

Distributional modeling on a diet: One-shot word learning from text only

Su Wang[♠] Stephen Roller[♣] Katrin Erk[♠]

[♠]Department of Linguistics, [♣]Department of Computer Science
The University of Texas at Austin

shrekwang@utexas.edu, roller@cs.utexas.edu, katrin.erk@mail.utexas.edu

Abstract

We test whether distributional models can do one-shot learning of definitional properties from text only. Using Bayesian models, we find that first learning overarching structure in the known data, regularities in textual contexts and in properties, helps one-shot learning, and that individual context items can be highly informative. Our experiments show that our model can learn properties from a single exposure when given an informative utterance.

1 Introduction

When humans encounter an unknown word in text, even with a single instance, they can often infer approximately what it means, as in this example from Lazaridou et al. (2014):

We found a cute, hairy *wampimuk* sleeping behind the tree.

People who hear this sentence typically guess that a *wampimuk* is an animal, or even that it is a mammal. Distributional models, which describe the meaning of a word in terms of its observed contexts (Turney and Pantel, 2010), have been suggested as a model for how humans learn word meanings (Landauer and Dumais, 1997). However, distributional models typically need hundreds of instances of a word to derive a high-quality representation for it, while humans can often infer a passable meaning approximation from one sentence only (as in the above example). This phenomenon is known as *fast mapping* (Carey and Bartlett, 1978). Our primary modeling objective in this paper is to explore a plausible model for fast-mapping learning from textual context.

While there is preliminary evidence that fast mapping can be modeled distributionally (Lazaridou et al., 2016), it is unclear what enables it.

How do humans infer word meanings from so little data? This question has been studied for *grounded* word learning, when the learner perceives an object in non-linguistic context that corresponds to the unknown word. The literature emphasizes the importance of learning general knowledge or overarching structure, which we define as the information that is learned by accumulation across concepts (e.g. regularities in property co-occurrence), across all concepts (Kemp et al., 2007). In grounded word learning, overarching structure that has been proposed includes knowledge about which properties. For example knowledge about which properties are most important to object naming (Smith et al., 2002; Colunga and Smith, 2005), or a taxonomy of concepts (Xu and Tenenbaum, 2007).

In this paper we study models for fast mapping in word learning¹ from textual context alone, using probabilistic distributional models. Our task differs from the grounded case in that we do not perceive any object labeled by the unknown word. In that context, learning *word meaning* means learning the associated definitional properties and their weights (see Section 3). For the sake of interpretability, we focus on learning definitional properties. We ask what kinds of overarching structure in distributional contexts and in properties will be helpful for one-shot word learning.

We focus on learning from syntactic context. Distributional representations of syntactic context are directly interpretable as selectional constraints, which in manually created resources are typically characterized through high-level taxonomy classes (Kipper-Schuler, 2005; Fillmore et al., 2003). So they should provide good evidence for

¹In this paper, we interchangeably use the terms *unknown word* and *unknown concept*, as we learn properties, and properties belong to concepts rather than words, and we learn them from text, where we observe words rather than concepts.

the meaning of role fillers. Also, it has been shown that selectional constraints can be learned distributionally (Erk et al., 2010; Ó Séaghdha and Korhonen, 2014; Ritter et al., 2010). However, our point will not be that syntax is needed for fast word learning, but that it helps to observe overarching structure, with syntactic context providing a clear test bed.

We test two types of overarching structure for their usefulness in fast mapping. First, we hypothesize that it is helpful to learn about commonalities among context items, which enables mapping from contexts to properties. For example the syntactic contexts *eat-dobj* and *cook-dobj* should prefer similar targets: things that are cooked are also things that are eaten (Hypothesis **H1**).

The second hypothesis is that it will be useful to learn co-occurrence patterns between properties. That is, we hypothesize that in learning an entity is a *mammal*, we may also infer it is *four-legged* (Hypothesis **H2**).

We do not intend to make *strong* cognitive claims, for which additional experimentation will be in order, and we leave this for future work. This work sets its goal on building a plausible computational model that models human fast-mapping in learning (i) well from limited grounded data, (ii) effectively from only one instance.

2 Background

Fast mapping and textual context. Fast mapping (Carey and Bartlett, 1978) is the human ability to construct provisional word meaning representations after one or few exposures. An important reason for why humans can do fast mapping is that they acquire overarching structure that constrains learning (Smith et al., 2002; Colunga and Smith, 2005; Kemp et al., 2007; Xu and Tenenbaum, 2007; Maas and Kemp, 2009). In this paper, we ask what forms of overarching structure will be useful for text-based word learning.

Lazaridou et al. (2014) consider fast mapping for grounded word learning, mapping image data to distributional representations, which is in a way the mirror image of our task. Lazaridou et al. (2016) were the first to explore fast mapping for text-based word learning, using an extension to word2vec with both textual and visual features. However, they model the unknown word simply by averaging the vectors of known words in the sentence, and do not explore what types of knowl-

edge enable fast mapping.

