

Predicting Semantic Relations using Global Graph Properties

Yuval Pinter

School of Interactive Computing
Georgia Institute of Technology
Atlanta, GA
uvp@gatech.edu

Jacob Eisenstein

School of Interactive Computing
Georgia Institute of Technology
Atlanta, GA
jacobe@gatech.edu

Abstract

Semantic graphs, such as WordNet, are resources which curate natural language on two distinguishable layers. On the local level, individual relations between synsets (semantic building blocks) such as hypernymy and meronymy enhance our understanding of the words used to express their meanings. Globally, analysis of graph-theoretic properties of the entire net sheds light on the structure of human language as a whole. In this paper, we combine global and local properties of semantic graphs through the framework of Max-Margin Markov Graph Models (M3GM), a novel extension of Exponential Random Graph Model (ERGM) that scales to large multi-relational graphs. We demonstrate how such global modeling improves performance on the local task of predicting semantic relations between synsets, yielding new state-of-the-art results on the WN18RR dataset, a challenging version of WordNet link prediction in which “easy” reciprocal cases are removed. In addition, the M3GM model identifies multirelational motifs that are characteristic of well-formed lexical semantic ontologies.

1 Introduction

Semantic graphs, such as WordNet (Fellbaum, 1998), encode the structural qualities of language as a representation of human knowledge. On the local level, they describe connections between specific semantic concepts, or **synsets**, through individual edges representing relations such as hypernymy (‘is-a’) or meronymy (‘is-part-of’); on the global level, they encode emergent regular properties in the induced relation graphs. Local properties have been subject to extensive study in recent years via the task of **relation prediction**, where individual edges are found based mostly on

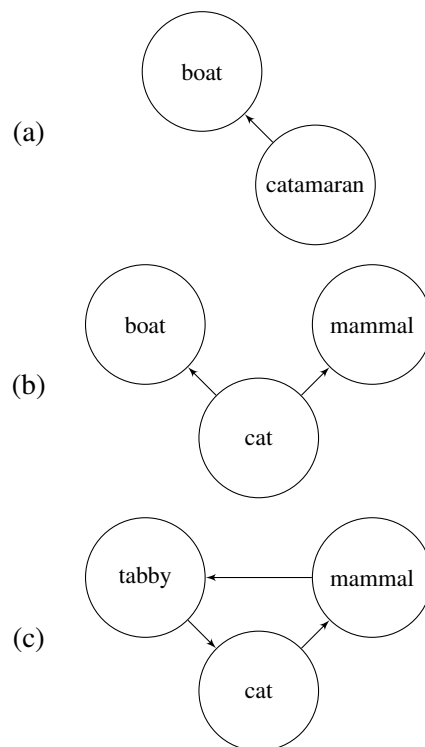


Figure 1: Probable (a) and improbable (b-c) structures in a hypothetical hypernym graph.

distributional methods that embed synsets and relations into a vector space (e.g. Socher et al., 2013; Bordes et al., 2013; Toutanova and Chen, 2015; Neelakantan et al., 2015). In contrast, while the structural regularity and significance of global aspects of semantic graphs is well-attested (Sigman and Cecchi, 2002), global properties have rarely been used in prediction settings. In this paper, we show how global semantic graph features can facilitate in local tasks such as relation prediction.

To motivate this approach, consider the hypothetical hypernym graph fragments in Figure 1: in (a), the semantic concept (synset) ‘catamaran’ has

a single hypernym, ‘boat’. This is a typical property across a standard hypernym graph. In (b), the synset ‘cat’ has two hypernyms, an unlikely event. While a local relation prediction model might mistake the relation between ‘cat’ and ‘boat’ to be plausible, for whatever reason, a high-order graph-structure-aware model should be able to discard it based on the knowledge that a synset should not have more than one hypernym. In (c), an impossible situation arises: a cycle in the hypernym graph leads each of the participating synsets to be predicted by transitivity as its own hypernym, contrary to the relation’s definition. However, a purely local model has no explicit mechanism for rejecting such an outcome.

