduce random noise, our choice of how to apply it is systematic, as it is important that the noising function applied to one container leaves the information in the remaining container intact, otherwise the results will not offer relevant insight.

**Ablating the Dimension Container:** The noise function for ablating the dimensions needs to remove its information completely, while leaving the norm intact. It should also not change the dimensionality of the vector, given that a change in the dimensionality of a feature also changes the chance of the probe finding a random or spurious hyper-plane that performs well on the data sample. Maintaining the dimensionality thus ensures that the probability of the model finding such a lucky split in the feature space remains unchanged.

Our noise function satisfies these constraints: for each embedding in a dataset, we generate a new, random vector of the same dimensionality, then scale the new dimension values to match the norm of the original vector. This invalidates any semantics assigned to a particular dimension as the values are replaced with meaningless noise, while retaining the original vector’s norm values.

**Ablating the Norm Container:** To remove information potentially carried by a vector’s norm while retaining dimension information, we apply a noising function analogous to the previous one: for

each embedding we generate a random norm value, and then scale the vector’s original dimension values to match the new norm. This randomises vector magnitudes, while the relative sizes of the dimensions remain unchanged. In other words, all vectors will keep pointing in the same directions, but any information encoded by differences in magnitude is removed.<sup>2</sup>

**Ablating Both Containers:** The two approaches are not mutually exclusive: applying both noising functions should have a compounding effect and ablate both information containers simultaneously, essentially generating a completely random vector with none of the original information.

## 2.2 Random Baselines

Even when no information is encoded in an embedding, the train set may contain class imbalance, and the probe can learn the distribution of classes. To account for this, as well as the possibility of a powerful probe detecting an empty signal (Zhang and Bowman, 2018), we need to establish informative random baselines against which we can compare the probe’s performance.

We employ two such baselines: (a) we assert a random prediction onto the test set, negating any information that a classifier could have learned, class distributions included; and (b) we train the probe on randomly generated vectors, establishing a baseline with access only to class distributions.

## 2.3 Confidence Intervals

Finally, we must account for the degrees of randomness, which stem from two sources: (1) the probe may contain a stochastic component, e.g. a random weight initialisation; (2) the noise functions are highly stochastic (i.e. sampling random norm/dimension values). Hence, evaluation scores will differ each time the probe is trained, making relative comparisons of scores problematic. To mitigate this, we retrain and evaluate each model 50 times reporting the average score of all runs, essentially bootstrapping over the random seeds.

To obtain statistical significance for the averages, we calculate a 99% confidence interval (CI) to confirm that observed differences in the averages of different model scores are significant. We use

---

<sup>2</sup>We are conscious that vectors have more than one kind of norm, so choosing which norm to scale to might not be trivial. We have explored this in supplementary experiments and found that in our framework there is no significant difference between scaling to the L1 norm vs. L2 norm.

the CI range when comparing evaluation scores of probes on any two noise models to determine whether they come from the same distribution: if there is overlap in the range of two possible averages they might belong to the same distribution and there is no statistically significant difference between them. Using CIs in this way gives us a clearly defined decision criterion on whether any model performances are different.

## 3 Data

In our experiments we use 10 established probing task datasets for the English language introduced by Conneau et al. (2018). The goal of the multi-class *Sentence Length* (SL) probing task is to predict the length of the sentence as binned in 6 possible categories, while *Word Content* (WC) is a task with 1000 words as targets, predicting which of the target words appears in a given sentence. The *Subject* and *Object Number* tasks (SN and ON) are binary classification tasks that predict the grammatical number of the subject/object of the main clause as being singular or plural, while the *Tense* (TE) task predicts whether the main verb of the sentence is in the present or past tense. The *Coordination Inversion* (CIN) task distinguishes between a sentence where the order of two coordinated clausal conjoints has been inverted or not. *Parse Tree Depth* (TD) is a multi-class prediction task where the goal is to predict the maximum depth of the sentence’s syntactic tree, while *Top Constituents* (TC) predicts one of 20-classes of the most common syntactic top-constituent sequences. In the *Bigram Shift* (BS) task, the goal is to predict whether two consecutive tokens in the sentence have been inverted, and *Semantic Odd Man Out* (SOMO) is a task predicting whether a noun or verb was replaced with a different noun or verb. We use these datasets as published in their totality, with no modifications.<sup>3</sup> We also consider these tasks to represent examples of different language domains: surface information (SL,WC), morphology (SN,ON,TE), syntax (TD,TC,CIN) and contextual incongruity (BS,SOMO). This level of abstraction can lend itself to interpreting the experimental results, as there may be similarities across tasks in the same domain (note that Durrani et al. (2020) follow a similar line of reasoning).

---

<sup>3</sup><https://github.com/facebookresearch/SentEval/tree/master/data/probing>

## 4 Experiments

### 4.1 Models and Implementation

Given the current prominence of contextual encoders, such as BERT (Devlin et al., 2019), ELMo (Peters et al., 2018b) and their derivatives, they are an obvious choice for the application of our method. However, rather than compare different contextual encoders, we prefer to draw a contrastive comparison with a static encoder, such as GloVe (Pennington et al., 2014), which is a distributed representation based on a word to word co-occurrence matrix. This provides insight into both models and demonstrates the method’s generalisability to more than one type of encoder. In our experiments we examine BERT and GloVe embeddings.

Note that all the probing datasets we use are framed as classification tasks at the sentence level (see §3), so our experiments require sentence representations. We use pretrained versions of BERT and GloVe to generate embeddings for each sentence. The BERT model generates 12 layers of embedding vectors with each layer containing a separate 768-dimensional embedding for each word, so we average the word embeddings in BERT’s final layer, resulting in a 768-dimensional sentence embedding. We take the same mean pooling approach with GloVe, which yields a 300-dimensional sentence embedding for each sentence. While BERT uses sub-word tokens to get around out of vocabulary tokens, in the rare instance of encountering an OOV with GloVe, we generate a random word embedding in its stead.