Definitional properties. Feature norms are definitional properties collected from human participants. Feature norm datasets are available from McRae et al. (2005) and Vigliocco et al. (2004). In this paper we use feature norms as our target representations of word meaning. There are several recent approaches that learn to map distributional representations to feature norms (Johns and Jones, 2012; Rubinstein et al., 2015; Făgărășan et al., 2015; Herbelot and Vecchi, 2015a). We also map distributional information to feature norms, but we do it based on a single textual instance (one-shot learning).

In the current paper we use the **Quantified McRae (QMR)** dataset (Herbelot and Vecchi, 2015b), which extends the McRae et al. (2005) feature norms by ratings on the proportion of category members that have a property, and the **Animal** dataset (Herbelot, 2013), which is smaller but has the same shape. For example, *most* alligators are dangerous. The quantifiers are given probabilistic interpretations, so if *most* alligators are dangerous, the probability for a random alligator to be dangerous would be 0.95. This makes this dataset a good fit for our probabilistic distributional model. We discuss QMR and the Animal data further in Section 4.

Bayesian models in lexical semantics. We use Bayesian models for the sake of interpretability and because the existing definitional property datasets are small. The Bayesian models in lexical semantics that are most related to our approach are Dinu and Lapata (2010), who represent word meanings as distributions over latent topics that approximate senses, and Andrews et al. (2009) and Roller and Schulte im Walde (2013), who use multi-modal extensions of Latent Dirichlet Allocation (LDA) models (Blei et al., 2003) to represent co-occurrences of textual context and definitional features. Ó Séaghdha (2010) and Ritter et al. (2010) use Bayesian approaches to model selectional preferences.

3 Models

In this section we develop a series of models to test our hypothesis that acquiring general knowledge is helpful to word learning, in particular knowledge about similarities between context items (H1) and co-occurrences between properties (H2). The count-based model will implement neither hypoth-

esis, while the bimodal topic model will implement both. To test the hypotheses separately, we employ two clustering approaches via Bernoulli Mixtures, which we use as extensions to the count-based model and bimodal topic model.

3.1 The Count-based Model

Independent Bernoulli condition. Let Q be a set of definitional properties, C a set of concepts that the learner knows about, and V a vocabulary of context items. For most of our models, context items $w \in V$ will be predicate-role pairs such as *eat-dobj*. The task is determine properties that apply to an unknown concept $u \notin C$. Any concept $c \in C$ is associated with a vector \mathbf{c}_{Ind} (where “Ind” stands for “independent Bernoulli probabilities”) of $|Q|$ probabilities, where the i -th entry of \mathbf{c}_{Ind} is the probability that an instance of concept c would have property q_i . These probabilities are independent Bernoulli probabilities. For instance, $\mathbf{alligator}_{\text{Ind}}$ would have an entry of 0.95 for *dangerous*. An *instance* $\underline{c} \in \{0, 1\}^{|Q|}$ of a concept $c \in C$ is a vector of zeros and ones drawn from \mathbf{c}_{Ind} , where an entry of 1 at position i means that this instance has the property q_i .

The model proceeds in two steps. First it learns property probabilities for context items $w \in V$. The model observes instances \underline{c} occurring textually with context item w , and learns property probabilities for w , where the probability that w has for a property q indicates the probability that w would appear as a context item with an instance that has property q . In the second step the model uses the acquired context item representations to learn property probabilities for an unknown concept u . When u appears with w , the context item w “imagines” an instance (samples it from its property probabilities), and uses this instance to update the property probabilities of u . Instead of making point estimates, the model represents its uncertainty about the probability of a property through a Beta distribution, a distribution over Bernoulli probabilities. As a Beta distribution is characterized by two parameters α and β , we associate each context item $w \in V$ with vectors $\mathbf{w}^\alpha \in \mathbb{R}^{|Q|}$ and $\mathbf{w}^\beta \in \mathbb{R}^{|Q|}$, where the i -th α and β values are the parameters of the Beta distribution for property q_i . When an instance \underline{c} is observed with context item w , we do a Bayesian update on w simply as

$$\begin{aligned} \mathbf{w}^\alpha &= \mathbf{w}^\alpha + \underline{c} \\ \mathbf{w}^\beta &= \mathbf{w}^\beta + (1 - \underline{c}) \end{aligned} \quad (1)$$

because the Beta distribution is the conjugate prior of the Bernoulli. To draw an instance from w , we draw it from the predictive posterior probabilities of its Beta distributions, $\mathbf{w}_{\text{Ind}} = \mathbf{w}^\alpha / (\mathbf{w}^\alpha + \mathbf{w}^\beta)$.

