

Probabilistic Embedding of Knowledge Graphs with Box Lattice Measures

Luke Vilnis* Xiang Li* Shikhar Murty Andrew McCallum

College of Information and Computer Sciences

University of Massachusetts Amherst

{luke, xiangli, smurty, mccallum}@cs.umass.edu

Abstract

Embedding methods which enforce a partial order or lattice structure over the concept space, such as Order Embeddings (OE) (Vendrov et al., 2016), are a natural way to model transitive relational data (e.g. entailment graphs). However, OE learns a deterministic knowledge base, limiting expressiveness of queries and the ability to use uncertainty for both prediction and learning (e.g. learning from expectations). Probabilistic extensions of OE (Lai and Hockenmaier, 2017) have provided the ability to somewhat calibrate these *denotational probabilities* while retaining the consistency and inductive bias of ordered models, but lack the ability to model the negative correlations found in real-world knowledge. In this work we show that a broad class of models that assign probability measures to OE can never capture negative correlation, which motivates our construction of a novel *box lattice* and accompanying probability measure to capture anticorrelation and even disjoint concepts, while still providing the benefits of probabilistic modeling, such as the ability to perform rich joint and conditional queries over arbitrary sets of concepts, and both learning from and predicting calibrated uncertainty. We show improvements over previous approaches in modeling the Flickr and WordNet entailment graphs, and investigate the power of the model.

1 Introduction

Structured embeddings based on regions, densities, and orderings have gained popularity in recent years for their inductive bias towards the essential asymmetries inherent in problems such as image captioning (Vendrov et al., 2016), lexical and textual entailment (Erk, 2009; Vilnis and McCallum, 2015; Lai and Hockenmaier, 2017; Athiwaratkun and Wilson, 2018), and knowledge graph completion and reasoning (He et al., 2015; Nickel and Kiela, 2017; Li et al., 2017).

Models that easily encode asymmetry, and related properties such as transitivity (the two components of commonplace relations such as partially ordered sets and lattices), have great utility in these applications, leaving less to be learned from the data than arbitrary relational models. At their best, they resemble a hybrid between embedding models and structured prediction. As noted by Vendrov et al. (2016) and Li et al. (2017), while the models learn sets of embeddings, these parameters obey rich structural constraints. The entire set can be thought of as one, sometimes provably consistent, structured prediction, such as an ontology in the form of a single directed acyclic graph.

While the structured prediction analogy applies best to Order Embeddings (OE), which embeds consistent partial orders, other region- and density-based representations have been proposed for the express purpose of inducing a bias towards asymmetric relationships. For example, the Gaussian Embedding (GE) model (Vilnis and McCallum, 2015) aims to represent the asymmetry and uncertainty in an object's relations and attributes by means of uncertainty in the representation. However, while the space of representations is a manifold of probability distributions, the model is not truly probabilistic in that it does not model asymmetries and relations in terms of prob-

* Equal contribution.

abilities, but in terms of asymmetric comparison functions such as the originally proposed KL divergence and the recently proposed *thresholded divergences* (Athiwaratkun and Wilson, 2018).

Probabilistic models are especially compelling for modeling ontologies, entailment graphs, and knowledge graphs. Their desirable properties include an ability to remain consistent in the presence of noisy data, suitability towards semi-supervised training using the expectations and uncertain labels present in these large-scale applications, the naturality of representing the inherent uncertainty of knowledge they store, and the ability to answer complex queries involving more than 2 variables. Note that the final one requires a true joint probabilistic model with a tractable inference procedure, not something provided by e.g. matrix factorization.

We take the dual approach to density-based embeddings and model uncertainty about relationships and attributes as explicitly probabilistic, while basing the probability on a latent space of geometric objects that obey natural structural biases for modeling transitive, asymmetric relations. The most similar work are the probabilistic order embeddings (POE) of Lai (Lai and Hockenmaier, 2017), which apply a probability measure to each order embedding’s forward cone (the set of points greater than the embedding in each dimension), assigning a finite and normalized volume to the unbounded space. However, POE suffers severe limitations as a probabilistic model, including an inability to model negative correlations between concepts, which motivates the construction of our box lattice model.

Our model represents objects, concepts, and events as high-dimensional products-of-intervals (*hyperrectangles* or *boxes*), with an event’s unary probability coming from the box volume and joint probabilities coming from overlaps. This contrasts with POE’s approach of defining events as the forward cones of vectors, extending to infinity, integrated under a probability measure that assigns them finite volume.