In this paper, we examine the effect of global graph properties on the link structure via the WordNet relation prediction task. Our hypothesis is that features extracted from the entire graph can help constrain local predictions to structurally sound ones (Guo et al., 2007). Such features are often manifested as aggregate counts of small sub-graph structures, known as **motifs**, such as the number of nodes with two or more outgoing edges, or the number of cycles of length 3. Returning to the example in Figure 1, each of these features will be affected when graphs (b) and (c) are evaluated, respectively.

To estimate weights on local and global graph features, we build on the **Exponential Random Graph Model** (ERGM), a log-linear model over networks utilizing global graph features (Holland and Leinhardt, 1981). In ERGMs, the likelihood of a graph is computed by exponentiating a weighted sum of the features, and then normalizing over all possible graphs. This normalization term grows exponentially in the number of nodes, and in general cannot be decomposed into smaller parts. Approximations are therefore necessary to fit ERGMs on graphs with even a few dozen nodes, and the largest known ERGMs scale only to thousands of nodes (Schmid and Desmarais, 2017). This is insufficient for WordNet, which has an order of 10^5 nodes.

We extend the ERGM framework in several ways. First, we replace the maximum likelihood objective with a margin-based objective, which compares the observed network against alternative networks; we call the resulting model the **Max-Margin Markov Graph Model** (M3GM), drawing on ideas from structured prediction (Taskar

et al., 2004). The gradient of this loss is approximated by importance sampling over candidate negative edges, using a local relational model as a proposal distribution. The complexity of each epoch of estimation is thus linear in the number of edges, making it possible to scale up to the 10^5 nodes in WordNet.¹ Second, we address the multi-relational nature of semantic graphs, by incorporating a combinatorial set of labeled motifs. Finally, we link graph-level relational features with distributional information, by combining the M3GM with a dyad-level model over word sense embeddings.

We train M3GM as a re-ranker, which we apply to a strong local-feature baseline on the WN18RR dataset (Dettmers et al., 2018). This yields absolute improvements of 3-4 points on all commonly-used metrics. Model inspection reveals that M3GM assigns importance to features from all relations, and captures some interesting inter-relational properties that lend insight into the overall structure of WordNet.²

2 Related Work

Relational prediction in semantic graphs. Recent approaches to relation prediction in semantic graphs generally start by embedding the semantic concepts into a shared space and modeling relations by some operator that induces a score for an embedding pair input. We use several of these techniques as base models (Nickel et al., 2011; Bordes et al., 2013; Yang et al., 2014); detailed description of these methods is postponed to Section 3.2. Socher et al. (2013) generalize over the approach of Nickel et al. (2011) by using a bilinear tensor which assigns multiple parameters for each relation; Shi and Weninger (2017) project the node embeddings in a translational model similar to that of Bordes et al. (2013); Dettmers et al. (2018) apply a convolutional neural network by reshaping synset embeddings to 2-dimensional matrices. None of these embedding-based approaches incorporate structural information; in general, improvements in embedding-based methods are expected to be complementary to our approach.

¹Although in principle the number of edges could grow quadratically with the number of nodes, Steyvers and Tenenbaum (2005) show that semantic graphs like WordNet tend to be very sparse, so that the number of observed edges grows roughly linearly with the number of nodes.

²Our code is available at <http://www.github.com/yuvalpinter/m3gm>.

Some recent works compose single edges into more intricate motifs, such as Guu et al. (2015), who define a task of **path prediction** and compose various functions to solve it. They find that compositionalized bilinear models perform best on WordNet. Minervini et al. (2017) train link-prediction models against an adversary that produces examples which violate structural constraints such as symmetry and transitivity. Another line of work builds on local neighborhoods of relation interactions and automatic detection of relations from syntactically parsed text (Riedel et al., 2013; Toutanova et al., 2015). Schlichtkrull et al. (2017) use Graph Convolutional Networks to predict relations while considering high-order neighborhood properties of the nodes in question. In general, these methods aggregate information over local neighborhoods, but do not explicitly model structural motifs.

Our model introduces interaction features between relations (e.g., hypernyms and meronyms) for the goal of relation prediction. To our knowledge, this is the first time that relation interaction is explicitly modeled into a relation prediction task. Within the ERGM framework, Lu et al. (2010) train a limited set of combinatory path features for social network link prediction.