Likewise, we associate an unknown concept u with vectors \mathbf{u}^α and \mathbf{u}^β . When the model observes u in the context of w , it draws an instance from \mathbf{w}_{Ind} , and performs a Bayesian update as in (1) on the vectors associated with u . After training, the property probabilities for u are again the posterior predictive probabilities $\mathbf{u}_{\text{Ind}} = \mathbf{u}^\alpha / (\mathbf{u}^\alpha + \mathbf{u}^\beta)$. The model can be used for multi-shot learning and one-shot learning in the same way.

Multinomial condition. We also test a multinomial variant of the count-based model, for greater comparability with the LDA model below. Here, the concept representation \mathbf{c}_{Mult} is a multinomial distribution over the properties in Q . (That is, all the properties compete in this model.) An instance of concept c is now a single property, drawn from c ’s multinomial. The representation of a context item w , and also the representation of the unknown concept u , is a Dirichlet distribution with $|Q|$ parameters. Bayesian update of the representation of w based on an occurrence with c , and likewise Bayesian update of the representation of u based on an occurrence with w , is straightforward again, as the Dirichlet distribution is the conjugate prior of the multinomial.

The two count-based models do not implement either of our two hypotheses. They compute separate selectional constraints for each context item, and do not attend to co-occurrences between properties. In the experiments below, the count-based models will be listed as **Count Independent** and **Count Multinomial**.

3.2 The Bimodal Topic Model

We use an extension of LDA (Blei et al., 2003) to implement our hypotheses on the usefulness of overarching structure, both commonalities in selectional constraints across predicates, and co-occurrence of properties across concepts. In particular, we build on Andrews et al. (2009) in using a *bimodal topic model*, in which a single topic simultaneously generates both a context item and a property. We further build on Dinu and Lapata (2010) in having a “pseudo-document” for each concept c to represent its observed occurrences. In our case, this pseudo-document contains pairs of a context item $w \in V$ and a property $q \in Q$,




Figure 1: Plate diagram for the Bimodal Topic Model (bi-TM)

meaning that w has been observed to occur with an instance of c that had q .

The generative story is as follows. For each known concept c , draw a multinomial θ_c over topics. For each topic z , draw a multinomial ϕ_z over context items $w \in V$, and a multinomial ψ_z over properties $q \in Q$. To generate an entry for c 's pseudo-document, draw a topic $z \sim \text{Mult}(\theta_c)$. Then, from z , simultaneously draw a context item from ϕ_z and a property from ψ_z . Figure 1 shows the plate diagram for this model.

To infer properties for an unknown concept u , we create a pseudo-document for u containing just the observed context items, no properties, as those are not observed. From this pseudo-document d_u we infer the topic distribution θ_u . Then the probability of a property q given d_u is

$$P(q|d_u) = \sum_z P(z|\theta_u)P(q|\psi_z) \quad (2)$$

For the one-shot condition, where we only observe a single context item w with u , this simplifies to

$$P(q|w) = \sum_z P(z|w)P(q|\psi_z) \quad (3)$$

We refer to this model as **bi-TM** below. The topics of this model implement our hypothesis H1 by grouping context items that tend to occur with the same concepts and the same properties. The topics also implement our hypothesis H2 by grouping properties that tend to occur with the same concepts and the same context items. By using multinomials ψ_z it makes the simplifying assumption that all properties compete, like the Count Multinomial model above.

3.3 Bernoulli Mixtures

With the Count models, we investigate word learning without any overarching structures. With the bi-TMs, we investigate word learning with both types of overarching structures at once. In order to evaluate each of the two hypotheses separately, we use clustering with Bernoulli Mixture models of either the context items or the properties.

A Bernoulli Mixture model (Juan and Vidal, 2004) assumes that a population of m -dimensional binary vectors \mathbf{x} has been generated by a set of mixture components K , each of which is a vector of m Bernoulli probabilities:

$$p(\mathbf{x}) = \sum_{k=1}^{|K|} p(k)p(\mathbf{x}|k) \quad (4)$$

A Bernoulli Mixture can represent co-occurrence patterns between the m random variables it models without assuming competition between them.

To test the effect of modeling *cross-predicate selectional constraints*, we estimate a Bernoulli Mixture model from n instances $\underline{\mathbf{w}}$ for each $w \in V$, sampled from \mathbf{w}_{Ind} (which is learned as in the Count Independent model). Given a Bernoulli Mixture model of $|K|$ components, we then assign each context item w to its closest mixture component as follows. Say the instances of w used to estimate the Bernoulli Mixture were $\{\underline{\mathbf{w}}_1, \dots, \underline{\mathbf{w}}_n\}$, then we assign w to the component

$$k_w = \text{argmax}_k \sum_{j=1}^n p(k|\underline{\mathbf{w}}_j) \quad (5)$$

We then re-train the representations of context items in the Count Multinomial condition, treating each occurrence of c with context w as an occurrence of c with k_w . This yields a Count Multinomial model called **Count BernMix H1**.