One desirable property of a structured representation for ordered data, originally noted in (Vendrov et al., 2016) is a “slackness” shared by OE, POE, and our model: when the model predicts an “edge” or lack thereof (i.e. $P(a|b) = 0$ or 1, or a zero constraint violation in the case of OE), being exposed to that fact again will not update

the model. Moreover, there are large degrees of freedom in parameter space that exhibit this slackness, giving the model the ability to embed complex structure with 0 loss when compared to models based on symmetric inner products or distances between embeddings, e.g. bilinear GLMs (Collins et al., 2002), Trans-E (Bordes et al., 2013), and other embedding models which must always be pushing and pulling parameters towards and away from each other.

Our experiments demonstrate the power of our approach to probabilistic ordering-biased relational modeling. First, we investigate an instructive 2-dimensional toy dataset that both demonstrates the way the model self organizes its box event space, and enables sensible answers to queries involving arbitrary numbers of variables, despite being trained on only pairwise data. We achieve a new state of the art in denotational probability modeling on the Flickr entailment dataset (Lai and Hockenmaier, 2017), and a matching state-of-the-art on WordNet hypernymy (Vendrov et al., 2016; Miller, 1995) with the concurrent work on thresholded Gaussian embedding of Athiwaratkun and Wilson (2018), achieving our best results by training on additional co-occurrence expectations aggregated from leaf types.

We find that the strong empirical performance of probabilistic ordering models, and our box lattice model in particular, and their endowment of new forms of training and querying, make them a promising avenue for future research in representing structured knowledge.

2 Related Work

In addition to the related work in structured embeddings mentioned in the introduction, our focus on directed, transitive relational modeling and ontology induction shares much with the rich field of directed graphical models and causal modeling (Pearl, 1988), as well as learning the structure of those models (Heckerman et al., 1995). Work in undirected structure learning such the Graphical Lasso (Friedman et al., 2008) is also relevant due to our desire to learn from pairwise joint/conditional probabilities and moment matrices, which are closely related in the setting of discrete variables.

Especially relevant research in Bayesian networks are applications towards learning taxonomic structure of relational data (Bansal et al.,

2014), although this work is often restricted towards tree-shaped ontologies, which allow efficient inference by Chu-Liu-Edmonds’ algorithm (Chu and Liu, 1995), while we focus on arbitrary DAGs.

As our model is based on populating a latent “event space” into boxes (products of intervals), it is especially reminiscent of the Mondrian process (Roy and Teh, 2009). However, the Mondrian process partitions the space as a high dimensional tree (a non-parametric kd-tree), while our model allows the arbitrary box placement required for DAG structure, and is much more tractable in high dimensions compared to the Mondrian’s Bayesian non-parametric inference.

Embedding applications to relational learning constitute a huge field to which it is impossible to do justice, but one general difference between our approaches is that e.g. a matrix factorization model treats the embeddings as objects to score relation links with, as opposed to POE or our model in which embeddings represent subsets of probabilistic event space which are directly integrated. They are full probabilistic models of the joint set of variables, rather than embedding-based approximations of only low-order joint and conditional probabilities. That is, any set of our parameters can answer any arbitrary probabilistic question (possibly requiring intractable computation), rather than being fixed to modeling only certain subsets of the joint.

Embedding-based learning’s large advantage over the combinatorial structure learning presented by classical PGM approaches is its applicability to large-scale probability distributions containing hundreds of thousands of events or more, as in both our WordNet and Flickr experiments.

3 Background

3.1 Partial Orders and Lattices

A non-strict *partial ordered set* (*poset*) is a set P equipped with a binary relation \preceq such that for all $a, b, c \in P$,

- $a \preceq a$ (reflexivity)
- $a \preceq b \preceq a$ implies $a = b$ (antisymmetry)
- $a \preceq b \preceq c$ implies $a \preceq c$ (transitivity)

This is simply a generalization of a totally ordered set that allows some elements to be incomparable,

and is a good model for the kind of acyclic directed graph data found in knowledge bases.

A *lattice* is a poset where any subset has a unique least upper and greatest lower bound, which will be true of all posets (lattices) considered in this paper. The least upper bound of two elements $a, b \in P$ is called the *join*, denoted $a \vee b$, and the greatest lower bound is called the *meet*, denoted $a \wedge b$.

Additionally, in a *bounded lattice* we have two extra elements, called *top*, denoted \top and *bottom*, denoted \perp , which are respectively the least upper bound and greatest lower bound of the entire space. Using the extended real number line (adding points at infinity), all lattices considered in this paper are bounded lattices.

3.2 Order Embeddings (OE)

Vendrov et al. (2016) introduced a method for embedding partially ordered sets and a task, *partial order completion*, an abstract term for things like hypernym or entailment prediction (learning transitive relations). The goal is to learn a mapping from the partially-ordered data domain to some other partially-ordered space that will enable generalization.