In each set of experiments, the sentence embeddings are used as input to a Multi-Layered Perceptron (MLP) classifier, which labels them according to the probing task. We evaluate the performance of all probes using the AUC-ROC score.<sup>4</sup> Regarding implementation and parameter details, we used the bert-base-uncased BERT model from the `pytorch_pretrained_bert` library<sup>5</sup> (Paszke et al., 2019), a pre-trained GloVe model<sup>6</sup> and for the MLP probe we used the scikit-learn MLP implementation (Pedregosa et al., 2011) using the default pa-

<sup>4</sup>[https://scikit-learn.org/stable/modules/generated/sklearn.metrics.roc\\_auc\\_score.html](https://scikit-learn.org/stable/modules/generated/sklearn.metrics.roc_auc_score.html)

<sup>5</sup><https://pypi.org/project/pytorch-pretrained-bert/>

<sup>6</sup>The larger common crawl vectors: <https://nlp.stanford.edu/projects/glove/>

rameters.<sup>7,8</sup>

### 4.2 Chosen Noise Models

As described in §2, we remove information from the norm by sampling random norm values and scaling the vector dimensions to the new norm. However, considering that vectors have more than one calculable norm, the scaling can be done to match more than one norm value. We have examined the effects of scaling to both the L1 and L2 norms, as they are most widely used in NLP, and found that applying our norm ablation noise function to scale to either norm removes information from both norms (see Table 3).<sup>9</sup> In order to streamline the results presentation, henceforth when discussing norm ablations we only report results pertaining to scaling to the L2 norm.

To ablate information encoded in the dimension container, we randomly sample dimension values and then scale them to match the original norm of the vector (see §2).<sup>10</sup> We expect this to fully remove all interpretable information encoded in the dimension values, making the norm the only information container available to the probe. Applying both noise functions together on the same vector should remove any information encoded in it.

Finally, we use the vanilla BERT and GloVe sentence embeddings in their respective evaluations as vanilla baselines against which the models with noise are compared. Here the probe has access to both information containers: dimensions and norm. However, it is also important to establish the vanilla baseline’s performance against the random baselines: we need to confirm whether the information is in fact encoded somewhere in the embeddings.

<sup>7</sup>activation='relu', solver='adam', max\_iter=200, hidden\_layer\_sizes=100, learning\_rate\_init=0.001, batch\_size=min(200,n\_samples), early\_stopping=False, weight\_init.  $W \sim \mathcal{N}\left(0, \sqrt{6/(fan_{in} + fan_{out})}\right)$  (scikit relu default). See: [https://scikit-learn.org/stable/modules/generated/sklearn.neural\\_network.MLPClassifier.html](https://scikit-learn.org/stable/modules/generated/sklearn.neural_network.MLPClassifier.html)

<sup>8</sup>Code available here: <https://github.com/GreenParachute/probing-with-noise>

<sup>9</sup>This contrasts with applying a normalisation function to the vector, where normalising to one of the norms removes information encoded in that norm, but retains, or even emphasises the information in the remaining norm, making normalisation an unsuitable ablation function (see §A for details).

<sup>10</sup>The random norm and dimension values are sampled uniformly from a range between the minimum and maximum norm/dimension values of the respective embeddings on all 10 datasets. BERT norm range: [7.1896,13.2854], BERT dimension range: [-5.427,1.9658]; GloVe norm range: [2.0041,8.0359], GloVe dimension range: [-2.5446,3.1976]



GloVe											Key
Model	SL		WC		SN		ON		TE		<i>Surface Info.</i> SL: Sentence Length WC: Word Content <i>Morphology</i> SN: Subject Number ON: Object Number TE: Tense <i>Syntax</i> CIN: Coordination Inversion TD: Parse Tree Depth TC: Top Constituents <i>Incongruity</i> BS: Bigram Shift SOMO: Semantic Odd Man Out
	auc	±CI	auc	±CI	auc	±CI	auc	±CI	auc	±CI	
rand. pred.	.5006	.0013	.4995	.001	.4996	.002	.4999	.0023	.4981	.0022	
rand. vec.	.4999	.0011	.5006	.0009	.499	.0022	.4998	.0024	.4997	.0024	
vanilla	.9475	.0005	.9974	.0001	.8114	.0014	.7805	.0013	.8632	.0014	
abl. N	.9384	.0005	.994	.0001	.8058	.0016	.7743	.0018	.8594	.0013	
abl. D	<b>.5481</b>	.0013	<b>.504</b>	.0011	.5003	.0022	.4994	.0024	.5013	.0025	
abl. D+N	.5001	.0011	.4999	.0008	.4987	.0024	.4994	.002	.4998	.0021	
Model	CIN		TD		TC		BS		SOMO		
	auc	±CI	auc	±CI	auc	±CI	auc	±CI	auc	±CI	
rand. pred.	.5004	.0022	.5005	.0012	.5005	.0009	.4998	.0022	.4999	.0026	
rand. vec.	.4993	.0022	.5002	.0014	.5004	.0009	.4989	.0023	.4991	.0023	
vanilla	.5493	.0019	.7799	.0012	.9512	.0004	.5017	.0021	.5291	.0021	
abl. N	.5437	.002	.7689	.001	.9438	.0004	.5034	.0024	.5235	.002	
abl. D	.5003	.0023	<b>.5137</b>	.0012	<b>.5331</b>	.0013	.499	.0026	.5005	.0021	
abl. D+N	.5004	.0021	.501	.0013	.4996	.0011	.4996	.0024	.5007	.0019	

Table 1: Experimental results on GloVe models and baselines. Reporting average AUC-ROC scores and confidence intervals (CI) of the average of all training runs. Cells shaded light grey belong to the same distribution as random baselines, dark grey cells share the vanilla baseline distribution, while scores significantly different from both the random and vanilla baselines are unshaded, while the most pertinent scores are marked in bold.