To test the effect of modeling *property co-occurrences*, we estimate a $|K|$ -component Bernoulli Mixture model from n instances of each known concept $c \in C$, sampled from \mathbf{c}_{Ind} . We then represent each concept c by a vector \mathbf{c}_{Mult} , a multinomial with $|K|$ parameters, as follows. Say the instances of c used to estimate the Bernoulli Mixture were $\{\underline{\mathbf{c}}_1, \dots, \underline{\mathbf{c}}_n\}$, then the k -th entry in \mathbf{c}_{Mult} is the average probability, over all $\underline{\mathbf{c}}_i$, of being generated by component k :

$$\mathbf{c}_k = \frac{1}{n} \sum_{j=1}^n p(k|\underline{\mathbf{c}}_j) \quad (6)$$

This can be used as a Count Multinomial model where the entries in c_{Mult} stand for Bernoulli Mixture components rather than individual properties. We refer to it as **Count BernMix H2**.²

Finally, we extend the bi-TM with the H2 Bernoulli Mixture in the same way as a Count Multinomial model, and list this extension as **bi-TM BernMix H2**. While the bi-TM already implements both H1 and H2, its assumption of competition between all properties is simplistic, and bi-TM BernMix H2 tests whether lifting this assumption will yield a better model. We do not extend the bi-TM with the H1 Bernoulli Mixture, as the assumption of competition between context items that the bi-TM makes is appropriate.

4 Data and Experimental Setup

Definitional properties. As we use probabilistic models, we need probabilities of properties applying to concept instances. So the QMR dataset (Herbelot and Vecchi, 2015b) is ideally suited. QMR has 532 concrete noun concepts, each associated with a set of quantified properties. The quantifiers have been given probabilistic interpretations, mapping *all*→1, *most*→0.95, *some*→0.35, *few*→0.05, *none*→0.³ Each concept/property pair was judged by 3 raters. We choose the majority rating when it exists, and otherwise the minimum proposed rating. To address sparseness, especially for the one-shot learning setting, we omit properties that are named for fewer than 5 concepts. This leaves us with 503 concepts and 220 properties. We intentionally choose this small dataset: One of our main objectives is to explore the possibility of learning effectively from very limited training data. In addition, while the feature norm dataset is small, our distributional dataset (the BNC, see below) is not. The latter essentially serves as a pivot for us to propagate the knowledge from the feature norm data to the wider semantic space.

It is a problem of both the original McRae et al. (2005) data and QMR that if a property is not named by participants, it is not listed, even if it applies. For example, the property *four-legged*

²We use the H2 Bernoulli Mixture as a soft clustering because it is straightforward to do this through concept representations. For the H1 mixture, we did not see an obvious soft clustering, so we use it as a hard clustering.

³The dataset also contains *KIND* properties that do not have probabilistic interpretations. Following Herbelot and Vecchi (2015a) we omit these properties.

is missing for *alligator* in QMR. So we additionally use the **Animal** dataset of Herbelot (2013), where every property has a rating for every concept. The dataset comprises 72 animal concepts with quantification information for 54 properties.

Distributional data. We use the British National Corpus (BNC) (The BNC Consortium, 2007), with dependency parses from Spacy.⁴ As context items, we use pairs $\langle \text{pred}, \text{dep} \rangle$ of predicates *pred* that are content words (nouns, verbs, adjectives, adverbs) but not stopwords, where a concept from the respective dataset (QMR, Animal) is a dependency child of *pred* via *dep*. In total we obtain a vocabulary of 500 QMR concepts and 72 Animal concepts that appear in the BNC, and 29,124 context items. We refer to this syntactic context as **Syn**. For comparison, we also use a baseline model with a bag-of-words (**BOW**) context window of 2 or 5 words, with stopwords removed.

Models. We test our probabilistic models as defined in the previous section. While our focus is on one-shot learning, we also evaluate a multi-shot setting where we learn from the whole BNC, as a sanity check on our models. (We do not test our models in an incremental learning setting that adds one occurrence at a time. While this is possible in principle, the computational cost is prohibitive for the bi-TM.) We compare to the Partial Least Squares (**PLS**) model of Herbelot and Vecchi (2015a)⁵ to see whether our models perform at state of the art levels. We also compare to a baseline that always predicts the probability of a property to be its relative frequency in the set C of known concepts (**Baseline**).

We can directly use the property probabilities in QMR and the Animal data as concept representations c_{Ind} for the Count Independent model. For the Count Multinomial model, we never explicitly compute c_{Mult} . To sample from it, we first sample an instance $\underline{c} \in \{0, 1\}^{|\mathcal{Q}|}$ from the independent Bernoulli vector of c , c_{Ind} . From the properties that apply to \underline{c} , we sample one (with equal probabilities) as the observed property. All priors for the count-based models (Beta priors or Dirichlet priors, respectively) are set to 1.

