

# Assessing the Limits of the Distributional Hypothesis in Semantic Spaces: Trait-based Relational Knowledge and the Impact of Co-occurrences

Mark Anderson

PIN Caerdydd

Prifysgol Caerdydd

AndersonM8@caerdydd.ac.uk

Jose Camacho-Collados

Cardiff NLP

Cardiff University

CamachoColladosJ@cardiff.ac.uk

## Abstract

The increase in performance in NLP due to the prevalence of distributional models and deep learning has brought with it a reciprocal decrease in interpretability. This has spurred a focus on *what* neural networks *learn* about natural language with less of a focus on *how*. Some work has focused on the data used to develop data-driven models, but typically this line of work aims to highlight issues with the data, e.g. highlighting and offsetting harmful biases. This work contributes to the relatively untrodden path of what is required in data for models to capture meaningful representations of natural language. This entails evaluating how well English and Spanish semantic spaces capture a particular type of relational knowledge, namely the traits associated with concepts (e.g. *bananas-yellow*), and exploring the role of co-occurrences in this context.

## 1 Introduction

Vector space models have been the main driving force behind progress in NLP. Most work in this area, either in the form of static or contextualised embeddings, has been based on co-occurrence statistics and largely driven by the distributional hypothesis (Harris, 1954; Firth, 1957). This has also resulted in these representations seemingly capturing certain relational knowledge, such as word analogies (Mikolov et al., 2013b; Gittens et al., 2017). In this context, Chiang et al. (2020) found that the ability of word embeddings to evaluate analogies was not greatly impaired by removing co-occurrences related to relational pairs. This suggests there are limits to how the distributional hypothesis impacts the encoding of relational knowledge. We extend this line of work by focusing on the relational knowledge of concepts and traits. We also creep beyond English by translating concept and traits used in one of our datasets into Spanish.

**Contributions:** (1) We show that there is no impact on the ability of semantic spaces to predict

whether a pair of embeddings corresponds to a trait-concept pair or to predict what traits a given concept has when removing co-occurrences of concepts and traits. (2) We developed a freely available dataset that can be used for further trait-based relational knowledge analyses for English and Spanish.<sup>1</sup>

## 2 Related work

**What models learn** Evaluation of *neural* semantic spaces has focused on what knowledge they capture with a slew of work showing that some knowledge of analogies can be seen by applying simple transformations (Mikolov et al., 2013b; Levy and Goldberg, 2014; Arora et al., 2016; Paperno and Baroni, 2016; Gittens et al., 2017; Ethayarajh et al., 2019). Others have investigated what syntactic information neural semantic spaces seem to capture with most showing that they do capture something deeper than surface patterns (Linzen et al., 2016; Gulordava et al., 2018; Giulianelli et al., 2018). However, they fail to exhaustively capture syntactic phenomena and specifically have been shown to struggle with polarity (Futrell et al., 2018; Jumelet and Hupkes, 2018) and certain *filler-gap* dependencies (Wilcox et al., 2018; Chowdhury and Zamparelli, 2018). Pretrained language models (PLMs) have been found to capture varying degrees of syntactic information (Peters et al., 2018; Tenney et al., 2019; Goldberg, 2019; Clark et al., 2019), however, they have also been shown to struggle to predict the grammaticality of sentences (Marvin and Linzen, 2018; Warstadt et al., 2019) and seem to depend on fragile heuristics rather than anything deeper (McCoy et al., 2019).

**Relational knowledge** More specifically with respect to relational knowledge and semantic spaces, for some time now work has shown that semantic

<sup>1</sup><https://github.com/cardiffnlp/trait-concept-datasets>

spaces could encode certain relational knowledge, e.g. knowledge of the relative positioning of geographical locations (Louwerse and Zwaan, 2009). Similarly, Gupta et al. (2015) found that embeddings capture something of relational knowledge associated with countries and cities, e.g. how countries related to one another with respect to GDP. Rubinstein et al. (2015) found that word embeddings captured some taxonomic relational knowledge but fared less well with respect to trait-based relational knowledge. Often analogy completion tasks are used to investigate what sort of relational knowledge a semantic space has captured with early work showing that simple linear transformations were enough to highlight analogies (Mikolov et al., 2013a; Vylomova et al., 2016). This method has drawn some criticism and has been challenged as a robust means of evaluating what relational knowledge models capture (Drozd et al., 2016; Gladkova et al., 2016; Schluter, 2018; Bouraoui et al., 2018). Attempts to evaluate what PLMs capture of relational knowledge have also been made, highlighting that these larger, more data-hungry models capture some but not all relational knowledge (Forbes et al., 2019; Bouraoui et al., 2020).

**Patterns in data** However, all the work cited above focuses work focuses on *what* models learn about relational knowledge and not *how*, or rather what are the salient signals in the data used in these techniques that manifest in relational knowledge. Some work has been done in this direction, with Pardos and Nam (2020) showing co-occurrences are not necessary in their distributional model of courses to predict similar or related courses. Chiang et al. (2020) evaluated this finding in neural semantic spaces, finding that the ability of a semantic space to complete analogies isn't impacted when removing co-occurrences

It is important to understand what aspects of the data result in what models learn because without this semblance of interpretability, problematic biases can creep in, e.g. gender biases in Word2Vec (Bolukbasi et al., 2016) or in BERT (Bhardwaj et al., 2021). Attempts have been made to mitigate certain biases in contextualised word embeddings (Kaneko and Bollegala, 2021), but in order to do so, the biases have to be known. Also, Shwartz and Choi (2020) discuss the issue of reporting bias in the data typically used in NLP, where rarer occurrences are more likely to be explicitly mentioned than common ones which results in models that can

generalise about under-reported phenomena but not temper the over-reported information. Therefore it is necessary to understand the nature of the data and how it impacts what models capture and how.

In this work, we aim to expand on the work of Chiang et al. (2020) in two main ways. First, we do not use analogies and analogy completion to evaluate the impact co-occurrences of concept-traits has on relational knowledge developed in neural semantic spaces, but instead use a dataset of different trait-based relations (e.g. *is-colour*, *has-component*) derived from the MCRAE and NORMS feature datasets. This allows us to more directly evaluate the ability of models to predict relational knowledge by casting the evaluation as a simple classification task (both in a multi class and binary class setting). And second, we extend the analysis by looking at Spanish data as well to evaluate whether the results extend beyond English.