Scaling exponential random graph models.

The problem of approximating the denominator of the ERGM probability has been an active research topic for several decades. Two common approximation methods exist in the literature. In **Maximum Pseudolikelihood Estimation (MPLE; Strauss and Ikeda, 1990)**, a graph’s probability is decomposed into a product of the probability for each edge, which in turn is computed based on the ERGM feature difference between the graph excluding the edge and the full graph. **Monte Carlo Maximum Likelihood Estimation (MCMLE; Snijders, 2002)** follows a sampling logic, where a large number of graphs is randomly generated from the overall space under the intuition that the sum of their scores would give a good approximation for the total score mass. The probability for the observed graph is then estimated following normalization conditioned on the sampling distribution, and its precision increases as more samples are gathered. Recent work found that applying a parametric bootstrap can increase the reliability of MPLE, while retaining its superiority in training speed (Schmid

and Desmarais, 2017). Despite this result, we opted for an MCMLE-based approach for M3GM, mainly due to the ability to keep the number of edges constant in each sampled graph. This property is important in our setup, since local edge scores added or removed to the overall graph score can occasionally dominate the objective function, giving unintended importance to the overall edge count.

3 Max-Margin Markov Graph Models

Consider a graph $G = (V, E)$, where V is a set of vertices and $E = \{(s_i, t_i)\}_{i=1}^{|E|}$ is a set of directed edges. The ERGM scoring function defines a probability over $\mathcal{G}_{|V|}$, the set of all graphs with $|V|$ nodes. This probability is defined as a log-linear function,

$$P_{\text{ERGM}}(G) \propto \psi_{\text{ERGM}}(G) = \exp(\boldsymbol{\theta}^T \mathbf{f}(G)), \quad (1)$$

where \mathbf{f} is a feature function, from graphs to a vector of feature counts. Features are typically counts of motifs — small subgraph structures — as described in the introduction. The vector $\boldsymbol{\theta}$ is the parameter to estimate.

In this section we discuss our adaptation of this model to the domain of semantic graphs, leveraging their idiosyncratic properties. Semantic graphs are composed of multiple relation types, which the feature space needs to accommodate; their nodes are linguistic constructs (semantic concepts) associated with complex interpretations, which can benefit the graph representation through incorporating their embeddings in \mathbb{R}^d into a new scoring model. We then present our M3GM framework to perform reliable and efficient parameter estimation on the new model.

3.1 Graph Motifs as Features

Based on common practice in ERGM feature extraction (e.g., Morris et al., 2008), we select the following graph features as a basis:

- Total edge count;
- Number of cycles of length k , for $k \in \{2, 3\}$;
- Number of nodes with exactly k outgoing (incoming) edges, for $k \in \{1, 2, 3\}$;
- Number of nodes with **at least** k outgoing (incoming) edges, for $k \in \{1, 2, 3\}$;
- Number of paths of length 2;

- **Transitivity:** the proportion of length-2 paths $u \rightarrow v \rightarrow w$ where an edge $u \rightarrow w$ also exists.

Semantic graphs are multigraphs, where multiple relationships (hypernymy, meronymy, derivation, etc.) are overlaid atop a common set of nodes. For each relation r in the relation inventory \mathcal{R} , we denote its edge set as E_r , and redefine $E = \bigcup_{r \in \mathcal{R}} E_r$, the union of all labeled edges. Some relations do not produce a connected graph, while others may coincide with each other frequently, possibly in regular but intricate patterns: for example, derivation relations tend to occur between synsets in the higher, more abstract levels of the hypernym graph. We represent this complexity by expanding the feature space to include relation-sensitive **combinatory motifs**. For each feature template from the basis list above, we extract features for all possible combinations of relation types existing in the graph. Depending on the feature type, these could be relation singletons, pairs, or triples; they may be order-sensitive or order-insensitive. For example:

- A combinatory ‘transitivity’ feature will be extracted for the proportion of paths $u \xrightarrow{\text{hypernym}} v \xrightarrow{\text{meronym}} w$ where an edge $u \xrightarrow{\text{has-part}} w$ also exists.
- A combinatory ‘2-outgoing’ feature will be extracted for the number of nodes with exactly one *derivation* and one *has-part*.