For the bi-TM, a pseudo-document for a known

⁴<https://spacy.io>

⁵Herbelot and Vecchi (2015a) is the only directly relevant previous work on the subject. Further, to the best of our knowledge, for one-shot property learning from text (only), our work has been the first attempt.

Models		QMR		Animal Syn
		BOW5	Syn	
Baseline		0.12	0.16	0.63
PLS		0.24	0.35	0.71
Count	Mult.	0.13	0.25	0.64
	Ind.	0.11	0.23	0.64
	BernMix H1	0.11	0.17	0.65
	BernMix H2	0.10	0.18	0.63
bi-TM	plain	0.23	0.36	0.80
	BernMix H2	0.20	0.34	0.81

Table 1: MAP scores, multi-shot learning on the QMR and Animal datasets

concept c is generated as follows: Given an occurrence of known concept c with context item w in the BNC, we sample a property q from c (in the same way as for the Count Multinomial model), and add $\langle w, q \rangle$ to the pseudo-document for c . For training the bi-TM, we use collapsed Gibbs sampling (Steyvers and Griffiths, 2007) with 500 iterations for burn-in. The Dirichlet priors are uniformly set to 0.1 following Roller and Schulte im Walde (2013). We use 50 topics throughout.

For all our models, we report the average performance from 5 runs. For the PLS benchmark, we use 50 components with otherwise default settings, following Herbelot and Vecchi (2015a).

Evaluation. We test all models using 5-fold cross validation and report average performance across the 5 folds. We evaluate performance using *Mean Average Precision* (MAP), which tests to what extent a model ranks definitional properties in the same order as the gold data. Assume a system that predicts a ranking of n datapoints, where 1 is the highest-ranked, and assume that each datapoint i has a gold rating of $I(i) \in \{0, 1\}$. This system obtains an Average Precision (AP) of

$$AP = \frac{1}{\sum_{i=1}^n I(i)} \sum_{i=1}^n \text{Prec}_i \cdot I(i)$$

where Prec_i is precision at a cutoff of i . Mean Average Precision is the mean over multiple AP values. In our case, $n = |Q|$, and we compare a model-predicted ranking of property probabilities with a binary gold rating of whether the property applies to any instances of the given concept. For the one-shot evaluation, we make a separate prediction for each occurrence of an unknown concept u in the BNC, and report MAP by averaging over the AP values for all occurrences of u .

5 Results and Discussion

Multi-shot learning. While our focus in this paper is on one-shot learning, we first test all models in a multi-shot setting. The aim is to see how well they perform when given ample amounts of training data, and to be able to compare their performance to an existing multi-shot model (as we will not have any related work to compare to for the one-shot setting.) The results are shown in Table 1, where *Syn* shows results that use syntactic context (encoding selectional constraints) and *BOW5* is a bag-of-words context with a window size of 5. We only compare our models to the baseline and benchmark for now, and do an in-depth comparison of our models when we get to the one-shot task, which is our main focus.

Across all models, the syntactic context outperforms the bag-of-words context. We also tested a bag-of-words context with window size 2 and found it to have a performance halfway between *Syn* and *BOW5* throughout. This confirms our assumption that it is reasonable to focus on syntactic context, and for the rest of this paper, we test models with syntactic context only.

Focusing on *Syn* conditions now, we see that almost all models outperform the property frequency baseline, though the MAP scores for the baseline do not fall far behind those of the weakest count-based models.⁶ The best of our models perform on par with the PLS benchmark of Herbelot and Vecchi (2015a) on QMR, and on the Animal dataset they outperform the benchmark. Comparing the two datasets, we see that all models show better performance on the cleaner (and smaller) *Animal* dataset than on QMR. This is probably because QMR suffers from many false negatives (properties that apply but were not mentioned), while *Animal* does not. The Count Independent model shows similar performance here and throughout all later experiments to the Count Multinomial (even though it matches the construction of the QMR and Animal datasets better), so to avoid clutter we do not report on it further below.

One-shot learning. Table 2 shows the perfor-

⁶This is because MAP gives equal credit for all properties correctly predicted as non-zero. When we evaluate with Generalized Average Precision (GAP) (Kishida, 2005), which takes gold weights into account, the baseline model is roughly 10 points below other models. This indicates our models learn approximate property distributions. We omit GAP scores because they correlate strongly with MAP for non-baseline models.