### 4.3 Results

Detailed experimental evaluation results for GloVe and BERT on each of the 10 probing tasks are presented in Tables 1 and 2 respectively. Note that all cells shaded light grey belong to the same distribution as random baselines on a given task, as there is no statistically significant difference between the different scores<sup>11</sup>; cells shaded dark grey belong to the same distribution as the vanilla baseline on a given task; and all cells that are not shaded contain a significantly different score than both the random and vanilla baselines, indicating that they belong to different distributions. The scores most pertinent to the result discussion are marked in bold.

**GloVe results:** The vanilla GloVe vectors outperform the random baselines on all tasks except BS. This is not surprising, as BS is essentially a local-context task, and GloVe does not encode context in such a localised manner. In all other tasks, at least some task-relevant information is encoded in the embeddings. Having established the vanilla results as a baseline for the ablations, we examine which information container encodes the relevant information: dimension or norm.

Generally, the results show that the answers are task-dependent. In the SN, ON, TE, CIN and

<sup>11</sup>We highlight that the *rand. vec.* baseline is equivalent to the scenario where both dimensions and norm are ablated (*abl. D+N*). While the two scenarios are arguably the exact same condition, we include both of them in the results presentation to demonstrate a consistent application of our methodology, where we consider *rand. vec.* to be a baseline, and the *abl. D+N* a sense-check of our ablation functions.

SOMO tasks, there is a substantial drop in the probe’s performance after ablating the dimension container and it is immediately comparable to random baselines. Furthermore, performance does not significantly change after also ablating the norm, indicating that for these tasks no pertinent information is stored in the norm, and that all the information the probe uses is stored in the dimensions.

However, the results for the SL, WC, TD and TC probes tell a different story. Once the dimension container is ablated from these vectors, although the performance drops markedly compared to vanilla, it does not quite reach the random baseline performance as observed in the above tasks.<sup>12</sup> These results indicate that for these tasks the relevant information is not contained *only* in the dimension container. Furthermore, when the dimension and norm ablation functions are applied together, this induces a further performance drop, and the resulting performance scores become comparable to the random baselines. This indicates that the vectors with ablated dimension information still contain residual information relevant to the task, which is removed when also ablating the norm, pointing to the fact that the norm contains some of the relevant information *regardless of what is encoded in the vector dimensions*.

We should note here that, while it is true that in

<sup>12</sup>This is true even in the case of WC, where the difference is really quite small, yet still statistically significant. Note that the WC task is a particularly unusual classification task, as there are 1000 possible classes to predict, which could explain the statistical significance of such a small difference.

BERT											Key <i>Surface Info.</i>
Model	SL		WC		SN		ON		TE		
	auc	±CI	auc	±CI	auc	±CI	auc	±CI	auc	±CI	
rand. pred.	.5002	.0006	.4996	.0012	.4995	.0021	.4988	.0022	.5007	.0021	SL: Sentence Length
rand. vec.	.5003	.0004	.4997	.0009	.5006	.002	.4996	.0024	.4993	.0021	WC: Word Content
vanilla	.9733	.0011	.982	.0003	.9074	.0008	.8674	.0019	.9135	.0008	<i>Morphology</i>
abl. N	.973	.0008	.9783	.0003	.9078	.0008	.8658	.0017	.9118	.0012	SN: Subject Number
abl. D	.5047	.0008	.5013	.0011	.4992	.0021	.5004	.0023	.5007	.0019	ON: Object Number
abl. D+N	.4997	.0008	.5	.0013	.5006	.0024	.4994	.0024	.4983	.0021	TE: Tense
Model	CIN		TD		TC		BS		SOMO		<i>Syntax</i>
	auc	±CI	auc	±CI	auc	±CI	auc	±CI	auc	±CI	CIN: Coordination
rand. pred.	.5007	.0022	.4999	.0012	.5001	.0013	.5011	.0020	.499	.0018	Inversion
rand. vec.	.5014	.0019	.4999	.0012	.5001	.0013	.5005	.0024	.5001	.0021	TD: Parse Tree Depth
vanilla	.7472	.0016	.7751	.0016	.9562	.0002	.9382	.0006	.6401	.0013	TC: Top Constituents
abl. N	.7492	.0018	.7709	.0016	.9547	.0004	.9371	.001	.6396	.0017	<i>Incongruity</i>
abl. D	.5049	.0021	.5004	.0013	<b>.5093</b>	.0019	<b>.556</b>	.0025	<b>.5272</b>	.002	BS: Bigram Shift
abl. D+N	.5015	.0035	.5	.0012	.5001	.001	.4972	.0035	.4997	.002	SOMO: Semantic Odd Man Out

Table 2: Experimental results on BERT models and baselines. Reporting average AUC-ROC scores and confidence intervals (CI) of the average of all training runs. Cells shaded light grey belong to the same distribution as random baselines, dark grey cells share the vanilla baseline distribution, while scores significantly different from both the random and vanilla baselines are unshaded, while the most pertinent scores are marked in bold.

all tasks ablating the norm alone causes a statistically significant drop in performance, this finding on its own should not be taken as an indicator that the norm encodes task-relevant information. Given how consistently small the drop is across all tasks ( $<0.1$ ), this is more likely an artefact of an interaction between the noising function and the GloVe vectors. The more reliable indicator of where the information is encoded is the experiment on dimension ablations compared to ablating both dimension and norm: if for a particular task performance remains above random after ablating dimensions, but drops to random when ablating both dimensions and norms, this shows that the norm is encoding at least part of the relevant information.