Definition 1. Vendrov et al. (2016)

A function $f : (X, \preceq_X) \rightarrow (Y, \preceq_Y)$ is an order-embedding if for all $u, v \in X$

$$u \preceq_X v \iff f(u) \preceq_Y f(v)$$

They choose Y to be a vector space, and the order \preceq_Y to be based on the *reverse product order* on \mathbb{R}_+^n , which specifies

$$x \preceq y \iff \forall i \in \{1..n\}, x_i \geq y_i$$

so an embedding is below another in the hierarchy if all of the coordinates are larger, and 0 provides a top element.

Although Vendrov et al. (2016) do not explicitly discuss it, their model does not just capture partial orderings, but is a standard construction of a vector (Hilbert) lattice, in which the operations of meet and join can be defined as taking the pointwise maximum and minimum of two vectors, respectively (Zaanen, 1997). This observation is also used in (Li et al., 2017) to generate extra constraints for training order embeddings.

As noted in the original work, these single point embeddings can be thought of as regions, i.e. the

cone extending out from the vector towards infinity. All concepts “entailed” by a given concept must lie in this cone.

This ordering is optimized from examples of ordered elements and negative samples via a max-margin loss.

3.3 Probabilistic Order Embeddings (POE)

Lai and Hockenmaier (2017) built on the “region” idea to derive a probabilistic formulation (which we will refer to as POE) to model entailment probabilities in a consistent, hierarchical way.

Noting that all of OE’s regions obviously have the same infinite area under the standard (Lebesgue) measure of \mathbb{R}_+^n , they propose a probabilistic interpretation where the Bernoulli probability of each concept a or joint set of concepts $\{a, b\}$ with corresponding vectors $\{x, y\}$ is given by its volume under the exponential measure:

$$p(a) = \exp\left(-\sum_i x_i\right) = \int_{z \preceq x} \exp(-\|z\|_1) dz$$

$$p(a, b) = p(x \wedge y) = \exp(-\|\max(x_i, y_i)\|_1)$$

since the meet of two vectors is simply the intersection of their area cones, and replacing sums with ℓ_1 norms for brevity since all coordinates are positive. While having the intuition of measuring the areas of cones, this also automatically gives a valid probability distribution over concepts since this is just the product likelihood under a coordinatewise exponential distribution.

However, they note a deficiency of their model — it can only model positive (Pearson) correlations between concepts (Bernoulli variables).

Consider two Bernoulli variables a and b , whose probabilities correspond to the areas of cones x and y . Recall the Bernoulli covariance formula (we will deal with covariances instead of correlations when convenient, since they always have the same sign):

$$\text{cov}(a, b) = p(a, b) - p(a)p(b) = \exp(-\|\max(x_i, y_i)\|_1) - \exp(-\|x_i + y_i\|_1)$$

Since the sum of two positive vectors can only be greater than the sum of their pointwise maximum, this quantity will always be nonnegative. This has real consequences for probabilistic modeling in KBs: conditioning on more concepts will only make probabilities higher (or unchanged), e.g. $p(\text{dog}|\text{plant}) \geq p(\text{dog})$.

3.4 Probabilistic Asymmetric Transitive Relations

Probabilistic models have pleasing consistency properties for modeling asymmetric transitive relations, in particular compared to density-based embeddings — a pairwise conditional probability table can almost always (in the technical sense) be asymmetricized to produce a DAG by simply taking an edge if $P(a|b) > P(b|a)$. A matrix of pairwise Gaussian KL divergences cannot be consistently asymmetricized in this manner. These claims are proven in Appendix C. While a high $P(a|b)$ does not always indicate an edge in an ontology due to confounding variables, existing graphical model structure learning methods can be used to further prune on the base graph without adding a cycle, such as Graphical Lasso or simple thresholding (Fattahi and Sojoudi, 2017).

4 Method

We develop a probabilistic model for lattices based on hypercube embeddings that can model both positive and negative correlations. Before describing this, we first motivate our choice to abandon OE/POE type cone-based models for this purpose.

4.1 Correlations from Cone Measures

Claim. For a pair of Bernoulli variables $p(a)$ and $p(b)$, $\text{cov}(a, b) \geq 0$ if the Bernoulli probabilities come from the volume of a cone as measured under any product (coordinate-wise) probability measure $p(x) = \prod_i^n p_i(x_i)$ on \mathbb{R}^n , where F_i , the associated CDF for p_i , is monotone increasing.

Proof. For any product measure we have

$$\int_{z \preceq x} p(z) dz = \prod_i^n \int_{x_i \leq z_i} p_i(z_i) dz_i = \prod_i^n 1 - F_i(x_i)$$

This is just the area of the unique box corresponding to $\prod_i^n [F_i(x_i), 1] \in [0, 1]^n$, under the uniform measure. This box is unique as a monotone increasing univariate CDF is bijective with $(0, 1)$ — cones in \mathbb{R}^n can be invertibly mapped to boxes of equivalent measure inside the unit hypercube $[0, 1]^n$. These boxes have only half their degrees of freedom, as they have the form $[F_i(x_i), 1]$ per dimension, (intuitively, they have one end “stuck at infinity” since the cone integrates to infinity).