The number of features thus scales in $O(|\mathcal{R}|^K)$ for a feature basis which involves up to K edges in any feature, and so our 17 basis features (with $K = 3$) generate a combinatory feature set with roughly 3,000 features for the 11-relation version of WordNet used in our experiments (see Section 4.1).

3.2 Local Score Component

In classical ERGM application domains such as social media or biological networks, nodes tend to have little intrinsic distinction, or at least little meaningful intrinsic information that may be extracted prior to applying the model. In semantic graphs, however, the nodes represent synsets, which are associated with information that is both valuable to predicting the graph structure and approximable using unsupervised techniques such as embedding into a common d -dimensional vector

space based on copious amounts of available data. We thus modify the traditional scoring function from eq. (1) to include node-specific information, by introducing a relation-specific **association operator** $\mathcal{A}^{(r)} : V \times V \rightarrow \mathbb{R}$:

$$\begin{aligned} \psi_{\text{ERGM}^+}(G) &= \\ &= \exp \left(\boldsymbol{\theta}^T \mathbf{f}(G) + \sum_{r \in \mathcal{R}} \sum_{(s,t) \in E_r} \mathcal{A}^{(r)}(s,t) \right). \end{aligned} \quad (2)$$

The association operator generalizes various models from the relation prediction literature:

TransE (Bordes et al., 2013) embeds each relation r into a vector in the shared space, representing a ‘difference’ between sources and targets, to compute the association score under a translational objective,

$$\mathcal{A}_{\text{TRANSE}}^{(r)}(s,t) = -\|e_s + e_r - e_t\|.$$

BiLin (Nickel et al., 2011) embeds relations into full-rank matrices, computing the score by a bilinear multiplication,

$$\mathcal{A}_{\text{BILIN}}^{(r)}(s,t) = e_s^T \mathbf{W}_r e_t.$$

DistMult (Yang et al., 2014) is a special case of **BiLin** where the relation matrices are diagonal, reducing the computation to a ternary dot product,

$$\mathcal{A}_{\text{DISTMULT}}^{(r)}(s,t) = \langle e_s, e_r, e_t \rangle = \sum_{i=1}^d e_{s_i} e_{r_i} e_{t_i}.$$

3.3 Parameter Estimation

The probabilistic formulation of ERGM requires the computation of a normalization term that sums over all possible graphs with a given number of nodes, \mathcal{G}_N . The set of such graphs grows at a rate that is super-exponential in the number of nodes, making exact computation intractable even for networks that are orders of magnitude smaller than semantic graphs like WordNet. One solution is to approximate probability using a variant of the Monte Carlo Maximum Likelihood Estimation (MCMLE) procedure,

$$\log P(G) \approx \log \psi(G) - \log \frac{|\mathcal{G}_{|V|}|}{M} \sum_{\tilde{G} \sim \mathcal{G}_{|V|}} \psi(\tilde{G}), \quad (3)$$

where M is the number of networks \tilde{G} sampled from $\mathcal{G}_{|V|}$, the space of all (multirelational) edge sets on nodes V . Each \tilde{G} is referred to as a *negative sample*, and the goal of estimation is to assign low scores to these samples, in comparison with the score assigned to the observed network G .

Network samples can be obtained using edge-wise negative sampling. For each edge $s \xrightarrow{r} t$ in the training network G , we remove it temporarily and consider T alternative edges, keeping the source s and relation r constant, and sampling a target \tilde{t} from a **proposal distribution** Q . Every such substitution produces a new graph \tilde{G} ,

$$\tilde{G} = G \cup \{s \xrightarrow{r} \tilde{t}\} \setminus \{s \xrightarrow{r} t\}. \quad (4)$$