### 3 Methodology

The methodology follows five sequential steps: the development of datasets that include concepts and their traits (Section 3.1); the selection and processing of large general-domain corpora (Section 3.2); the transformation of the selected corpora based on the concept-trait datasets to test our hypothesis (Section 3.3); training of word embeddings on the original and adapted corpora (Section 3.4); and finally the evaluation of the embeddings based on the trait-based datasets (Section 3.5).

#### 3.1 Datasets

The datasets were based on the MCRAE features dataset (McRae et al., 2005). This is a collection of semantics features associated with a large set of concepts (541) generated from features given by human participants. A secondary trait-based dataset was also collated for English based on the NORMS dataset (Devereux et al., 2014). This is developed in the same way as MCRAE and is partially an extension of that dataset with 638 concepts. We wanted to avoid value judgements (such as *is-feminine*) and to collate more trait-based relations, that is pairs of words related by an inherent attribute of a concept.

**MCRAE-EN** The first step in developing the datasets used in this work was to collate certain features into subsets of similar traits. This was done in a partially manual way by splitting data into 5 subsets. Each feature in MCRAE has the number

	trait type	$N_C$	$N_T$	Traits
<b>M CRAE-EN</b>	colour	148	7	green (32), brown (32), black (24), white (21), red (16), yellow (13), orange (10)
	components	110	6	handle (39), legs (19), wheels (14), leaves (14), seeds (13), doors (11)
	materials	144	4	metal (79), wood (43), cotton (11), leather (11)
	size & shape	234	4	small (83), large (70), long (44), round (37)
	tactile	117	7	heavy (21), soft (19), furry (18), sharp (17), hard (16), juicy (16), slimy (10)
<b>NORMS</b>	colour	133 (78)	5	green (35), brown (32), white (30), black (22), yellow (14)
	components	35 (26)	2	handle (25), sugar (10)
	materials	94 (62)	5	metal (46), wood (16), water (11), paper (11), bones (10)
	size & shape	242 (138)	4	small (109), large (73), long (31), round (29)
	tactile	106 (70)	6	heavy (28), sharp (26), liquid (14), light (13), juicy (13), soft (12)
<b>M CRAE-ES</b>	colour	140	7	verde (31), marrón (31), blanco (21), negro (20), rojo (16), amarillo (12), naranja (9)
	components	100	6	mango (33), piernas (18), ruedas (14), hojas (14), semillas (11), puertas (10)
	materials	131	4	métal (72), madera (38), algodón (11), cuero (10)
	size & shape	216	4	pequeño (75), grande (66), largo (41), redondo (34)
	tactile	101	6	pesado (19), suave (19), peludo (17), duro (16), afilado (16), jugoso (14)

Table 1: Dataset statistics:  $N_C$  is the number of concepts,  $N_T$  is the number of unique features, NORMS  $N_C$  includes unique count in parenthesis, and the number in parenthesis for traits is the number of concepts with that trait.

of participants who specified that feature for that concept, so initially a frequency cut of 10 was applied to the features. From this set, we observed a number of similar traits that broadly fit into trait categories. A series of simple heuristics were then applied to extract all potential concept-feature pairs for each subset. For some trait types this was trivial with the MCRAE dataset, e.g. colour relations could be found using the feature classification in MCRAE of `visual-colour`. The full details of the heuristics can be seen in Appendix A.

This process resulted in 5 trait-based subsets: **colours**, **components**, **materials**, **size & shape**, and **tactile**. From each subset, we removed duplicates (e.g. ambulance has the features `is-white`, `is-red`, and `is-orange` in the colour subset).<sup>2</sup> And from the remaining concept-feature pairs, we cut on 10+ concepts per trait to ensure a suitable number of instances per target in our evaluation. The resulting statistics associated with this dataset can be seen in the top section of Table 1.

**MCRAE-ES** The set of concepts and trait words that occur across all 5 subsets were manually translated. The translators consisted of one native English speaker with some knowledge of Spanish and one native Spanish speaker who is fluent in English.

As might be expected, issues occurred when undertaking the translation that required judgements to be made. When there was a one to many translation, we used the translation that was *Iberian* if

multiple translations were due to regional variants. Otherwise we chose the most common or most canonical. However, we also chose single word alternatives to avoid multiword concepts when this wouldn't have resulted in using an obscure word. We also made some choices to avoid having duplicate/competing concepts, i.e. *boat* was translated as *barca* and *ship* as *barco*. Further, we tried to match the intended use in English, i.e. translated *sledgehammer* to *almádena* rather than more generic term in Spanish *mazo* as heavy metal version is more standard in English. Otherwise we tried to use more generic options. A variety of resources were used to aid this including bilingual dictionaries, Wikipedia, and RAE (Real Academia Española). Despite our best efforts to maintain as many concept-trait pairs as possible, certain concepts just don't work in Spanish, typically many to one translations, e.g. *dove* translates to *paloma* which also means normal mangy pigeons. A more common issue was the tendency to use multi-word expressions in Spanish for certain concepts, such as *goldfish* (*pez dorado*) and *escalator* (*escalera mecánica*) with no single-word alternatives. The statistics resulting to the trait subsets for MCRAE-ES are shown in the bottom section of Table 1.

**NORMS-EN** To make our experiments more robust, we also used the NORMS dataset. In order to use this dataset, we manually classified features in this dataset based on the subset from our MCRAE trait dataset. First, we cut the features in NORMS that occurred less than 10 times and then took the set of remaining features and classified them as one of the five subsets and then automatically cast each

<sup>2</sup>A multi-label version of these subsets are included at <https://github.com/cardiffnlp/trait-concept-datasets> for MCRAE-EN and NORMS-EN.

	Corpus	Sentences	Tokens
English	UMBC	135M	3.4B
	Wiki	114M	2.5B
	Wee-Wiki	71M	1.6B
Spanish	ES1B	62M	1.4B
	Wiki	28M	0.6B
	Wee-Wiki	19M	0.4B

Table 2: Basic statistics of corpora used.

concept-trait pair into their respective subset. We manually checked to see if any features not used had been erroneously omitted due to annotation issues and folded those features into the relative subsets. This entailed adding *is-liquid* and *is-furry* to the tactile subset after some consideration (with *is-furry* subsequently being removed due to the minimum frequency cut after removing duplicates). The resulting subsets had duplicate concepts removed and then a minimum frequency cut on the remaining features of 10. The statistics of the resulting subsets can be seen in the middle section of Table 1 with the number of new unique concepts added to each subset shown in parenthesis in the concept count ( $N_C$ ) column.