**BERT results:** The vanilla BERT vectors outperform random baselines across all tasks, including the BS task. When ablating the dimensions on most tasks, the probe’s performance drops dramatically and is comparable to random baselines. It does not change after also ablating the norm, indicating that no pertinent information is stored in BERT’s norm container for these tasks. However, the BS and SOMO tasks show that some of the task information is stored in BERT’s norm, as the performance drop when ablating dimensions is not comparable to random baselines, and only reaches that once the norm is also ablated. The same is true for the syntactic TC task, which is also the only BERT result that shows a similar trend as GloVe, though it seems that BERT stores far less TC information in the norm than GloVe does.

Ultimately, our experimental results allow us to make a number of general inferences: (a) the norm is indeed a separate information container, (b) on most tasks the vast majority of the relevant information is encoded in the dimension values, but can be supplemented with information from the norm, (c) though the information contained in the norm is not always very impactful, it is not negligible, (d) different encoders use the norm to carry different types of information, (e) specifically BERT stores information pertinent to the BS, SOMO and TC tasks in the norm, (f) while GloVe uses it to store SL, WC, TC and TD information.

#### 4.4 Norm Correlation Analysis

While we have demonstrated that information can be encoded in the norm, we wish to also understand the relationship between the norms and the probed information. We explore this with a Pearson correlation analysis: we test the correlation between each vector norm and the sentence labels on each probing task dataset.<sup>13</sup> The correlation results are presented in Table 3, and largely support our result interpretations from §4.3,<sup>14</sup> including that applying

<sup>13</sup>The Pearson test only works on continuous variables, but it is still possible to calculate with categorical variables if they are binary, by simply converting the categories to 0 and 1.

<sup>14</sup>In cases such as WC and TC where there are more than two categorical variables we can perform a Kruskal-Wallis test to determine a statistically significant difference between the categories. This does not quantify the difference in the same way as a Pearson test, and does not allow us to determine whether the correlation is positive or not, nor how strong it is. Instead we can only say that we performed the test and found

Task	Vectors	GloVe		BERT	
		L1	L2	L1	L2
SL	Vanilla	-0.7278	-0.3758	-0.1564	-0.1039
	Abl. norm	-0.1893	-0.0025	-0.0417	-0.0013
SN	Vanilla	0.0360	0.0268	0.0071	0.0146
	Abl. norm	0.0036	-0.0033	-0.0035	-0.0021
ON	Vanilla	0.0013	0.0008	-0.0736	-0.0583
	Abl. norm	0.0009	0.0013	-0.0181	-0.0010
TE	Vanilla	0.1152	0.0571	0.0542	0.0413
	Abl. norm	0.0277	-0.0031	0.0097	-0.0030
TD	Vanilla	-0.0817	0.1908	-0.0415	-0.0251
	Abl. norm	-0.0665	0.0016	-0.0163	-0.0045
CIN	Vanilla	-0.0019	-0.0094	-0.0755	-0.0638
	Abl. norm	0.0029	0.0018	-0.0152	-0.0015
BS	Vanilla	0.0040	0.0002	-0.3866	-0.3238
	Abl. norm	0.0022	0.0006	-0.0978	-0.0005
SO MO	Vanilla	-0.0464	-0.0222	-0.2414	-0.2305
	Abl. norm	-0.0105	0.0000	-0.0420	0.0021

Table 3: Pearson correlation coefficients between the class labels and vector norms for vanilla vectors and vectors with ablated norms.

our noise function to ablate the norm fully removes the information from the norms: the correlation between either norm and the class labels drops to  $\approx 0$ ,<sup>15</sup> indicating that information encoded by the norm and any distinguishing properties it may have had have been removed.

The data shows that most task labels do not exhibit a correlation with the vanilla GloVe norm. There is a moderate positive correlation between TD and the L2 norm, but not the L1 norm, and a weak positive correlation between TE and the L1 norm, but not the L2 norm. There is a high correlation between the SL labels and both norms, showing that GloVe uses the norm to encode sentence length, as reflected in our experiments in §4.

When it comes to vanilla BERT, most task labels do not exhibit a correlation with the norms. However, both norms have a weak negative correlation with SL, and a moderate negative correlation with BS and SOMO. The latter two are most highly correlated with BERT’s norm, which also aligns with our experimental findings in §4.

## 5 Discussion

The correlation coefficients in Table 3 can be interpreted in terms of how these linguistic phenomena are encoded in vector space. A negative correlation coefficient means that larger norms indicate a

the results to be significant, indicating some correlation.

<sup>15</sup>Except in GloVe-SL-L1 where the coefficient ‘only’ drops from strongly correlated to weakly correlated.

negative class, while a positive coefficient means that larger norms indicate a positive class. For example, the negative correlation in SL-GloVe and SL-BERT indicates that longer sentences are positioned closer to the origin. The same interpretation holds for BERT embeddings on the BS and SOMO tasks; e.g. in SOMO a sentence containing an out of context word is positioned closer to the origin.

It is interesting that BERT’s norm stores information on the BS and SOMO tasks specifically. Their common thread is a violation of the local context of the affected words: though the overall context and structure of the sentence is unaffected, there is a small, localised disruption in co-occurrences. Hence, these tasks capture contextual incongruity. Given that we know that BERT is a contextual encoder, and that its self-attention uses the vector norm to control the levels of contribution from less informative words (Kobayashi et al., 2020), we suspect that this gives it the capabilities to accurately model these short-distance dependencies and word co-occurrence probabilities, concepts which strongly correspond to local contextual incongruity. BERT is evidently capable of encoding this signal well, and seems to be using its norm to supplement the encoding of the phenomenon in such a way that it positions sentences exhibiting local contextual incongruity closer to the origin, relative to sentences that do not contain it. Furthermore, BERT’s ability to model incongruity via the norm could essentially be frequency-based, similar to how some word embeddings encode word frequency in the norm. In contrast, GloVe is a static encoder and exhibits no indication that it stores this information in the norm, or indeed any ability to accurately model this phenomenon at all, but uses the norm to store surface-level and syntactic information.