Large-margin objective. Rather than approximating the log probability, as in MCMLE estimation, we propose a margin loss objective: the log score for each negative sample \tilde{G} should be below the log score for G by a margin of at least 1. This motivates the hinge loss,

$$\mathcal{L}(\Theta, \tilde{G}; G) = \left(1 - \log \psi_{\text{ERGM}^+}(G) + \log \psi_{\text{ERGM}^+}(\tilde{G})\right)_+, \quad (5)$$

where $(x)_+ = \max(0, x)$. Recall that the scoring function ψ_{ERGM^+} includes both the local association score for the alternative edge and the global graph features for the resulting graph. However, it is not necessary to recompute all association scores; we need only subtract the association score for the deleted edge $s \xrightarrow{r} t$, and add the association score for the sampled edge $s \xrightarrow{r} \tilde{t}$.

The overall loss function is the sum over $N = |E| \times T$ negative samples, $\{\tilde{G}^{(i)}\}_{i=1}^N$, plus an L_2 regularizer on the model parameters,

$$\mathcal{L}(\Theta; G) = \lambda \|\Theta\|_2^2 + \sum_{i=1}^N \mathcal{L}(\Theta, \tilde{G}^{(i)}). \quad (6)$$

Proposal distribution. The proposal distribution Q used to sample negative edges is defined to be proportional to the local association scores of edges not present in the training graph:

$$Q(\tilde{t} \mid s, r, G) \propto \begin{cases} 0 & s \xrightarrow{r} \tilde{t} \in G \\ \mathcal{A}^{(r)}(s, \tilde{t}) & s \xrightarrow{r} \tilde{t} \notin G. \end{cases} \quad (7)$$

By preferring edges that have high association scores, the negative sampler helps push the M3GM parameters away from likely false positives.

4 Relation Prediction

We evaluate M3GM on the relation graph edge prediction task.³ Data for this task consists of a set of labeled edges, i.e. tuples of the form (s, r, t) , where s and t denote source and target entities, respectively. Given an edge from an evaluation set, two prediction instances are created by hiding the source and target side, in turn. The predictor is then evaluated on its ability to predict the hidden entity, given the other entity and the relation type.⁴

4.1 WN18RR Dataset

A popular relation prediction dataset for WordNet is the subset curated as WN18 (Bordes et al., 2013, 2014), containing 18 relations for about 41,000 synsets extracted from WordNet 3.0. It has been noted that this dataset suffers from considerable leakage: edges from reciprocal relations such as hypernym / hyponym appear in one direction in the training set and in the opposite direction in dev / test (Socher et al., 2013; Dettmers et al., 2018). This allows trivial rule-based baselines to achieve high performance. To alleviate this concern, Dettmers et al. (2018) released the **WN18RR** set, removing seven relations altogether. However, even this dataset retains four symmetric relation types: *also see*, *derivationally related form*, *similar to*, and *verb group*. These symmetric relations can be exploited by defaulting to a simple rule-based predictor.

4.2 Metrics

We report the following metrics, common in ranking tasks and in relation prediction in particular: **MR**, the Mean Rank of the desired entity; **MRR**, Mean Reciprocal Rank, the main evaluation metric; and **H@k**, the proportion of Hits (true entities) found in the top k of the lists, for $k \in \{1, 10\}$. Unlike some prior work, we do not type-restrict the possible relation predictions (so, e.g., a *verb group* link may select a noun, and that would count against the model).

4.3 Systems

We evaluate a single-rule baseline, three association models, and two variants of the M3GM re-

³Sometimes referred to as Knowledge Base Completion, e.g. in Socher et al. (2013).

⁴We follow prior work in excluding the following from the ranked lists: the known entity (no self loops); entities from the training set which fit the instance; other entities in the evaluation set.

ranker trained on top of the best-performing association baseline.

4.3.1 RULE

We include a single-rule baseline that predicts a relation between s and t in the evaluation set if the same relation was encountered between t and s in the training set. All other models revert to this baseline for the four symmetric relations.

4.3.2 Association Models

The next group of systems compute local scores for entity-relation triplets. They all encode entities into embeddings e . Each of these systems, in addition to being evaluated as a baseline, is also used for computing association scores in M3GM, both in the proposal distribution (see Section 3.3) and for creating lists to be re-ranked (see below): **TRANSE**, **BiLIN**, **DISTMULT**. For detailed descriptions, see Section 3.2.