Models			all	oracle top20	AvgCos top20
QMR	Count	Mult.	0.16	0.37	0.28
		BernMix H1	0.14	0.33	0.21
		BernMix H2	0.15	0.31	0.22
	bi-TM	plain	0.21	0.47	0.35
		BernMix H2	0.18	0.45	0.34
Animal	Count	Mult.	0.58	0.77	0.61
		BernMix H1	0.60	0.80	0.57
		BernMix H2	0.59	0.81	0.59
	bi-TM	plain	0.64	0.88	0.63
		BernMix H2	0.65	0.89	0.66

Table 2: MAP scores, one-shot learning on the QMR and Animal datasets

mance of our models on the one-shot learning task. We cannot evaluate the benchmark PLS as it is not suitable for one-shot learning. The baseline is the same as in Table 1. The numbers shown are Average Precision (AP) values for learning from a single occurrence. Column *all* averages over all occurrences of a target in the BNC (using only context items that appeared at least 5 times in the BNC), and column *oracle top-20* averages over the 20 context items that have the highest AP for the given target. As can be seen, AP varies widely across sentences: When we average over all occurrences of a target in the BNC, performance is close to baseline level.⁷ But the most *informative* instances yield excellent information about an unknown concept, and lead to MAP values that are much higher than those achieved in multi-shot learning (Table 1). We explore this more below.

Comparing our models, we see that the bi-TM does much better throughout than any of the count-based models. Since the bi-TM model implements both cross-predicate selectional constraints (H1) and property co-occurrence (H2), we find both of our hypotheses confirmed by these results. The Bernoulli mixtures improved performance on the Animal dataset, with no clear pattern of which one improved performance more. On QMR, adding a Bernoulli mixture model harms performance across both the count-based and bi-TM models. We suspect that this is because of the false negative entries in QMR; an inspection of Bernoulli mixture H2 components supports this intuition, as the QMR ones were found to be of poorer quality than those for the Animal data.

Comparing Tables 1 and 2 we see that they show

⁷Context items with few occurrences in the corpus perform considerably worse than baseline, as their property distributions are dominated by the small number of concepts with which they appear.

Count Mult.	clothing, made_of_metal, different_colours, an_animal, is_long
bi-TM	clothing, made_of_material, has_sleeves, different_colours, worn_by_women
bi-TM one-shot	clothing, is_long, made_of_material, different_colours, has_sleeves

Table 3: QMR: top 5 properties of *gown*. Top 2 entries: multi-shot. Last entry: one-shot, context *undo-dobj*

Top	<i>undo-dobj</i> (0.70), <i>nylon-nmod</i> (0.66), <i>pink-amod</i> (0.65), <i>retie-dobj</i> (0.64), <i>silk-amod</i> (0.64)
Bottom	<i>sport-nsubj</i> (0.01), <i>contemplate-dobj</i> (0.01), <i>comic-amod</i> (0.01), <i>wait-nsubj</i> (0.01), <i>fibrous-amod</i> (0.01)

Table 4: QMR one-shot: AP for top and bottom 5 context items of *gown*

the same patterns of performance: Models that do better on the multi-shot task also do better on the one-shot task. This is encouraging in that it suggests that it should be possible to build incremental models that do well both in a low-data and an abundant-data setting.

Table 3 looks in more detail at what it is that the models are learning by showing the five highest-probability properties they are predicting for the concept *gown*. The top two entries are multi-shot models, the third shows the one-shot result from the context item with the highest AP. The bi-TM results are very good in both the multi-shot and the one-shot setting, giving high probability to some quite specific properties like *has_sleeves*. The count-based model shows a clear frequency bias in erroneously giving high probabilities to the two overall most frequent properties, *made_of_metal* and *an_animal*. This is due to the additive nature of the Count model: In updating unknown concepts from context items, frequent properties are more likely to be sampled, and their effect accumulates as the model does not take into account interactions among context items. The bi-TM, which models these interactions, is much more robust to the effect of property frequency.

Informativity. In Table 2 we saw that one-shot performance averaged over all context items in the whole corpus was quite bad, but that good, *informative* context items can yield high-quality property information. Table 4 illustrates this point fur-

	Model		Freq.	Entropy	AvgCos
QMR	Count	Mult.	0.09	-0.12	0.18
	Count	BernMix H1	0.07	-0.10	0.17
	Count	BernMix H2	0.10	-0.09	0.17
	bi-TM	plain	0.15	-0.09	0.41*
	bi-TM	BernMix H2	0.16	-0.10	0.39*
Ani.	bi-TM	plain	0.25	-0.40	0.49*
	bi-TM	BernMix H2	0.23*	-0.37*	0.52*

Table 5: Correlation of informativity with AP, Spearman’s ρ . * and • indicate significance at $p < 0.05$ and $p < 0.1$

ther. For the concept *gown*, it shows the five context items that yielded the highest AP values, at the top *undo-obj*, with an AP as high as 0.7.