We emphasise the importance of the norm as it expands our understanding of the way information is encoded in vector space, but it could also have important implications for downstream tasks involving operations on vectors: e.g. the calculation of a cosine similarity measure normalises the vectors being compared. This nullifies the information in the norm, reducing the comparison to one of directions (i.e. dimensions), and any linguistic information encoded in the norm will be lost and unaccounted for when making the comparison.

## 6 Related Work

Probing has been proposed seemingly independently by different groups of NLP researchers (Ettinger et al., 2016; Shi et al., 2016; Veldhoen et al., 2016; Adi et al., 2017) and has gained significant momentum in the community, helping to explore different aspects of text encodings (e.g. Hupkes et al. (2018); Giulianelli et al. (2018); Krasnowska-Kieraś and Wróblewska (2019); Tenney et al. (2019a); Lin et al. (2019); Şahin et al. (2020); Liu et al. (2021); Arps et al. (2022)). Probes trained on various representations successfully predict surface properties of sentences (Adi et al., 2017; Conneau et al., 2018), POS and morphological information (Belinkov et al., 2017a; Liu et al., 2019), as well as syntactic (Zhang and Bowman, 2018; Peters et al., 2018a; Tenney et al., 2019b), semantic (Belinkov et al., 2017b; Ahmad et al., 2018; Conia and Navigli, 2022), and even number (Wallace et al., 2019), emotions (Qian et al., 2016), idiomaticity (Salton et al., 2016; Nedumpozhimana and Kelleher, 2021; Garcia et al., 2021; Nedumpozhimana et al., 2022) and world knowledge information (Ettinger, 2020), among others (Belinkov and Glass, 2019; Rogers et al., 2020; Koto et al., 2021; Ousidhoum et al., 2021; Aghazadeh et al., 2022).

Furthermore, some dichotomies have emerged in the literature, due to nuanced differences in the pre-suppositions behind probing approaches. Ravichander et al. (2020) distinguish varying points of view on embeddings, highlighting a difference between *instrumentative* and *agentive* probing. Vig et al. (2020) view probing as a method of analysis and distinguish two types of methods: *structural* and *behavioural*. Additionally, Pimentel et al. (2020) and Voita and Titov (2020) take an information-theoretic perspective on embeddings, highlighting the tension between probing identifying the mere *presence* of information, versus its *extractability*. We position our work as being **instrumentative**, i.e. we view embeddings as tools that extract and store knowledge from text; we consider our probing method to be **structural**, i.e. it provides insight into how information is encoded within the representation and the vector space; and the goal of our work is to identify the **presence** of information in embedding components. It is important to clearly signpost this position in order to avoid confusion and emphasise that our chosen approach is sufficient to address our research questions.

Meanwhile, recent work calls for greater rigor in evaluation approaches in NLP (McCoy et al., 2020; Sadeqi Azer et al., 2020; Card et al., 2020), advocating for more widespread use of statistical tests on common benchmarks. Probing has attracted similar criticism: Hewitt and Liang (2019) have shown that under certain conditions, above-random probing accuracy can be achieved even when probing for linguistically-meaningless noise. Recent work addresses some of these problems by constructing counterfactual representations in order to compare the performance of the probe with and without the pertinent information (Feder et al., 2020). Similarly, Elazar et al. (2020) remove the relevant information from the representation, allowing a comparison of probe performance with and without the removed information; not unlike the intrinsic probe of Torroba Hennigen et al. (2020) who focus on isolating the dimensions that encode relevant information. In essence, these recent efforts address the issue of relativising probe interpretations by removing information from the encoding; in that sense, our work finds its place alongside them. However, our method is not meant to remove specific information, but is more exploratory in nature, with a focus on understanding where within an embedding certain information is encoded. Our use of confidence intervals gives us a way to claim statistically significant differences in our evaluations, offering a more principled basis for result interpretation.

Our work also contributes to the relatively scarce study of the role of the norm: Adi et al. (2017) explain its correlation with SL information due to the central limit theorem (which we see does not apply to BERT as its vector values are not centred around zero). Hewitt and Manning (2019) show that the squared L2 norm of BERT and ELMo corresponds to the depth of the word in a parse tree (a finding we could not confirm as they probe embeddings at the word level, unlike our work). In contrast, work on the role of dimensions as carriers of specific types of information is plentiful (e.g. Karpathy et al. (2015); Qian et al. (2016); Bau et al. (2019); Dalvi et al. (2019); Lakretz et al. (2019)). Work complementary to ours (Torroba Hennigen et al., 2020) which focuses on the dimension container also highlights the need for an intrinsic probe of embedding models, and shows that most linguistic properties are reliably encoded by only a handful of dimensions, a finding consistent with Durrani et al. (2020) and Durrani et al. (2022).



## 7 Conclusion

We have developed a method of enquiry that provides geometric insights into embeddings and show experimental evidence that both BERT and GloVe embeddings use two separate information containers to store different types of linguistic information. Our findings show that BERT primarily uses the norm to store contextual incongruity information and positions incongruous sentences closer to the origin. Meanwhile, GloVe stores much more syntactic information in its norm than BERT, but does not store contextual information at all, and mainly stores surface-level information in the norm.

*Probing with noise* can shift perspectives and broaden our understanding of embeddings, demonstrated by our experiments which provide novel insights into contextual and static encoders. However, they are by no means exhaustive: deeper and further applications of the method, such as exploring a host of other representations, different pooling strategies or tracking behavior across embedding layers, exploring word-level tasks or folding in additional datasets, are all fruitful avenues for future work. Fortunately, the method is robust enough to be applied to any encoder and any dataset, whether it is at the word or sentence level, which allows for systematic further study.