4.3.3 Max-Margin Markov Graph Model

The M3GM is applied as a re-ranker. For each relation and source (target), the top K candidate targets (sources) are retrieved based on the local association scores. Each candidate edge is introduced into the graph, and the score $\psi_{\text{ERGM}^+}(G)$ is used to re-rank the top- K list.

We add a variant to this protocol where the graph score and association score are weighted by α and $1 - \alpha$, respectively, before being summed. We tune a separate α_r for each relation type, using the development set’s mean reciprocal rank (MRR). These hyperparameter values offer further insight into where the M3GM signal benefits relation prediction most (see Section 6).

Since we do not apply the model to the symmetric relations (scored by the RULE baseline), they are excluded from the sampling protocol described in eq. (5), although their edges do contribute to the combinatory graph feature vector f .

Our default setting backpropagates loss into only the graph weight vector θ . We experiment with a model variant which backpropagates into the association model and synset embeddings as well.

4.4 Synset Embeddings

For the association component of our model, we require embedding representations for WordNet synsets. While unsupervised word embedding techniques go a long way in representing wordforms (Collobert et al., 2011; Mikolov et al., 2013;

Pennington et al., 2014), they are not immediately applicable to the semantically-precise domain of synsets. We explore two methods of transforming pre-trained word embeddings into synset embeddings.

Averaging. A straightforward way of using word embeddings to create synset embeddings is to collect the words representing the synset as surface form within the WordNet dataset and average their embeddings (Socher et al., 2013). We apply this method to pre-trained GloVe embeddings (Pennington et al., 2014) and pre-trained FastText embeddings (Bojanowski et al., 2017), averaging over the set of all wordforms in all lemmas for each synset, and performing a case-insensitive query on the embedding dictionary. For example, the synset ‘*determine.v.01*’ lists the following lemmas: ‘determine’, ‘find’, ‘find_out’, ‘ascertain’. Its vector is initialized as

$$\frac{1}{5}(e_{\text{determine}} + 2 \cdot e_{\text{find}} + e_{\text{out}} + e_{\text{ascertain}}).$$

AutoExtend retrofitting + Mimick. AutoExtend is a method developed specifically for embedding WordNet synsets (Rothe and Schütze, 2015), in which pre-trained word embeddings are retrofitted to the tripartite relation graph connecting wordforms, lemmas, and synsets. The resulting synset embeddings occupy the same space as the word embeddings. However, some WordNet senses are not represented in the underlying set of pre-trained word embeddings.⁵ To handle these cases, we trained a character-based model called MIMICK, which learns to predict embeddings for out-of-vocabulary items based on their spellings (Pinter et al., 2017). We do not modify the spelling conventions of WordNet synsets before passing them to Mimick, so e.g. ‘*mask.n.02*’ (the second synset corresponding to ‘mask’ as a noun) acts as the input character sequence as is.

Random initialization. In preliminary experiments, we attempted training the association models using randomly-initialized embeddings. These proved to be substantially weaker than distributionally-informed embeddings and we do not report their performance in the results section. We view this finding as strong evidence to support the necessity of a distributional signal in a type-level semantic setup.

⁵We use the out-of-the-box vectors supplied in <http://www.cis.lmu.de/~sascha/AutoExtend>.

	System	MR	MRR	H@10	H@1
	RULE	13396	35.26	35.27	35.23
1	DISTMULT	1111	43.29	50.73	39.67
2	BILIN	738	45.36	52.93	41.37
3	TRANSE	2231	46.07	55.65	41.41
4	M3GM	2231	47.94	57.72	43.26
5	M3GM _{α_r}	2231	48.30	57.59	43.78

Table 1: Results on development set (all metrics except MR are $\times 100$). M3GM lines use TRANSE as their association model. In M3GM _{α_r} , the graph component is tuned post-hoc against the local component per relation.