### 3.2 Corpora


For the statistics of the corpora used see Table 2.

**UMBC** The University of Maryland, Baltimore County (UMBC) webbase corpus is the resulting collection of paragraphs from a webcrawl in 2007 over millions of webpages (Han et al., 2013).

**ES1B** The Spanish Billion Words Corpus (ES1B) is a collection of unannotated sentences taken from the web which span difference sources from Europarl to books. It also include data from a Wikipedia dump from 2015, so has some crossover with the Spanish Wikipedia corpus (Cardellino, 2019).

**Wiki** We used English Wikipedia dump from 1st October 2021 and Spanish Wikipedia dump from 1st January 2022. They were extracted and cleaned using the WikiExtractor tool from Attardi (2015). This left document ID HTML tags in the data which we removed with a simple heuristic.

**Wee-Wiki** Similar to the standard pre-processing of the Wikipedia data, but we also cut articles with very little views as these tend to be stub articles and automatically generated articles. The idea behind this is to cultivate a *cleaner* and more natural version of the data. We used Wikipedia’s official



Original text: *Cerró la puerta del granero*  
English: *She/he closed the barn door*

Figure 1: *granero* (highlighted in red) is a concept in MCRAE-ES with a component trait of *puerta* (highlighted in blue). In the example here they are linked by an *nmod* edge (highlighted in blue). For the syntactic removal method this sentence would be removed.

viewing statistics for 1st December 2021.<sup>3</sup> Articles with less than 10 views were removed.

### 3.3 Removing co-occurrences

We used 3 methods to remove co-occurrences with different levels of granularity to find co-occurrences. The first step in the process was to segment the corpora by sentence and to lemmatise the tokens. This was done using the spaCy library and the corresponding pre-trained models for English and Spanish (Montani et al., 2022). We used lemmas to handle gender of adjectives and nouns in Spanish and for plural forms in both languages. The segmented version of each corpus was then split into two separate corpora with 80% of the sentences in the first, which were used as the standard corpora in our experiments, and with 20%, which were used as reserves for replacing sentence with co-occurrences when creating input data without co-occurrences. When an instance was removed based on the criteria specified below, a random sentence was selected from the reserves, so as to balance the total number of sentences in each set.<sup>4</sup> The resulting number of instances removed is shown in Table 3 (English) and in Table 4 (Spanish).

**Sentence** The simplest method used was to merely remove any sentence where a concept and its corresponding trait was observed. The lemmatised version of the data was used to search for co-occurrences to be more thorough, especially with respect to the Spanish data. This entails using the lemmatised version of the concepts and traits to

<sup>3</sup><https://dumps.wikimedia.org/other/pageviews/2021/2021-12/>

<sup>4</sup>Chiang et al. (2020) observed only a small difference when using this methodology and when using one where instances were replaced with sentences containing the relative concepts (and as is shown in §4 this holds for our work).



		UMBC			Wiki			Wee-Wiki		
		instances removed			instances removed			instances removed		
trait type		sentence	window	syntactic	sentence	window	syntactic	sentence	window	syntactic
MCRAE	colour	76,800	70,159	8,974	105,614	97,728	13,397	70,194	64,594	9,083
	components	33,284	23,347	9,745	22,307	15,500	5,915	15,553	10,987	4,544
	material	28,061	19,171	6,030	29,695	20,477	5,771	21,239	14,669	4,431
	size & shape	104,478	68,697	18,213	131,165	88,453	26,612	90,280	60,516	17,452
	tactile	18,881	13,845	4,632	14,437	10,658	3,657	11,413	8,529	2,981
NORMS	colour	25,106	18,737	7,422	26,378	19,824	8,360	19,561	14,777	6,581
	components	5,270	3,793	1,291	4,463	3,110	1,005	3,637	2,483	766
	material	51,898	34,484	12,150	30,916	20,441	7,338	21,051	13,823	4,694
	size & shape	105,895	68,162	18,372	117,210	79,041	22,812	83,329	55,933	15,814
	tactile	17,965	13,040	4,264	13,683	10,156	3,533	11,048	8,307	2,929

Table 3: Total instances removed and replaced for English Corpora (UMBC, Wiki, Wee-Wiki) for each dataset (MCRAE and NORMS) by trait type and removal method (sentence, window, and syntactic as described in §3.3).

		ES1B			Wiki			Wee-Wiki		
		instances removed			instances removed			instances removed		
trait type		sentence	window	syntactic	sentence	window	syntactic	sentence	window	syntactic
MCRAE	colour	31,267	25,121	208	19,424	15,804	2,729	12,473	10,129	1,836
	components	11,855	7,680	1,551	6,628	4,048	1,873	4,318	2,716	1,317
	material	8,473	6,087	1,344	6,704	4,698	2,200	4,353	3,091	1,501
	size & shape	34,416	19,276	248	23,224	13,513	4,001	15,584	9,157	2,798
	tactile	3,508	2,404	185	2,459	1,743	782	1,787	1,291	584

Table 4: Total instances removed and replaced for each Spanish Corpora (ES1B, Wiki, Wee-Wiki) for the MCRAE dataset broken down by trait type and removal method (sentence, window, and syntactic as described in §3.3).

match them in the lemmatised instances in the data. This was done independently for each trait type.

**Window** The second method used removed instances when the concept and its relative trait occurred within a given window, again using lemmatised forms. The window size used was 10 to match the size used during the training of the embeddings.

**Syntactic** Finally, used the Stanza library and the corresponding pre-trained models available for English and Spanish to parse the instances where a concept and its relative trait occurred (Qi et al., 2020). If an edge between the concept and the trait was predicted after finding a co-occurrence using the lemmas, this was removed, otherwise the instance was left. This method tests whether co-occurrences which are syntactically related are more impactful than haphazard co-occurrences. An example is shown in Figure 1.