## Limitations

While our insights into how linguistic information can be encoded in embeddings are valuable on their own merit, our experiments mainly serve the purpose of validating the *probing with noise* method, in demonstrating that it can produce relevant insights on different types of embeddings. Hence we did not have scope to more thoroughly pursue many of the topics touched upon in the paper.

One example is our choice in generating sentence embeddings needed to probe for sentence-level information. The encoders we have used generate word-level embeddings, so we average the word embeddings in each sentence, as this is one of the most popular ways to generate sentence representations. However, there are other known approaches available to choose from, such as max pooling and min pooling, or, when it comes to BERT, using the CLS token.<sup>16</sup> Indeed, rather than a pooling strategy, using direct sentence-level

<sup>16</sup>Presumably, we may have observed a crisper effect in BERT encoding incongruity using min or max pooling, given that the BS task mainly affects only a few vectors in a sentence.

representations such as doc2vec (Le and Mikolov, 2014) or SentenceBERT (Reimers and Gurevych, 2019) might also be prudent, as well as applying the method to word-level representations, for which this paper did not allow scope.

Similarly, we have consistently used only one probing classifier, an MLP with default parameters, and we cannot say whether parameter tuning or different probes would yield different results. These choices were made consciously, in order to avoid adding more variables to our line of enquiry and increasing the complexity of our experiments, yet it is still a limitation in the sense that we do not know whether the findings generalise to other probes.

It is also worth noting that the correlation study in §4 comes with the limitation of only describing linear relationships, whereas it is possible that connections between variables can be non-linear. We argue that this demonstrates the value of our method, which allows for a non-linear probe to test for non-linear relationships. While even this limited correlation test can provide interesting insights, much more can be done to study both the norm and the dimension container—we have just barely scratched the surface. Indeed, we have considered only the most fundamental geometric properties of vectors, yet vectors have other (distributed) properties that could potentially be considered distinct information containers in their own right, such as the vector’s minimum and maximum value, their ratio, the entropy in the vector etc. Thankfully the principles underpinning our method can be expanded to include other types of noise that help discriminate other possible geometric properties of embeddings as information containers.

These points speak to the more general limitations of our research: like any empirical work, we measure behaviours on a number of data points and draw conclusions from these measurements. Thus there is a risk that our findings hold only for the datasets on which we measured or the models which were used to measure, be it encoders, probes or probing tasks, and it is possible that our findings might not generalise to other settings. While this issue is more epistemological than it is specific to our work, we must keep it in mind. Now, having demonstrated that a signal is detectable in our particular setting, a more comprehensive host of studies is needed to draw more general conclusions.

Another source of uncertainty stems from our use of off-the-shelf GloVe and BERT embeddings:

they have been trained on completely different datasets of dramatically varying sizes and content. To truly test the interaction of their architectures with our method, the training data used to train their word embeddings should be identical between both encoders, however implementing this was not feasible in practice. Granted, using off-the-shelf varieties does provide insight into the functioning of well-known and commonly used embeddings, but it consequently limits the comparability of their results as we cannot confidently distinguish whether differences in probe performance are due to differences in encoder architecture or training data.

While we acknowledge a number of the work’s limitations, we stress that all our choices have been made in a sound, informed and methodologically consistent manner. Here we simply highlight just how many choices have been made along the way, and how quickly the number of alternative paths grows the further back up the decision tree we look. While we believe that the work is fundamentally sound, each choice could have made for a drastically different suite of experiments and could potentially have yielded different results. In fact, we find this to be a very exciting motivator for future work, as this long list of “missed opportunities” only goes to show how young and rich this research area still is and how many more avenues there are to explore.

## Acknowledgements

This research was conducted with the financial support of Science Foundation Ireland under Grant Agreements No. 13/RC/2106 and 13/RC/2106\_P2 at the ADAPT SFI Research Centre at Technological University Dublin. ADAPT, the SFI Research Centre for AI-Driven Digital Content Technology, is funded by Science Foundation Ireland through the SFI Research Centres Programme, and is co-funded under the European Regional Development Fund.

## References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. In *Proceedings of ICLR, 2017*.
- Ehsan Aghazadeh, Mohsen Fayyaz, and Yadollah Yaghoobzadeh. 2022. [Metaphors in pre-trained language models: Probing and generalization across datasets and languages](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2037–2050, Dublin, Ireland. Association for Computational Linguistics.
- Wasi Uddin Ahmad, Xueying Bai, Zhechao Huang, Chao Jiang, Nanyun Peng, and Kai-Wei Chang. 2018. [Multi-task learning for universal sentence embeddings: A thorough evaluation using transfer and auxiliary tasks](#).
- Howard Anton and Chris Rorres. 2013. *Elementary linear algebra: applications version*. John Wiley & Sons.
- David Arps, Younes Samih, Laura Kallmeyer, and Hassan Sajjad. 2022. [Probing for constituency structure in neural language models](#).
- Kaspars Balodis and Daiga Deksnė. 2018. Intent detection system based on word embeddings. In *Artificial Intelligence: Methodology, Systems, and Applications*, pages 25–35, Cham. Springer International Publishing.
- Anthony Bau, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2019. [Identifying and controlling important neurons in neural machine translation](#). In *International Conference on Learning Representations*.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017a. [What do neural machine translation models learn about morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–872, Vancouver, Canada. Association for Computational Linguistics.
- Yonatan Belinkov and James Glass. 2019. [Analysis methods in neural language processing: A survey](#). *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2017b. [Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1–10, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Dallas Card, Peter Henderson, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, and Dan Jurafsky. 2020. [With little power comes great responsibility](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9263–9274, Online. Association for Computational Linguistics.
- Simone Conia and Roberto Navigli. 2022. [Probing for predicate argument structures in pretrained language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4622–4632, Dublin, Ireland. Association for Computational Linguistics.