4.5 Setup

Following tuning experiments, we train the association models on synset embeddings with $d = 300$, using a negative log-likelihood loss function over 10 negative samples and iterating over symmetric relations once every five epochs. We optimize the loss using AdaGrad with $\eta = 0.01$, and perform early stopping based on the development set mean reciprocal rank. M3GM is trained in four epochs using AdaGrad with $\eta = 0.1$. We set M3GM’s re-rank list size $K = 100$ and, following tuning, the regularization parameter $\lambda = 0.01$ and negative sample count per edge $T = 10$. Our models are all implemented in DyNet (Neubig et al., 2017).

5 Results

Table 1 presents the results on the development set. Lines 1-3 depict the results for local models using averaged FastText embedding initialization, showing that the best performance in terms of MRR and top-rank hits is achieved by TRANSE. Mean Rank does not align with the other metrics; this is an interpretable tradeoff, as both BILIN and DISTMULT have an inherent preference for correlated synset embeddings, giving a stronger fallback for cases where the relation embedding is completely off, but allowing less freedom for separating strong cases from correlated false positives, compared to a translational objective.

Effect of global score. There is a clear advantage to re-ranking the top local candidates using the score signal from the M3GM model (line 4). These results are further improved when the graph score is weighted against the association component per relation (line 5). We obtain similar improvements when re-ranking the predictions from DISTMULT and BILIN.

	System	MR	MRR	H@10	H@1
	RULE	13396	35.26	35.26	35.26
	COMPLEX [†]	5261	44	51	41
	CONVE [†]	5277	46	48	39
	CONVKB [†]	2554	24.8	52.5	
	TRANSE	2195	46.59	55.55	42.26
	M3GM _{α_r}	2193	49.83	59.02	45.37

Table 2: Main results on test set. [†] These models were not re-implemented, and are reported as in Nguyen et al. (2018) and in Dettmers et al. (2018).

The M3GM training procedure is not useful in fine-tuning the association model via backpropagation: this degrades the association scores for true edges in the evaluation set, dragging the re-ranked results along with them to about a 2-point drop relative to the untuned variant.

Table 2 shows that our main results transfer onto the test set, with even a slightly larger margin. This could be the result of the greater edge density of the combined training and dev graphs, which enhance the global coherence of the graph structure captured by M3GM features. To support this theory, we tested the M3GM model trained on only the training set, and its test set performance was roughly one point worse on all metrics, as compared with the model trained on the training+dev data.

Synset embedding initialization. We trained association models initialized on AutoExtend+Mimick vectors (see Section 4.4). Their performance, inferior to averaged FastText vectors by about 1-2 MRR points on the dev set, is somewhat at odds with findings from previous experiments on WordNet (Guu et al., 2015). We believe the decisive factor in our result is the size of the training corpus used to create FastText embeddings, along with the increase in resulting vocabulary coverage. Out of 124,819 lemma tokens participating in 41,105 synsets, 118,051 had embeddings available (94.6%; type-level coverage 88.1%). Only 530 synsets (1.3%) finished this initialization process with no embedding and were assigned random vectors. AutoExtend, fit for embeddings from Mikolov et al. (2013) which were trained on a smaller corpus, offers a weaker signal: 13,377 synsets (32%) had no vector and needed Mimick initialization.

Positive	
1	$s \xrightarrow{\text{member_meronym}} \textcircled{t}$
2	$s \xrightarrow{\text{has_part}} \textcircled{t}$
3	$s \xrightarrow{\text{hypernym}} \textcircled{t} \xrightarrow{\text{derivationally_related_form}} u$
Negative	
4	$s \xrightarrow{\text{hypernym}} \textcircled{t}$
5	$\textcircled{s} \xleftarrow{\text{hypernym}} \textcircled{t}$
6	$s \xrightarrow{\text{member_meronym}} \textcircled{t} \xrightarrow{\text{instance_hypernym}} u$
7	$s_1 \xrightarrow{\text{has_part}} \textcircled{t} \xleftarrow{\text{verb_group}} s_2$

Table 3: Select heavyweight features (motifs) following best dev set training using M3GM. Circled nodes count towards the motif.