### 3.4 Word embeddings

The models used to evaluate the impact of co-occurrences were trained using the Gensim library (Řehůřek and Sojka, 2010). We used CBOW Word2Vec embedding models (Mikolov et al., 2013a) as they are quicker to train than skip-gram models which was paramount considering the num-

ber of models that were required. Further, Chiang et al. (2020) found no significant differences between CBOW and Skip-gram models with respect to the differences observed in analogy completion between models trained with and without co-occurrences. We used the default hyperparameters in Gensim except for embedding size which was set to 300 and window size which was set to 10, i.e. the same settings from Chiang et al. (2020). For each trait-type and for each corpus a model was trained on the data containing co-occurrences (**with** or **w/** in tables) and the data not containing co-occurrences (**without** or **w/o** in tables). We trained multiple models for the data including co-occurrences — once per trait type — giving us a robust measurement of those models’ performance. This means that results for each **with** for each trait type across the extraction methods are trained on the same data and are reported to show the variation seen training models on the same data.<sup>5</sup>

### 3.5 Classifiers

Trait-based relational knowledge was evaluated by casting it as a classification problem.

<sup>5</sup>Variation could also be due to slightly different datasets if **without** data doesn’t contain any occurrences of a concept.

		UMBC					Wiki					Wee-Wiki							
MCRAE	trait type	sentence		window		syntactic		sentence		window		syntactic		sentence		window		syntactic	
		w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o
	colour	0.35	<b>0.35</b>	0.34	0.34	<b>0.36</b>	0.35	<b>0.38</b>	0.36	<b>0.38</b>	0.30	<b>0.41</b>	0.36	<b>0.39</b>	0.38	<b>0.39</b>	0.38	<b>0.41</b>	0.35
	components	<b>0.82</b>	0.81	<b>0.81</b>	0.80	<b>0.82</b>	0.81	0.78	<b>0.80</b>	0.75	<b>0.77</b>	<b>0.79</b>	0.76	<b>0.75</b>	0.74	0.75	<b>0.79</b>	0.77	0.77
	materials	0.65	<b>0.69</b>	<b>0.67</b>	0.65	0.67	<b>0.68</b>	<b>0.68</b>	0.67	0.65	<b>0.69</b>	0.65	<b>0.67</b>	<b>0.71</b>	0.65	<b>0.67</b>	0.66	0.67	<b>0.68</b>
	size & shape	<b>0.57</b>	0.53	0.55	<b>0.58</b>	0.54	<b>0.58</b>	<b>0.60</b>	0.58	<b>0.58</b>	0.56	<b>0.61</b>	<b>0.58</b>	0.56	<b>0.59</b>	0.56	<b>0.57</b>	0.56	
	tactile	0.61	<b>0.62</b>	<b>0.64</b>	0.60	<b>0.65</b>	0.64	0.54	<b>0.55</b>	0.56	<b>0.59</b>	<b>0.58</b>	0.55	0.50	<b>0.51</b>	0.50	<b>0.54</b>	0.51	<b>0.55</b>
NORMS	colour	<b>0.40</b>	0.38	0.41	0.41	0.38	<b>0.40</b>	0.39	0.39	<b>0.41</b>	0.39	<b>0.44</b>	0.40	<b>0.43</b>	0.39	0.37	<b>0.39</b>	0.37	<b>0.41</b>
	components	0.89	0.89	0.89	0.89	0.89	0.89	0.91	0.91	0.89	<b>0.91</b>	0.91	0.91	0.89	<b>0.91</b>	0.89	<b>0.91</b>	0.94	0.94
	materials	<b>0.88</b>	0.87	<b>0.87</b>	0.85	0.87	<b>0.88</b>	<b>0.86</b>	0.84	0.85	0.85	0.86	0.86	0.84	0.84	<b>0.83</b>	0.82	<b>0.86</b>	0.82
	size & shape	<b>0.59</b>	0.57	0.58	<b>0.60</b>	<b>0.61</b>	0.58	0.59	0.59	<b>0.62</b>	0.57	0.59	<b>0.61</b>	<b>0.62</b>	0.59	<b>0.58</b>	0.55	<b>0.62</b>	0.57
	tactile	0.69	<b>0.72</b>	<b>0.68</b>	0.66	<b>0.70</b>	0.66	0.61	<b>0.65</b>	0.63	0.63	0.65	<b>0.67</b>	0.60	<b>0.61</b>	0.61	0.61	0.63	0.63

Table 5: Multi-class SVM results for English corpora and datasets by trait type and extraction method for models trained on data with (w/) and without (w/o) co-occurrences. Average accuracy across 3-fold cross validation is reported with best performing model between paired w/ and w/o models highlighted in bold.

		ES1B					Wiki					Wee-Wiki							
MCRAE	trait type	sentence		window		syntactic		sentence		window		syntactic		sentence		window		syntactic	
		w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o
	colour	0.29	<b>0.30</b>	<b>0.33</b>	0.31	<b>0.33</b>	0.30	<b>0.32</b>	0.29	<b>0.34</b>	0.32	<b>0.35</b>	0.33	0.31	<b>0.32</b>	0.30	<b>0.31</b>	<b>0.31</b>	0.29
	components	<b>0.77</b>	0.71	<b>0.81</b>	0.77	<b>0.74</b>	0.73	0.71	<b>0.75</b>	0.70	<b>0.75</b>	0.72	<b>0.74</b>	<b>0.73</b>	0.72	<b>0.71</b>	0.66	0.71	0.71
	materials	0.63	<b>0.67</b>	<b>0.67</b>	0.65	0.66	<b>0.67</b>	0.63	<b>0.64</b>	<b>0.70</b>	0.63	0.61	<b>0.66</b>	0.59	0.59	0.59	<b>0.61</b>	<b>0.63</b>	0.59
	size & shape	0.50	<b>0.52</b>	0.48	<b>0.53</b>	<b>0.46</b>	0.45	<b>0.52</b>	0.49	0.49	0.49	<b>0.50</b>	0.47	<b>0.48</b>	<b>0.49</b>	0.48	0.46	<b>0.53</b>	
	tactile	0.54	<b>0.58</b>	<b>0.55</b>	0.53	<b>0.55</b>	0.53	<b>0.60</b>	0.58	0.60	<b>0.62</b>	<b>0.60</b>	0.59	<b>0.51</b>	0.50	<b>0.52</b>	0.51	<b>0.49</b>	0.48

Table 6: Multi-class SVM results for Spanish corpora and datasets by trait type and extraction method for models trained on data with (w/) and without (w/o) co-occurrences. Average accuracy across 3-fold cross validation is reported with best performing model between paired w/ and w/o models highlighted in bold.