- 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 *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. Association for Computational Linguistics.
- Bhavana Dalvi, Niket Tandon, Antoine Bosselut, Wentaoh Yih, and Peter Clark. 2019. [Everything happens for a reason: Discovering the purpose of actions in procedural text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4496–4505, Hong Kong, China. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nadir Durrani, Fahim Dalvi, and Hassan Sajjad. 2022. Linguistic correlation analysis: Discovering salient neurons in deepnlp models. *arXiv preprint arXiv:2206.13288*.
- Nadir Durrani, Hassan Sajjad, Fahim Dalvi, and Yonatan Belinkov. 2020. [Analyzing individual neurons in pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4865–4880, Online. Association for Computational Linguistics.
- Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2020. When bert forgets how to POS: Amnesic probing of linguistic properties and MLM predictions. *arXiv preprint arXiv:2006.00995*.
- Allyson Ettinger. 2020. What bert is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Allyson Ettinger, Ahmed Elgohary, and Philip Resnik. 2016. [Probing for semantic evidence of composition by means of simple classification tasks](#). In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 134–139, Berlin, Germany. Association for Computational Linguistics.
- Amir Feder, Nadav Oved, Uri Shalit, and Roi Reichart. 2020. CausaLM: Causal model explanation through counterfactual language models. *arXiv preprint arXiv:2005.13407*.
- Marcos Garcia, Tiago Kramer Vieira, Carolina Scarton, Marco Idiart, and Aline Villavicencio. 2021. [Probing for idiomaticity in vector space models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3551–3564, Online. Association for Computational Linguistics.
- 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. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 240–248.
- Yoav Goldberg. 2017. Neural network methods for natural language processing. *Synthesis Lectures on Human Language Technologies*, 10(1):117.
- Jim Hefferon. 2018. *Linear Algebra*. openintro.org.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dieuwke Hupkes, Sara Veldhoen, and Willem Zuidema. 2018. Visualisation and 'diagnostic classifiers' reveal how recurrent and recursive neural networks process hierarchical structure. *Journal of Artificial Intelligence Research*, 61:907–926.
- Andrej Karpathy, Justin Johnson, and Fei-Fei Li. 2015. [Visualizing and understanding recurrent networks](#). *CoRR*, abs/1506.02078.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. [Attention is not only a weight: Analyzing transformers with vector norms](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7057–7075, Online. Association for Computational Linguistics.
- Fajri Koto, Jey Han Lau, and Timothy Baldwin. 2021. [Discourse probing of pretrained language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3849–3864, Online. Association for Computational Linguistics.



- Katarzyna Krasnowska-Kieraś and Alina Wróblewska. 2019. [Empirical linguistic study of sentence embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5729–5739, Florence, Italy. Association for Computational Linguistics.
- Yair Lakretz, German Kruszewski, Theo Desbordes, Dieuwke Hupkes, Stanislas Dehaene, and Marco Baroni. 2019. [The emergence of number and syntax units in LSTM language models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 11–20, Minneapolis, Minnesota. Association for Computational Linguistics.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196. PMLR.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. [Open sesame: Getting inside bert’s linguistic knowledge](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zeyu Liu, Yizhong Wang, Jungo Kasai, Hannaneh Hajishirzi, and Noah A. Smith. 2021. [Probing across time: What does RoBERTa know and when?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 820–842, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- R. Thomas McCoy, Junghyun Min, and Tal Linzen. 2020. [BERTs of a feather do not generalize together: Large variability in generalization across models with similar test set performance](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 217–227, Online. Association for Computational Linguistics.
- Vasudevan Nedumpozhi and John Kelleher. 2021. [Finding BERT’s idiomatic key](#). In *Proceedings of the 17th Workshop on Multiword Expressions (MWE 2021)*, pages 57–62, Online. Association for Computational Linguistics.
- Vasudevan Nedumpozhi, Filip Klubička, and John D. Kelleher. 2022. [Shapley idioms: Analysing bert sentence embeddings for general idiom token identification](#). *Frontiers in Artificial Intelligence*, 5.
- Nedjma Ousidhoum, Xinran Zhao, Tianqing Fang, Yangqiu Song, and Dit-Yan Yeung. 2021. [Probing toxic content in large pre-trained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4262–4274, Online. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Luke Zettlemoyer, and Wen-tau Yih. 2018a. Dissecting contextual word embeddings: Architecture and representation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018b. [Deep contextualized word representations](#). *CoRR*, abs/1802.05365.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4609–4622, Online. Association for Computational Linguistics.
- Peng Qian, Xipeng Qiu, and Xuanjing Huang. 2016. [Investigating language universal and specific properties in word embeddings](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1478–1488, Berlin, Germany. Association for Computational Linguistics.