So W.L.O.G. we can consider two transformed cones x and y corresponding to our Bernoulli

variables a and b , and letting $F_i(x_i) = u_i$ and $F_i(y_i) = v_i$, their intersection in the unit hypercube is $\prod_i^n [\max(u_i, v_i), 1]$.

Pairing terms in the right-hand product, we have

$$p(a, b) - p(a)p(b) = \prod_i^n (1 - \max(u_i, v_i)) - \prod_i^n (1 - u_i)(1 - v_i) \geq 0$$

since the right contains all the terms of the left and can only grow smaller. This argument is easily modified to the case of the nonnegative orthant, *mutatis mutandis*. \square

An open question for future work is what non-product measures this claim also applies to. Note that some non-product measures, such as multivariate Gaussian, can be transformed into product measures easily (*whitening*) and the above proof would still apply. It seems probable that some measures, nonlinearly entangled across dimensions, could encode negative correlations in cone volumes. However, it is not generally tractable to integrate high-dimensional cones under arbitrary non-product measures.

4.2 Box Lattices

The above proof gives us intuition about the possible form of a better representation. Cones can be mapped into boxes within the unit hypercube while preserving their measure, and the lack of negative correlation seems to come from the fact that they always have an overly-large intersection due to “pinning” the maximum in each dimension to 1. To remedy this, we propose to learn representations in the space of *all* boxes (axis-aligned hyperrectangles), gaining back an extra degree of freedom. These representations can be learned with a suitable probability measure in \mathbb{R}^n , the non-negative orthant \mathbb{R}_+^n , or directly in the unit hypercube with the uniform measure, which we elect.

We associate each concept with 2 vectors, the minimum and maximum value of the box at each dimension. Practically for numerical reasons these are stored as a minimum, a positive offset plus an ϵ term to prevent boxes from becoming too small and underflowing.

Let us define our box embeddings as a pair of vectors in $[0, 1]^n$, (x_m, x_M) , representing the maximum and minimum at each coordinate.

Then we can define a partial ordering by inclusion of boxes, and a lattice structure as

$x \wedge y = \perp$ if x and y disjoint, else

$$x \wedge y = \prod_i [\max(x_{m,i}, y_{m,i}), \min(x_{M,i}, y_{M,i})]$$

$$x \vee y = \prod_i [\min(x_{m,i}, y_{m,i}), \max(x_{M,i}, y_{M,i})]$$

where the meet is the intersecting box, or bottom (the empty set) where no intersection exists, and join is the smallest enclosing box. This lattice, considered on its own terms as a non-probabilistic object, is strictly more general than the order embedding lattice in any dimension, which is proven in Appendix B.

However, the finite sizes of all the lattice elements lead to a natural probabilistic interpretation under the uniform measure. Joint and marginal probabilities are given by the volume of the (intersection) box. For concept a with associated box (x_m, x_M) , probability is simply $p(a) = \prod_i^n (x_{M,i} - x_{m,i})$ (under the uniform measure). $p(\perp)$ is of course zero since no probability mass is assigned to the empty set.

It remains to show that this representation can represent both positive and negative correlations.

Claim. *For a pair of Bernoulli variables $p(a)$ and $p(b)$, $\text{corr}(a, b)$ can take on any value in $[-1, 1]$ if the probabilities come from the volume of associated boxes in $[0, 1]^n$.*

Proof. Boxes can clearly model disjointness (exactly -1 correlation if the total volume of the boxes equals 1). Two identical boxes give their concepts exactly correlation 1. The area of the meet is continuous with respect to translations of intersecting boxes, and all other terms in correlation stay constant, so by continuity of the correlation function our model can achieve all possible correlations for a pair of variables. \square

This proof can be extended to boxes in \mathbb{R}^n with product measures by the previous reduction.

Limitations: Note that this model cannot perfectly describe all possible probability distributions or concepts as embedded objects. For example, the complement of a box is not a box. However, queries about complemented variables can be calculated by the Inclusion-Exclusion principle, made more efficient by the fact that all non-negated terms can be grouped and calculated exactly. We show some toy exact calculations with

negated variables in Appendix A. Also, note that in a knowledge graph often true complements are not required — for example *mortal* and *immortal* are not actually complements, because the concept *color* is neither.