This raises the question of whether we can predict the informativity of a context item.⁸ We test three measures of informativity. The first is simply the **frequency** of the context item, with the rationale that more frequent context items should have more stable representations. Our second measure is based on **entropy**. For each context item w , we compute a distribution over properties as in the count-independent model, and measure the entropy of this distribution. If the distribution has few properties account for a majority of the probability mass, then w will have a low entropy, and would be expected to be more informative. Our third measure is based on the same intuition, that items with more “concentrated” selectional constraints should be more informative. If a context item w has been observed to occur with known concepts c_1, \dots, c_n , then this measure is the average cosine (**AvgCos**) of the property distributions (viewed as vectors) of any pair of $c_i, c_j \in \{c_1, \dots, c_n\}$.

We evaluate the three informativity measures using Spearman’s rho to determine the correlation of the informativity of a context item with the AP it produces for each unknown concept. We expect frequency and AvgCos to be positively correlated with AP, and entropy to be negatively correlated with AP. The result is shown in Table 5. Again, all measures work better on the Animal data than on QMR, where they at best approach significance. The correlation is much better on the bi-TM models than on the count-based models, which is probably due to their higher-quality predictions. Overall, AvgCos emerges as the most robust indicator

⁸Lazaridou et al. (2016), who use a bag-of-words context in one-shot experiments, propose an informativity measure based on the number of context that constitute properties. We cannot do that with our syntactic context.

Type	MAP
Function	0.45
Taxonomic	0.62
Visual	0.34
Encyclopaedic	0.35
Perc	0.40

Table 6: QMR, bi-TM, one-shot: MAP by property type over (oracle) top 20 context items

for informativity.⁹ We now test AvgCos, as our best informativity measure, on its ability to select good context items. The last column of Table 2 shows MAP results for the top 20 context items based on their AvgCos values. The results are much below the oracle MAP (unsurprisingly, given the correlations in Table 5), but for QMR they are at the level of the multi-shot results of Table 1, showing that it is possible to some extent to automatically choose informative examples for one-shot learning.

Properties by type. McRae et al. (2005) classify properties based on the brain region taxonomy of Cree and McRae (2003). This enables us to test what types of properties are learned most easily in our fast-mapping setup by computing average AP separately by property type. To combat sparseness, we group property types into five groups, *function* (the function or use of an entity), *taxonomic*, *visual*, *encyclopaedic*, and *other perceptual* (e.g., sound). Intuitively, we would expect our contexts to best reflect *taxonomic* and *function* properties: Predicates that apply to noun target concepts often express functions of those targets, and manually specified selectional constraints are often characterized in terms of taxonomic classes. Table 6 confirms this intuition. Taxonomic properties achieve the highest MAP by a large margin, followed by functional properties. Visual properties score the lowest.

6 Conclusion

We have developed several models for one-shot learning word meanings from single textual contexts. Our models were designed learn word properties using distributional contexts (H1) or about co-occurrences of properties (H2). We find evidence that both kinds of general knowledge are

⁹We also tested a binned variant of the frequency measure, on the intuition that medium-frequency context items should be more informative than either highly frequent or rare ones. However, this measure did not show better performance than the non-binned frequency measure.

helpful, especially when combined (in the bi-TM), or when used on clean property data (in the Animal dataset). We further saw that some contexts are highly informative, and preliminary experiments in informativity measures found that average pairwise similarity of seen role fillers (AvgCos) achieves some success in predicting which contexts are most useful.

In the future, we hope to test with other types of general knowledge, including a taxonomy of known concepts (Xu and Tenenbaum, 2007); wider-coverage property data (Baroni and Lenci, 2010, Type-DM); and alternative modalities (Lazaridou et al., 2016, image features as “properties”). We expect our model will scale to these larger problems easily.

We would also like to explore better informativity measures and improvements for AvgCos. Knowledge about informative examples can be useful in human-in-the-loop settings, for example a user aiming to illustrate classes in an ontology with a few typical corpus examples. We also note that the bi-TM cannot be used in for truly incremental learning, as the cost of global re-computation after each seen example is prohibitive. We would like to explore probabilistic models that support incremental word learning, which would be interesting to integrate with an overall probabilistic model of semantics (Goodman and Lassiter, 2014).

Acknowledgments

This research was supported by the DARPA DEFT program under AFRL grant FA8750-13-2-0026 and by the NSF CAREER grant IIS 0845925. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the view of DARPA, DoD or the US government. We acknowledge the Texas Advanced Computing Center for providing grid resources that contributed to these results.