**Multi-class** First we used a multi-class evaluation. Using the datasets described in Section 3.1, given a concept (e.g. *banana*), the task consisted of selecting the most appropriate trait for a given trait type (e.g. *yellow* in the colour dataset). We used a support vector machine (SVM) as our classifier from the Scikit-learn library (Pedregosa et al., 2011) with the word embeddings learned in the previous step as the only input. For each model we used 3-fold cross-validation and report the mean score across the splits.<sup>6</sup> For each pair of models (i.e. with and without co-occurrences for a given trait-type and for a given corpus), we checked to see if concepts appeared in both semantic spaces. When a concept was missing in one or both, it was removed from the dataset for both, such that the comparison of results is robust between the two models we are interested in comparing, however, this was not common. It brought up an issue with *orange* and

*naranja*, namely that it occurs as a concept and as trait, so that in our extraction method for sentence and window occurrences of these are always removed from the corpora and so were removed from the evaluation datasets.

**Binary** We also use binary classification by exploiting the earlier findings suggesting that differences between embeddings can be used as a proxy to capture semantic relations (Mikolov et al., 2013b; Vylomova et al., 2016). Again, we used SVM models, but this time the input features were the differences between concepts and their respective traits (i.e.  $e_c - e_t$ , where  $e_c$  is the concept embedding and  $e_t$  is the trait embedding) and the model predicted whether a pair was related or not. This required developing negative samples. This was done by randomly selecting words from the vocab space of the union of vocabs between each pair of model (i.e. with and without co-occurrences for a given trait type and a given corpus). These words then underwent a modicum of a control check by using lexical databases: WordNet (Fellbaum, 2000) for English and the Multilingual Central Repository version 3.0 for Spanish (Gonzalez-Agirre et al.,

<sup>6</sup>The full results for each model can be found at <https://github.com/cardiffnlp/trait-relations-and-co-occurrences>, including the number of concepts and features used for each model’s evaluation and the standard deviations which are all very small.

2012) via the Natural Language Toolkit (Bird et al., 2009). Once a word was randomly selected from the vocab space (excluding the concepts in the given dataset), the respective lexical database was checked to see if it contained the word and if so whether the synonyms associated with it were at least sometimes nouns (that is the synonym set of nouns contained at least one item). This was so that the selected word could in theory be something akin to a concept and not just gobbledygook. This procedure was done so the number of concepts in the negative sample set matched the number in the positive sample set (which had instances removed that didn't appear in one or both of the paired models similar to the multi-class setup). Then each randomly extracted negative *concept* was ascribed a trait from the given trait space. Similar to the multi-class SVM setup, 3-fold cross-validation was used and the mean score across the splits is reported.<sup>7</sup>

## 4 Results

**Multi-class results** The results for the multi-class experiments can be seen in Table 5 for the English corpora and in Table 6 for the Spanish corpora. The highest performing model for each pair of models, i.e. with (**w/**) and without (**w/o**) co-occurrences is highlighted in bold for clarity. Across the board, it is clear that there is no consistent pattern as to whether a model trained with co-occurrences outperforms a model trained without them or vice versa. This holds for all three co-occurrence extraction techniques, for all trait types, for all datasets, and for all corpora across both languages. This is similar to the findings of Chiang et al. (2020) where little effect was observed on analogy completion whether co-occurrences were included or not, however, a systemic decrease was observed in that context despite it being small. While there are some differences between some models, the differences that would be required to make claims of one model being superior to another are much larger than what are observed here as the experimental setup isn't robust enough to verify that a difference of 0.01-0.02 is significant or not. A visualisation of the differences between each corresponding with and without model for MCRAE-EN by trait type can be seen in Figure 2 (equivalent visualisations

<sup>7</sup>Full results for the binary classifier can be found at <https://github.com/cardiffnlp/trait-relations-and-co-occurrences>, including the number of instances for each model and the standard deviations.




Figure 2: Distributions of delta accuracy ( $\Delta\text{Acc}$ ) for corresponding pairs for each trait type in MCRAE-EN.

for NORMS-EN and MCRAE-ES are shown in Figure 4 and 5, respectively, in Appendix B). Figure 2 does highlight a slight difference with respect to colour traits, where a modest increase in performance is seen on average when training the models with co-occurrences, however, this isn't consisted across corpora and datasets as this increase is not observed in Figures 4 and 5 in Appendix B.

**Binary results** The results from the binary classification experiments substantiate these findings. They can be seen in Table 7 for English and in Table 8 for Spanish. Again, no pattern emerges across the different experimental dimensions that would suggest the removal of co-occurrences has impacted a model's ability to predict whether a pair is related or not. The overall high performance on the binary classification experiment for both English and Spanish suggests these models manage to encode meaningful information about these trait relations. But how this emerges is not clear. The simplest explanation is that suitably accurate representations are learnt due to the amount of data, but it could be for any number of other reasons not investigated here.




Figure 3: Distributions of delta accuracy ( $\Delta\text{Acc}$ ) for pairs for each extraction method in MCRAE-EN.