Additionally, requiring the total probability mass covered by boxes to equal 1, or exactly matching marginal box probabilities while modeling all correlations is a difficult box-packing-type problem and not generally possible. Modeling limitations aside, the union of boxes having mass < 1 can be seen as an open-world assumption on our KB (not all points in space have corresponding concepts, yet).

4.3 Learning

While inference (calculation of pairwise joint, unary marginal, and pairwise conditional probabilities) is quite straightforward by taking intersections of boxes and computing volumes (and their ratios), learning does not appear easy at first glance. While the (sub)gradient of the joint probability is well defined when boxes intersect, it is non-differentiable otherwise. Instead we optimize a lower bound.

Clearly $p(a \vee b) \geq p(a \cup b)$, with equality only when $a = b$, so this can give us a lower bound:

$$\begin{aligned} p(a \wedge b) &= p(a) + p(b) - p(a \cup b) \\ &\geq p(a) + p(b) - p(a \vee b) \end{aligned}$$

Where probabilities are always given by the volume of the associated box. This lower bound always exists and is differentiable, even when the joint is not. It is guaranteed to be nonpositive except when a and b intersect, in which case the true joint likelihood should be used.

While a negative bound on a probability is odd, inspecting the bound we see that its gradient will push the enclosing box to be smaller, while increasing areas of the individual boxes, until they intersect, which is a sensible learning strategy.

Since we are working with small probabilities it is advisable to negate this term and maximize the negative logarithm:

$$-\log(p(a \vee b) - p(a) - p(b))$$

This still has an unbounded gradient as the lower bound approaches 0, so it is also useful to add a constant within the logarithm function to avoid numerical problems.

Since the likelihood of the full data is usually intractable to compute as a conjunction of many negations, we optimize binary conditional and unary marginal terms separately by maximum likelihood.

In this work, we parametrize the boxes as $(\min, \Delta = \max - \min)$, with Euclidean projections after gradient steps to keep our parameters in the unit hypercube and maintain the minimum/delta constraints.

Now that we have the ability to compute probabilities and (surrogate) gradients for arbitrary marginals in the model, and by extension conditionals, we will see specific examples in the experiments.

5 Experiments

5.1 Warmup: 2D Embedding of a Toy Lattice

We begin by investigating properties of our model in modeling a small toy problem, consisting of a small hand constructed ontology over 19 concepts, aggregated from atomic synthetic examples first into a probabilistic lattice (e.g. some rabbits are brown, some are white), and then a full CPD. We model it using only 2 dimensions to enable visualization of the way the model self-organizes its “event space”, training the model by minimize weighted cross-entropy with both the unary marginals and pairwise conditional probabilities. We also conduct a parallel experiment with POE as embedded in the unit cube, where each representation is constrained to touch the faces $x = 1, y = 1$. In Figure 2, we show the representation of lattice structures by POE and the box lattice model as compared to the abstract probabilistic lattice used to construct the data, shown in Figure 1, and compare the conditional probabilities produced by our model to the ground truth, demonstrating the richer capacity of the box model in capturing strong positive and negative correlations. In Table 1, we perform a series of multivariable conditional queries and demonstrate intuitive results on high-order queries containing up to 4 variables, despite the model being trained on only 2-way information.

5.2 WordNet

We experiment on WordNet hypernym prediction, using the same train, development and test split as Vendrov et al. (2016), created by randomly taking 4,000 hypernym pairs from the 837,888-