References

- Mark Andrews, Gabriella Vigliocco, and David Vinson. 2009. Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, 116(3):463–498.
- Marco Baroni and Alexandero Lenci. 2010. Distributional memory: a general framework for corpus-based semantics. *Computational Linguistics*, 36(4):673–721.
- David Blei, Andrew Ng, and Michael Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3(4-5):993–1022.
- Susan Carey and Elsa Bartlett. 1978. Acquiring a single new word. *Papers and Reports on Child Language Development*, 15:17–29.
- Eliana Colunga and Linda B. Smith. 2005. From the lexicon to expectations about kinds: A role for associative learning. *Psychological Review*, 112(2):347–382.
- George S. Cree and Ken McRae. 2003. Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, 132:163–201.
- Georgiana Dinu and Mirella Lapata. 2010. Measuring distributional similarity in context. In *Proceedings of EMNLP*, Cambridge, MA.
- Katrin Erk, Sebastian Padó, and Ulrike Padó. 2010. A flexible, corpus-driven model of regular and inverse selectional preferences. *Computational Linguistics*, 36(4).
- Luana Făgărășan, Eva Maria Vecchi, and Stephen Clark. 2015. From distributional semantics to feature norms: Grounding semantic models in human perceptual data. In *Proceedings of IWCS*, London, Great Britain.
- C. J. Fillmore, C. R. Johnson, and M. Petruck. 2003. Background to FrameNet. *International Journal of Lexicography*, 16:235–250.
- Noah D. Goodman and Daniel Lassiter. 2014. Probabilistic semantics and pragmatics: Uncertainty in language and thought. In Shalom Lappin and Chris Fox, editors, *Handbook of Contemporary Semantics*. Wiley-Blackwell.
- Aurélie Herbelot. 2013. What is in a text, what isn’t and what this has to do with lexical semantics. *Proceedings of IWCS*.
- Aurélie Herbelot and Eva Vecchi. 2015a. Building a shared world:mapping distributional to model-theoretic semantic spaces. In *Proceedings of EMNLP*.
- Aurélie Herbelot and Eva Maria Vecchi. 2015b. Many speakers, many worlds. *Linguistic Issues in Language Technology*, 12(4):1–20.
- Brendan T Johns and Michael N Jones. 2012. Perceptual inference through global lexical similarity. *Topics in Cognitive Science*, 4(1):103–120.

- Alfons Juan and Enrique Vidal. 2004. Bernoulli mixture models for binary images. In *Proceedings of ICPR*.
- Charles Kemp, Amy Perfors, and Joshua B. Tenenbaum. 2007. Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*, 10(3):307–321.
- Karin Kipper-Schuler. 2005. *VerbNet: A broad-coverage, comprehensive verb lexicon*. Ph.D. thesis, Computer and Information Science Dept., University of Pennsylvania, Philadelphia, PA.
- Kazuaki Kishida. 2005. Property of average precision and its generalization: An examination of evaluation indicator for information retrieval experiments. *NII Technical Reports*, 2005(14):1–19.
- Thomas Landauer and Susan Dumais. 1997. A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, pages 211–240.
- Angeliki Lazaridou, Elia Bruni, and Marco Baroni. 2014. Is this a wampimuk? Cross-modal mapping between distributional semantics and the visual world. In *Proceedings of ACL*.
- Angeliki Lazaridou, Marco Marelli, and Marco Baroni. 2016. Multimodal word meaning induction from minimal exposure to natural text. *Cognitive Science*, pages 1–30.
- Andrew L. Maas and Charles Kemp. 2009. One-shot learning with Bayesian networks. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, Amsterdam, The Netherlands.
- 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(4):547–559.
- Diarmuid Ó Séaghdha. 2010. Latent variable models of selectional preference. In *Proceedings of ACL*.
- Diarmuid Ó Séaghdha and Anna Korhonen. 2014. Probabilistic distributional semantics with latent variable models. *Computational Linguistics*, 40(3):587–631.
- Alan Ritter, Mausam, and Oren Etzioni. 2010. A Latent Dirichlet Allocation method for selectional preferences. In *Proceedings of ACL*.
- Stephen Roller and Sabine Schulte im Walde. 2013. A multimodal lda model integrating textual, cognitive and visual modalities. In *Proceedings of EMNLP*.
- Dana Rubinstein, Effi Levi, Roy Schwartz, and Ari Rappoport. 2015. How well do distributional models capture different types of semantic knowledge? In *Proceedings of ACL*, volume 2, pages 726–730.
- Linda B. Smith, Susan S. Jones, Barbara Landau, Lisa Gershkoff-Stowe, and Larissa Samuelson. 2002. Object name learning provides on-the-job training for attention. *Psychological Science*, 13(1):13–19.
- Mark Steyvers and Tom Griffiths. 2007. Probabilistic topic models. In *T. Landauer, D.S. McNamara, S. Dennis, and W. Kintsch, eds., Handbook of Latent Semantic Analysis*.
- The BNC Consortium. 2007. *The British National Corpus, version 3 (BNC XML Edition)*. Oxford University Computing Services, URL: <http://www.natcorp.ox.ac.uk/>.
- Peter Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37:141–188.
- Gabriella Vigliocco, David Vinson, William Lewis, and Merrill Garrett. 2004. Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology*, 48:422–488.
- Fei Xu and Joshua B. Tenenbaum. 2007. Word learning as Bayesian inference. *Psychological Review*, 114(2):245–272.