		UMBC					Wiki					Wee-Wiki							
trait type		sentence		window		syntactic		sentence		window		syntactic		sentence		window		syntactic	
		w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o
MCRAE	colour	<b>0.90</b>	0.88	<b>0.88</b>	0.86	0.86	0.86	0.84	<b>0.85</b>	<b>0.88</b>	0.86	0.88	0.88	0.89	<b>0.90</b>	0.87	0.87	0.87	0.87
	components	<b>0.90</b>	0.88	0.90	0.90	0.90	0.90	<b>0.88</b>	0.87	<b>0.87</b>	<b>0.88</b>	0.89	<b>0.90</b>	0.86	0.86	0.92	0.92	0.89	0.89
	materials	<b>0.93</b>	0.92	0.92	0.92	0.90	0.90	<b>0.90</b>	0.88	0.88	<b>0.89</b>	0.89	0.89	<b>0.86</b>	0.85	0.88	<b>0.89</b>	<b>0.90</b>	0.89
	size & shape	0.88	0.88	0.85	0.85	0.85	0.85	0.86	0.86	0.83	0.83	0.84	0.84	<b>0.88</b>	0.87	0.87	<b>0.88</b>	0.86	0.86
	tactile	0.89	<b>0.90</b>	0.88	0.88	0.86	<b>0.88</b>	0.88	0.88	0.82	0.82	0.87	<b>0.88</b>	<b>0.84</b>	0.82	<b>0.84</b>	0.81	<b>0.82</b>	0.81
NORMS	colour	0.86	0.86	<b>0.84</b>	0.83	0.83	0.83	0.86	0.86	<b>0.85</b>	0.83	<b>0.84</b>	0.83	0.84	0.84	<b>0.87</b>	0.86	<b>0.86</b>	0.85
	components	0.80	<b>0.82</b>	<b>0.87</b>	0.77	0.80	0.80	0.90	0.90	<b>0.87</b>	0.78	<b>0.86</b>	0.84	0.86	<b>0.87</b>	<b>0.93</b>	0.90	<b>0.84</b>	0.83
	materials	0.84	<b>0.85</b>	<b>0.88</b>	0.86	0.88	<b>0.91</b>	<b>0.89</b>	0.87	0.86	<b>0.89</b>	0.85	<b>0.87</b>	0.85	<b>0.88</b>	0.85	<b>0.88</b>	0.90	<b>0.91</b>
	size & shape	0.84	<b>0.87</b>	0.88	<b>0.89</b>	<b>0.87</b>	0.86	<b>0.88</b>	0.86	0.84	0.84	0.85	0.85	0.88	0.88	<b>0.88</b>	0.87	<b>0.87</b>	0.85
	tactile	<b>0.87</b>	0.84	<b>0.86</b>	0.84	0.84	<b>0.86</b>	<b>0.83</b>	0.82	0.82	<b>0.84</b>	<b>0.86</b>	0.84	<b>0.78</b>	0.77	0.80	<b>0.81</b>	0.86	0.86

Table 7: Binary SVM results for English corpora and datasets by trait type and extraction method for models trained on data with (w/) and without (w/o) co-occurrences. Average accuracy across 3-fold cross validation is reported with best performing model between paired w/ and w/o models highlighted in bold.

		ES1B					Wiki					Wee-Wiki							
trait type		sentence		window		syntactic		sentence		window		syntactic		sentence		window		syntactic	
		w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o
MCRAE	colour	0.81	<b>0.82</b>	<b>0.85</b>	0.83	<b>0.81</b>	0.80	<b>0.84</b>	0.81	<b>0.81</b>	0.78	<b>0.87</b>	0.84	<b>0.83</b>	0.81	<b>0.83</b>	0.80	0.79	<b>0.82</b>
	components	<b>0.88</b>	0.87	<b>0.83</b>	0.80	<b>0.81</b>	0.80	0.86	<b>0.89</b>	0.78	0.78	<b>0.79</b>	<b>0.82</b>	0.77	<b>0.78</b>	0.82	<b>0.84</b>	<b>0.76</b>	0.74
	materials	0.81	<b>0.82</b>	0.86	0.86	0.84	0.84	<b>0.78</b>	0.76	0.75	0.75	<b>0.70</b>	0.67	0.75	<b>0.80</b>	<b>0.75</b>	0.74	0.74	0.74
	size & shape	<b>0.82</b>	0.81	<b>0.75</b>	0.73	<b>0.76</b>	0.74	0.79	<b>0.80</b>	<b>0.76</b>	0.75	<b>0.79</b>	0.78	0.79	0.79	0.75	<b>0.78</b>	0.82	<b>0.83</b>
	tactile	<b>0.72</b>	0.71	0.75	<b>0.79</b>	<b>0.81</b>	0.78	<b>0.74</b>	0.73	0.71	<b>0.72</b>	0.74	<b>0.75</b>	<b>0.78</b>	0.72	<b>0.77</b>	0.75	0.78	<b>0.80</b>

Table 8: Binary SVM results for Spanish corpora and datasets by trait type and extraction method for models trained on data with (w/) and without (w/o) co-occurrences. Average accuracy across 3-fold cross validation is reported with best performing model between paired w/ and w/o models highlighted in bold.

## 5 Discussion

The results highlight some tentatively interesting patterns with respect to trait types. In both English and Spanish, models perform consistently well on component traits, although for NORMS this turned out to be only over 2 traits, effectively casting it as a binary classification. Materials is the next consistently highest performing trait type across corpora and language with size & shape and tactile not far behind for English, but with a bigger gap in Spanish. The performance on colour traits is low across all settings and languages. This doesn't appear to be based on the size of the trait subset, e.g. the component subset is one of the smaller sets, yet has high performance and the performance of the other trait types don't vary with respect to the number of instance and unique features.

The number of removed sentences, as shown in Tables 3 and 4, gives a vague indication of the occurrences of the concepts in the dataset and the occurrence of their traits with colour sentence removals being the second highest for MCRAE-EN across all three English corpora, the third highest for NORMS-EN, and the highest for MCRAE-ES

across all Spanish corpora. These rankings are consistent across extraction methods. Therefore, it is unlikely that the embeddings for the colours and the corresponding concepts (often concepts that occur in the other datasets) are somehow low quality due to low occurrences of these words. More likely is that the colour relation is more difficult than the other trait types as the other types are more tangible and more specific. Although this doesn't necessarily hold for size & shape traits, specifically sizes which tend to be relative, e.g. in MCRAE a *plane* can be *large* (which it is, relative to most things) but so too can a *bathtub* (which it is, relative to a mouse or other such timorous beasties, but not relative to a house). However, size & shape is consistently one of the traits that models perform worst on especially with NORMS-EN and MCRAE-ES.