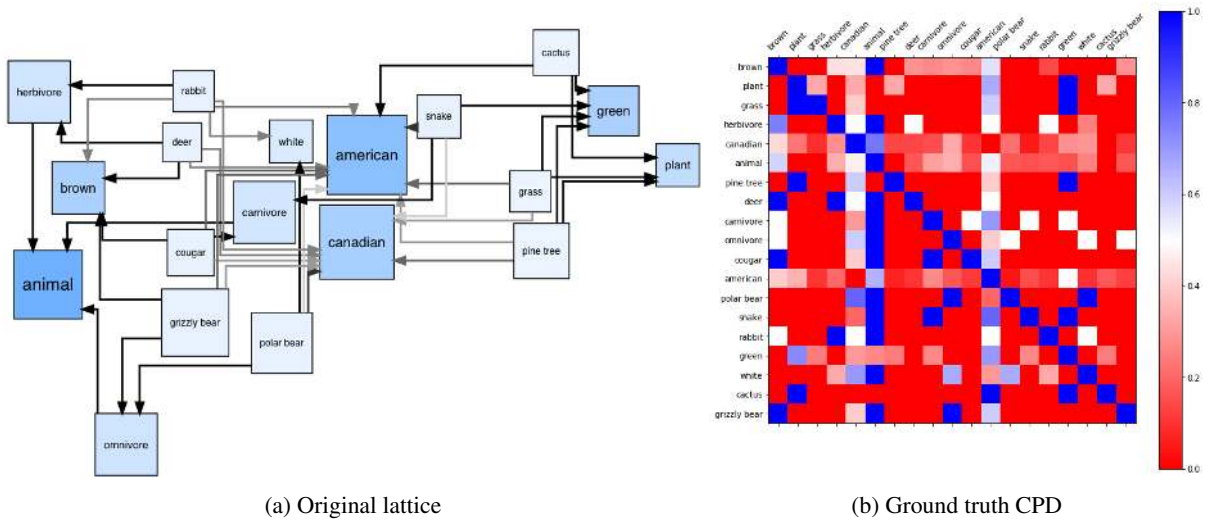


Figure 1: Representation of the toy probabilistic lattice used in Section 5.1. Darker color corresponds to more unary marginal probability. The associated CPD is obtained by a weighted aggregation of leaf elements.

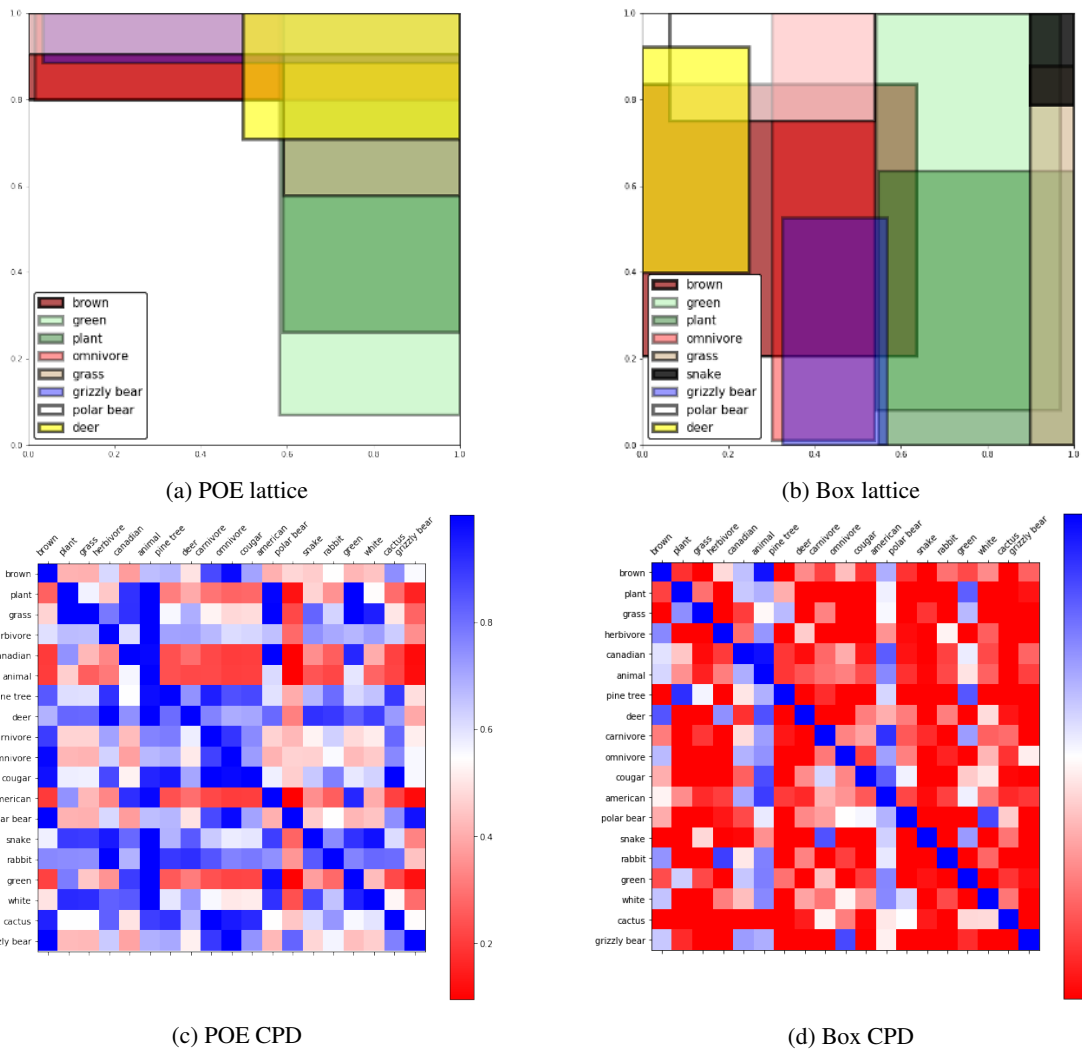


Figure 2: Lattice representations and conditional probabilities from POE vs. box lattice. Note how the box lattice model’s lack of “anchoring” to a corner allows it vastly more expressivity in matching the ground truth CPD seen in Figure 1.