- Abhilasha Ravichander, Eduard Hovy, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2020. On the systematicity of probing contextualized word representations: The case of hypernymy in bert. In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 88–102.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. **A primer in BERTology: What we know about how BERT works**. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Erfan Sadeqi Azer, Daniel Khashabi, Ashish Sabharwal, and Dan Roth. 2020. **Not all claims are created equal: Choosing the right statistical approach to assess hypotheses**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5715–5725, Online. Association for Computational Linguistics.
- Gözde Gül Şahin, Clara Vania, Iliia Kuznetsov, and Iryna Gurevych. 2020. **LINSPECTOR: Multilingual probing tasks for word representations**. *Computational Linguistics*, 46(2):335–385.
- Giancarlo Salton, Robert Ross, and John Kelleher. 2016. **Idiom token classification using sentential distributed semantics**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 194–204, Berlin, Germany. Association for Computational Linguistics.
- Adriaan MJ Schakel and Benjamin J Wilson. 2015. Measuring word significance using distributed representations of words. *arXiv preprint arXiv:1508.02297*.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. **Does string-based neural MT learn source syntax?** In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1526–1534, Austin, Texas. Association for Computational Linguistics.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. **BERT rediscovers the classical NLP pipeline**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel Bowman, Dipanjan Das, et al. 2019b. What do you learn from context? probing for sentence structure in contextualized word representations. In *7th International Conference on Learning Representations, ICLR 2019*.
- Lucas Torroba Hennigen, Adina Williams, and Ryan Cotterell. 2020. **Intrinsic probing through dimension selection**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 197–216, Online. Association for Computational Linguistics.
- Sara Veldhoen, Dieuwke Hupkes, and Willem H Zuidema. 2016. Diagnostic classifiers revealing how neural networks process hierarchical structure. In *Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches (at NIPS)*.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Simas Sakenis, Jason Huang, Yaron Singer, and Stuart Shieber. 2020. Causal mediation analysis for interpreting neural nlp: The case of gender bias. *arXiv preprint arXiv:2004.12265*.
- Elena Voita and Ivan Titov. 2020. **Information-theoretic probing with minimum description length**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 183–196, Online. Association for Computational Linguistics.
- Eric Wallace, Yizhong Wang, Sujian Li, Sameer Singh, and Matt Gardner. 2019. **Do NLP models know numbers? probing numeracy in embeddings**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5307–5315, Hong Kong, China. Association for Computational Linguistics.
- Kelly Zhang and Samuel Bowman. 2018. **Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 359–361, Brussels, Belgium. Association for Computational Linguistics.

## A Appendix A

### A.1 Analysis of L1 and L2 Normalised Embeddings

Table 4 presents an extended Pearson correlation analysis that includes correlations between class labels and the norms of L1- and L2-normalised vectors, in addition to vanilla vectors and vectors with ablated norm information using our noising function as described in §2.

As supported by Goldberg (2017, page 117), the results show that normalising the vectors removes information encoded in the norm. This also comes

Task	Vectors	GloVe		BERT	
		L1	L2	L1	L2
SL	Vanilla	-0.7278	-0.3758	-0.1564	-0.1039
	L1 normal.	-0.0013	0.7161	0.0032	0.2195
	L2 normal.	-0.7027	0.0001	-0.2223	0.0001
	Abl. norm	-0.1893	-0.0025	-0.0417	-0.0013
SN	Vanilla	0.0360	0.0268	0.0071	0.0146
	L1 normal.	0.0028	-0.0228	-0.0010	0.0087
	L2 normal.	0.0255	-0.0019	-0.0086	-0.0003
	Abl. norm	0.0036	-0.0033	-0.0035	-0.0021
ON	Vanilla	0.0013	0.0008	-0.0736	-0.0583
	L1 normal.	-0.0016	0.0048	-0.0015	0.0892
	L2 normal.	-0.0004	-0.0015	-0.0901	0.0037
	Abl. norm	0.0009	0.0013	-0.0181	-0.0010
TE	Vanilla	-0.1152	-0.0571	-0.0542	-0.0413
	L1 normal.	-0.0020	0.1040	-0.0023	0.0659
	L2 normal.	-0.1071	-0.0006	-0.0691	-0.0018
	Abl. norm	-0.0317	-0.0007	-0.0116	0.0010
TD	Vanilla	-0.0817	0.1908	-0.0415	-0.0251
	L1 normal.	0.0005	0.3133	0.0021	0.0645
	L2 normal.	-0.3159	-0.0026	-0.0652	0.0000
	Abl. norm	-0.0665	0.0016	-0.0163	-0.0045
CIN	Vanilla	-0.0019	-0.0094	-0.0755	-0.0638
	L1 normal.	0.0000	-0.0062	-0.0047	0.0846
	L2 normal.	0.0065	0.0064	-0.0850	0.0034
	Abl. norm	0.0029	0.0018	-0.0152	-0.0015
BS	Vanilla	0.0040	0.0002	-0.3866	-0.3238
	L1 normal.	-0.0015	-0.0048	0.0004	0.4333
	L2 normal.	0.0056	-0.0019	-0.4357	0.0024
	Abl. norm	0.0022	0.0006	-0.0978	-0.0005
SO MO	Vanilla	-0.0464	-0.0222	-0.2414	-0.2305
	L1 normal.	0.0031	0.0401	0.0035	0.2213
	L2 normal.	-0.0392	-0.0014	-0.2219	0.0023
	Abl. norm	-0.0105	0.0000	-0.0420	0.0021

Table 4: Pearson correlation coefficients between the class labels and vector norms for vanilla vectors, L1 and L2 normalised vectors, as well as vectors with ablated L2 norm containers.

with a caveat: normalisation only removes information from the same order norm as the normalisation algorithm. We can observe this in the table: applying an L1 normalisation algorithm to the vectors seems to completely remove any information encoded in the L1 norm, as the correlation drops to  $\approx 0$ . The same happens to the correlation with the L2 norm when applying L2 normalisation. However, surprisingly, it seems that a given normalisation algorithm impacts the other norm as well. For example, in the BS task L2 normalisation nullifies the L2 norm’s correlation with the class labels, but in turn strengthens that correlation for the L1 norm, which intensifies from -0.39 to -0.44. On the other hand, L1 normalisation causes the same strengthening of correlation in the L2 norm, but also changes the sign—the L2 norm’s correlation with BS class

labels increases from -0.32 to 0.43.

This shows that on certain tasks, not only is the other norm unaffected by a normalisation procedure, but its correlation with the task labels increases. We observe this to varying degrees in SL, ON, TE and BS. Furthermore, while the correlation weakens in SOMO, it still exhibits the latter behaviour—the sign changes when the vectors are L1 normalised, but not when they are L2 normalised. This is prevalent across all datasets, even in cases where the correlation between norm and class labels is  $\approx 0$ .

This analysis supports our decision from §2 to use a different noising function to remove information from the norm container, as only the vectors with fully ablated norms have an  $\approx 0$  correlation with both the L1 and L2 norms.