As a final note, the different extraction methods yield no differences when compared to one another. This can be observed clearly in Figure 3 in the main text and Figure 6 in Appendix B. While the number of extracted instances using the syntactically related co-occurrences is very low and so difficult to draw any major conclusions, the num-



ber of sentence-based and window-based instances removed are quite high and are similar in magnitude. From this, we can deduce that the proximity of the words also doesn't have a major impact on the ability of a semantic space to encode relational knowledge. It could still be the case that if the data used to train models contained more syntactically related concept-trait pairs, they would encode *more* relational knowledge, but it is clear that their absence doesn't result in the models losing what relational knowledge they can capture. Many questions remain on how these distributional models encode relational knowledge. We have merely presented results which *do not* support the hypothesis that direct co-occurrence are the major signal for this process as related to trait-based relational knowledge.

### Language models and wider impact of findings.

Whether the results observed here for static embeddings would hold for PLMS isn't a given. While they are still based on the same distributional hypothesis and adopt statistical methods to encode salient features of language, they could potentially be more sensitive to the loss of co-occurrences in the training data. But this is an open research question that requires specific experimentation which has its own difficulties, i.e. prompting language models often includes lexical *clues* which cloud our ability to say with any great certainty if they have captured some phenomenon or not, see [Kassner and Schütze \(2020\)](#) for sensitivity of PLMs to mispriming).

The results do suggest that merely increasing the amount of data used likely won't result in any major improvements in the ability of models to encode relational knowledge or commonsense knowledge more generally, which is attested to by recent work in [Li et al. \(2021\)](#). Potentially, we need to look to more complex methods to augment NLP systems with commonsense knowledge potentially using multimodal systems, e.g. language models trained with visual cues as was done in [Paik et al. \(2021\)](#) to offset reporting bias with respect to colours. Alternatively, we can focus on the linguistic input and consider how to add stronger signals in the data used to train NLP systems.

## 6 Conclusion

We have contributed to the emerging interest in how neural semantic models encode linguistic information, focusing on trait-based relational knowl-

edge. We have extended findings which showed that co-occurrences of relational pairs didn't have a major impact on a model's ability to encode knowledge of analogies by complementing this analysis with an evaluation of trait-based relational knowledge. We extended the analysis to include different extraction methods to evaluate whether a more fine-grained approach would highlight any differences in performance and found that this is not the case. The work presented here also expands beyond English and includes results in Spanish which follow the same trend. Finally, we have cultivated a set of datasets for different trait types in both English and Spanish (based on MCRAE and NORMS) which are available at <https://github.com/cardiffnlp/trait-concept-datasets>.

## Acknowledgements

Mark and Jose are supported by a UKRI Future Leaders Fellowship (MR/T042001/1).

## References

- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. A latent variable model approach to pmi-based word embeddings. *Transactions of the Association for Computational Linguistics*, 4:385–399.
- Giuseppe Attardi. 2015. Wikiextractor. <https://github.com/attardi/wikiextractor>.
- Rishabh Bhardwaj, Navonil Majumder, and Soujanya Poria. 2021. Investigating gender bias in bert. *Cogn. Comput.*, 13:1008–1018.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in Neural Information Processing Systems*, 29.
- Zied Bouraoui, Jose Camacho-Collados, and Steven Schockaert. 2020. Inducing relational knowledge from bert. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(5):7456–7463.
- Zied Bouraoui, Shoaib Jameel, and Steven Schockaert. 2018. [Relation induction in word embeddings revisited](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1627–1637, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

- Cristian Cardellino. 2019. [Spanish Billion Words Corpus and Embeddings](#).
- Hsiao-Yu Chiang, Jose Camacho-Collados, and Zachary Pardo. 2020. [Understanding the source of semantic regularities in word embeddings](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 119–131, Online. Association for Computational Linguistics.
- Shammur Absar Chowdhury and Roberto Zamparelli. 2018. Rnn simulations of grammaticality judgments on long-distance dependencies. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 133–144.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? An analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, Florence, Italy. Association for Computational Linguistics.
- Barry Devereux, Lorraine K. Tyler, Jeroen Geertzen, and Billi Randall. 2014. The centre for speech, language and the brain (cslb) concept property norms. *Behavior Research Methods*, 46:1119 – 1127.
- Aleksandr Drozd, Anna Gladkova, and Satoshi Matsuo. 2016. [Word embeddings, analogies, and machine learning: Beyond king - man + woman = queen](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3519–3530, Osaka, Japan. The COLING 2016 Organizing Committee.
- Kawin Ethayarajh, David Duvenaud, and Graeme Hirst. 2019. [Towards understanding linear word analogies](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3253–3262, Florence, Italy. Association for Computational Linguistics.
- Christiane D. Fellbaum. 2000. Wordnet : an electronic lexical database. *Language*, 76:706.
- John R Firth. 1957. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*.
- Maxwell Forbes, Ari Holtzman, and Yejin Choi. 2019. Do neural language representations learn physical commonsense? In *CogSci*.
- Richard Futrell, Ethan Wilcox, Takashi Morita, and Roger Levy. 2018. Rnns as psycholinguistic subjects: Syntactic state and grammatical dependency. *arXiv preprint arXiv:1809.01329*.
- Alex Gittens, Dimitris Achlioptas, and Michael W Mahoney. 2017. Skip-gram- zipf+ uniform= vector additivity. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 69–76.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Anna Gladkova, Aleksandr Drozd, and Satoshi Matsuo. 2016. Analogy-based detection of morphological and semantic relations with word embeddings: what works and what doesn’t. In *Proceedings of the Student Research Workshop at NAACL*, pages 8–15.
- Yoav Goldberg. 2019. Assessing BERT’s syntactic abilities. *arXiv preprint arXiv:1901.05287*.
- Aitor Gonzalez-Agirre, Egoitz Laparra, and German Rigau. 2012. Multilingual central repository version 3.0: upgrading a very large lexical knowledge base. In *Proceedings of the 6th Global WordNet Conference (GWC 2012)*, Matsue.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. [Colorless green recurrent networks dream hierarchically](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1195–1205, New Orleans, Louisiana. Association for Computational Linguistics.
- Abhijeet Gupta, Gemma Boleda, Marco Baroni, and Sebastian Padó. 2015. [Distributional vectors encode referential attributes](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 12–21, Lisbon, Portugal. Association for Computational Linguistics.
- Lushan Han, Abhay L. Kashyap, Tim Finin, James Mayfield, and Jonathan Weese. 2013. [UMBC\\_EBIQUITY-CORE: Semantic textual similarity systems](#). In *Second Joint Conference on Lexical and Computational Semantics (\*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 44–52, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Zellig S. Harris. 1954. Distributional structure. *Word*, 10:146–162.
- Jaap Jumelet and Dieuwke Hupkes. 2018. Do language models understand anything? On the ability of LSTMs to understand negative polarity items. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 222–231.
- Masahiro Kaneko and Danushka Bollegala. 2021. Debiasing pre-trained contextualised embeddings. *ArXiv*, abs/2101.09523.