P(grizzly bear ...)		P(cactus ...)		P(plant ...)	
P(grizzly bear)	0.12	P(cactus)	0.10	P(plant)	0.20
omnivore	0.29	green	0.16	green	0.37
white	0.00	plant	0.39	snake	0.00
brown	0.30	american, green	0.19	carnivore	0.00
omnivore, white	0.00	plant, green, american	0.40	cactus	0.78
omnivore, brown	0.38	american, carnivore	0.00	american, cactus	0.85

Table 1: Multi-way queries: conditional probabilities adjust when adding additional evidence or contradiction. In contrast, POE can only raise or preserve probability when conditioning.

term1	term2
craftsman.n.02	shark.n.03
homogenized_milk.n.01	apple_juice.n.01
tongue_depresser.n.01	paintbrush.n.01
deerstalker.n.01	bathing_cap.n.01
skywriting.n.01	transcript.n.01

Table 2: Negatively correlated variables produced by the model.

Method	Test Accuracy %
transitive	88.2
word2gauss	86.6
OE	90.6
Li et al. (2017)	91.3
DOE (KL)	92.3
POE	91.6
POE (100 dim)	91.7
Box	92.2
Box + CPD	92.3

Table 3: Classification accuracy on WordNet test set.

edge transitive closure of the WordNet hypernym hierarchy as positive training examples for the development set, 4,000 for the test set, and using the rest as training data. Negative training examples are created by randomly corrupting a train/development/test edge (u, v) by replacing either u or v with a randomly chosen negative node. We use their specific train/dev/test split, while Athiwaratkun and Wilson (2018) use a different train/dev split with the same test set (personal communication) to examine the effect of different negative sampling techniques. We cite their best performing model, called DOE (KL).

Since our model is probabilistic, we would like a sensible value for $P(n)$, where n is a node. We assign these marginal probabilities by looking at

the number of descendants in the hierarchy under a node, and normalizing over all nodes, taking $P(n) = \frac{|descendants(n)|}{|nodes|}$.

Furthermore, we use the graph structure (only of the subset of edges in the training set to avoid leaking data) to augment the data with approximate conditional probabilities $P(x|y)$. For each leaf, we consider all of its ancestors as pairwise co-occurrences, then aggregate and divide by the number of leaves to get an approximate joint probability distribution, $P(x, y) = \frac{|x, y \text{ co-occur in ancestor set}|}{|leaves|}$. With this and the unary marginals, we can create a conditional probability table, which we prune based on the difference of $P(x|y)$ and $P(y|x)$ and add cross entropy with these conditional “soft edges” to the training data. We refer to experiments using this additional data as Box + CPD in Table 3.

We use 50 dimensions in our experiments. Since our model has 2 parameters per dimension, we also perform an apples-to-apples comparison with a 100D POE model. As seen in Table 3, we outperform POE significantly even with this added representational power. We also observe sensible negatively correlated examples, shown in 2, in the trained box model, while POE cannot represent such relationships. We tune our models on the development set, with parameters documented in Appendix D.1. We observe that not only does our model outperform POE, it beats all previous results on WordNet, aside from the concurrent work of Athiwaratkun and Wilson (2018) (using different train/dev negative examples), the baseline POE model does as well. This indicates that probabilistic embeddings for transitive relations are a promising avenue for future work. Additionally, the ability of the model to learn from the expected “soft edges” improves it to state-of-the-art level. We expect that co-occurrence counts gathered from real textual corpora, rather than merely

aggregating up the WordNet lattice, would further strengthen this effect.

5.3 Flickr Entailment Graph

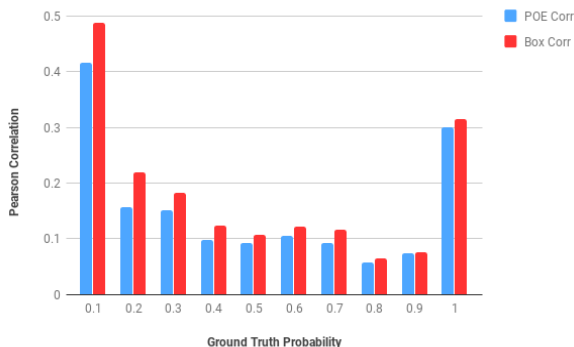


Figure 3: R between model and gold probabilities.

<i>Full test data</i>	$P(x y)$	
	KL	Pearson R
POE	0.031	0.949
POE*	0.031	0.949
Box	0.020	0.967
<i>Unseen pairs</i>		
POE	0.048	0.920
POE*	0.046	0.925
Box	0.025	0.957
<i>Unseen words</i>		
POE	0.127	0.696
POE*	0.084	0.854
Box	0.050	0.900

Table 4: KL and Pearson correlation between model and gold probability.