- Nora Kassner and Hinrich Schütze. 2020. **Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.
- Omer Levy and Yoav Goldberg. 2014. **Linguistic regularities in sparse and explicit word representations.** In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 171–180, Ann Arbor, Michigan. Association for Computational Linguistics.
- Xiang Lorraine Li, Adhiguna Kuncoro, Cyprien de Masson d’Autume, Phil Blunsom, and Aida Nematzadeh. 2021. **Do language models learn commonsense knowledge?** *ArXiv*.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. **Assessing the ability of LSTMs to learn syntax-sensitive dependencies.** *Transactions of the Association for Computational Linguistics*, 4(1):521–535.
- Max M. Louwerse and Rolf A. Zwaan. 2009. **Language encodes geographical information.** *Cognitive science*, 33 1:51–73.
- Rebecca Marvin and Tal Linzen. 2018. **Targeted syntactic evaluation of language models.** In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019. **Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448.
- Ken McRae, George S. Cree, Mark S. Seidenberg, and Chris McNorgan. 2005. **Semantic feature production norms for a large set of living and nonliving things.** *Behavior Research Methods*, 37:547–559.
- Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013a. **Efficient estimation of word representations in vector space.** In *ICLR*.
- Tomas Mikolov, Wen tau Yih, and Geoffrey Zweig. 2013b. **Linguistic regularities in continuous space word representations.** In *NAACL*.
- Ines Montani, Matthew Honnibal, Sofie Van Landeghem, Adriane Boyd, et al. 2022. **explosion/spaCy: v3.3.0: Improved speed, new trainable lemmatizer, and pipelines for Finnish, Korean and Swedish.**
- Cory Paik, Stéphane Aroca-Ouellette, Alessandro Roncone, and Katharina Kann. 2021. **The World of an Octopus: How Reporting Bias Influences a Language Model’s Perception of Color.** In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 823–835, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Denis Paperno and Marco Baroni. 2016. **When the whole is less than the sum of its parts: How composition affects pmi values in distributional semantic vectors.** *Computational Linguistics*, 42:345–350.
- Zachary A. Pardos and Andrew Joo Hun Nam. 2020. **A university map of course knowledge.** *PLoS ONE*, 15.
- 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.
- Matthew E. Peters, Mark Neumann, Luke Zettlemoyer, and Wen tau Yih. 2018. **Dissecting contextual word embeddings: Architecture and representation.** In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. **Stanza: A Python natural language processing toolkit for many human languages.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Radim Řehůřek and Petr Sojka. 2010. **Software Framework for Topic Modelling with Large Corpora.** In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta. ELRA. <http://is.muni.cz/publication/884893/en>.
- Dana Rubinstein, Effi Levi, Roy Schwartz, and Ari Rapoport. 2015. **How well do distributional models capture different types of semantic knowledge?** In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 726–730, Beijing, China. Association for Computational Linguistics.
- Natalie Schluter. 2018. **The word analogy testing caveat.** In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 242–246.
- Vered Shwartz and Yejin Choi. 2020. **Do neural language models overcome reporting bias?** In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6863–6870, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. **BERT rediscovers the classical NLP pipeline.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601.

Ekaterina Vylomova, Laura Rimell, Trevor Cohn, and Timothy Baldwin. 2016. [Take and took, gaggle and goose, book and read: Evaluating the utility of vector differences for lexical relation learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1671–1682, Berlin, Germany. Association for Computational Linguistics.

Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.

Ethan Wilcox, Roger Levy, Takashi Morita, and Richard Futrell. 2018. What do rnn language models learn about filler-gap dependencies? In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 211–221.

## A MCRAE-EN trait subset extraction heuristics

Here we describe the full heuristics used to develop the trait-based subsets from MCRAE used in our experiments

Some traits were trivial to extract. Colour relations were the simplest as they could be found using the feature classification in MCRAE of visual-colour. Component relations were shortlisted cutting on the WB feature classification (this is simply a classification of trait types where W and B refer to the practitioners who classified the concept-features pairs in unpublished work) in MCRAE using `external_component` and `internal_component` and then by extracting features beginning with `has_`. Similarly for material relations, the WB classification of `made_of` was used. Some manual corrections were applied to the components to extend the number of instances in the dataset and to make certain traits fit our experimental setup better. This involved casting features such as `has-4-legs` and `has-4-wheels` as simply `has-legs` and `has-wheels`, respectively. The feature `made-of-material` was cut from the material subset, the feature `has-an-inside` from the components subset, and the features `is-colourful` and `different-colours` were removed from the colour subset.

We then looked at the WB label `external_surface_property` (excluding features that fit into the colour, concept, or material subset) as this fit our desired trait-based feature space. The majority of concepts in this subset

tended to have features relating to their shape or to their size, so we opted to use this pair (size & shape) as another subset. This required manually removing features that didn't fit this trait-type, e.g. `is-smelly`, `is-shiny`, and so on. In this process, a final possible subset of tactile-based traits became apparent which was cut using the BR feature classification (this is simply a classification of trait types from different practitioners than WB) `tactile` and then manually removing certain features which were more value judgements than traits, such as `is-comfortable` or `is-warm`.

## B Visualisations of NORMS-EN and MCRAE-ES results




Figure 4: Distributions of delta accuracy ( $\Delta\text{Acc}$ ) for corresponding pairs for each trait type in NORMS-EN.




Figure 5: Distributions of delta accuracy ( $\Delta\text{Acc}$ ) for corresponding pairs for each trait type in MCRAE-ES.






Figure 6: Distributions of delta accuracy ( $\Delta\text{Acc}$ ) for corresponding pairs for extraction method in NORMS-EN.




Figure 7: Distributions of delta accuracy ( $\Delta\text{Acc}$ ) for corresponding pairs for each extraction method in MCRAE-ES.