We conduct experiments on the large-scale Flickr entailment dataset of 45 million image caption pairs. We use the exactly same train/dev/test from Lai and Hockenmaier (2017). We use a slightly different unseen word pairs and unseen words test data, obtained from the author. We include their published results and also use their published code, marked *, for comparison.

For these experiments, we relax our boxes from the unit hypercube to the nonnegative orthant and obtain probabilities under the exponential measure, $p(x) = \exp(-x)$. We enforce the nonnegativity constraints by clipping the LSTM-generated embedding (Hochreiter and Schmidhuber, 1997) for the box minimum with a ReLU,

and parametrize our Δ embeddings using a soft-plus activation to prevent dead units. As in Lai and Hockenmaier (2017), we use 512 hidden units in our LSTM to compose sentence vectors. We then apply two single-layer feed-forward networks with 512 units applied to the final LSTM state to produce the embeddings.

As we can see from Table 4, we note large improvements in KL and Pearson correlation to the ground truth entailment probabilities. In further analysis, Figure 3 demonstrates that while the box model outperforms POE in nearly every regime, the highest gains come from the comparatively difficult to calibrate small entailment probabilities, indicating the greater capability of our model to produce fine-grained distinctions.

6 Conclusion and Future Work

We have only scratched the surface of possible applications. An exciting direction is the incorporation of multi-relational data for general knowledge representation and inference. Secondly, more complex representations, such as $2n$ -dimensional products of 2-dimensional convex polyhedra, would offer greater flexibility in tiling event space. Improved inference of the latent boxes, either through better optimization or through Bayesian approaches is another natural extension. Our greatest interest is in the application of this powerful new tool to the many areas where other structured embeddings have shown promise.

7 Acknowledgments

We thank Alice Lai for making the code from her original paper public, and for providing the additional unseen pairs and unseen words data. We also thank Haw-Shiuan Chang, Laurent Dinh, and Ben Poole for helpful discussions. We also thank the anonymous reviewers for their constructive feedback.

This work was supported in part by the Center for Intelligent Information Retrieval and the Center for Data Science, in part by the Chan Zuckerberg Initiative under the project Scientific Knowledge Base Construction., and in part by the National Science Foundation under Grant No. IIS-1514053. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

- Ben Athiwaratkun and Andrew Gordon Wilson. 2018. [On modeling hierarchical data via probabilistic order embeddings](#). In *International Conference on Learning Representations*.
- Mohit Bansal, David Burkett, Gerard De Melo, and Dan Klein. 2014. Structured learning for taxonomy induction with belief propagation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1041–1051.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795.
- Y. J. Chu and T. H. Liu. 1995. On the shortest arborescence of a directed graph. *Science Sinica*, 20.
- Michael Collins, Sanjoy Dasgupta, and Robert E Schapire. 2002. A generalization of principal components analysis to the exponential family. In *Advances in neural information processing systems*, pages 617–624.
- Katrin Erk. 2009. [Representing words as regions in vector space](#). In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning*, CoNLL '09, pages 57–65, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Salar Fattahi and Somayeh Sojoudi. 2017. Graphical lasso and thresholding: Equivalence and closed-form solutions. *arXiv preprint arXiv:1708.09479*.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. 2008. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256.
- Shizhu He, Kang Liu, Guoliang Ji, and Jun Zhao. 2015. [Learning to represent knowledge graphs with gaussian embedding](#). In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, CIKM '15, pages 623–632, New York, NY, USA. ACM.
- David Heckerman, Dan Geiger, and David M Chickering. 1995. Learning bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20(3):197–243.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Alice Lai and Julia Hockenmaier. 2017. Learning to predict denotational probabilities for modeling entailment. In *EACL*.
- Xiang Li, Luke Vilnis, and Andrew McCallum. 2017. Improved representation learning for predicting commonsense ontologies. *NIPS Workshop on Structured Prediction*.
- George A Miller. 1995. WordNet: a lexical database for English. *Communications of the ACM*.
- Maximillian Nickel and Douwe Kiela. 2017. [Poincaré embeddings for learning hierarchical representations](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6338–6347. Curran Associates, Inc.
- Judea Pearl. 1988. *Probabilistic reasoning in intelligent systems*.
- Daniel M Roy and Yee W Teh. 2009. The mondrian process. In *Advances in neural information processing systems*, pages 1377–1384.
- Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2016. Order-embeddings of images and language. In *ICLR*.
- Luke Vilnis and Andrew McCallum. 2015. Word representations via gaussian embedding. In *ICLR*.
- Adriaan C. Zaanen. 1997. *Introduction to Operator Theory in Riesz Spaces*. Springer Berlin Heidelberg.