6 Graph Analysis

As a consequence of the empirical experiment, we aim to find out what M3GM has learned about WordNet. Table 3 presents a sample of top-weighted motifs. Lines 1 and 2 demonstrate that the model prefers a broad scattering of targets for the *member_meronym* and *has_part* relations⁶, which are flat and top-downwards hierarchical, respectively, while line 4 shows that a multitude of unique *hypernyms* is undesired, as expected from a bottom-upwards hierarchical relation. Line 5 enforces the asymmetry of the hypernym relation.

Lines 3, 6, and 7 hint at deeper interactions between the different relation types. Line 3 shows that the model assigns positive weights to hypernyms which have derivationally-related forms, suggesting that the derivational equivalence classes in the graph tend to exist in the higher, more abstract levels of the hypernym hierarchy, as noted in Section 3.1. Line 6 captures a semantic conflict: synsets located in the lower, specific levels of the graph can be specified either as *instances* of abstract concepts⁷, or as *members* of less specific concrete classes, but not as both. Line 7 may have captured a nodal property – since *part_of* is a relation which holds between nouns, and *verb_group* holds between verbs, this negative weight assignment may be the manifestation of a part-of-speech uniqueness constraint. In addition, in features 3 and 7 we see the importance of symmetric relations (here *derivationally_related_form*

⁶Example edges: ‘America’ → ‘American’, ‘face’ → ‘mouth’, respectively.

⁷Example *instance_hypernym* edge: ‘Rome’ → ‘national capital’.

Source	Relation	Correct target	Outranking local target(s)
indian lettuce	<i>hypernym</i>	herb	garden lettuce
austria	<i>has_part</i>	vienna	germany, hungary, france, european union

Table 4: Successful M3GM re-ranking examples.

Relation r	α_r	Relation r	α_r
mem. of domain usage	0.78	hypernym	0.64
mem. of domain region	0.77	domain topic of	0.38
member meronym	0.67	has part	0.33
instance hypernym	0.65		

Table 5: Graph score weights found for relations on the dev set. Zero means graph score is not considered at all for this relation, one means only it is considered.

and *verb_group*, respectively), which manage to be represented in the graph model despite not being directly trained on.

Table 4 presents examples of relation targets successfully re-ranked thanks to these features. The first false connection created a new unique hypernym, ‘garden lettuce’, downgraded by the graph score through incrementing the count of negatively-weighted feature 4. In the second case, ‘vienna’ was brought from rank 10 to rank 1 since it incremented the count for the positively-weighted feature 2, whereas all targets ranked above it by the local model were already *has_part*s, mostly of ‘europe’.

The α_r values weighing the importance of M3GM scores in the overall function, found per relation through grid search over the development set, are presented in Table 5. It appears that for all but two relations, the best-performing model preferred the signal from the graph features to that from the association model ($\alpha_r > 0.5$). Based on the surface properties of the different relation graphs, the decisive factor seems to be that *synset_domain_topic_of* and *has_part* pertain mostly to very common concepts, offering good local signal from the synset embeddings, whereas the rest include many long-tail, low-frequency synsets that require help from global features to detect regularity.

7 Conclusion

This paper presents a novel method for reasoning about semantic graphs like WordNet, combining the distributional coherence between individual entity pairs with the structural coherence of network motifs. Applied as a re-ranker, this method substantially improves performance on link prediction. Our analysis of results from Table 3, lines 6 and 7, suggests that adding graph motifs which qualify their adjacent nodes in terms of syntactic function or semantic category may prove useful.

From a broader perspective, M3GM can do more as a probabilistic model than predict individual edges. For example, consider the problem of linking a new entity into a semantic graph, given only the vector embedding. This task involves adding multiple edges simultaneously, while maintaining structural coherence. Our model is capable of scoring bundles of new edges, and in future work, we plan to explore the possibility of combining M3GM with a search algorithm, to automatically extend existing knowledge graphs by linking in one or more new entities.

We also plan to explore multilingual applications. To some extent, the structural parameters estimated by M3GM are not specific to English: for example, hypernymy cannot be symmetric in any language. If the structural parameters estimated from English WordNet are transferable to other languages, then the combination of M3GM and multilingual word embeddings could facilitate the creation and extension of large-scale semantic resources across many languages (Fellbaum and Vossen, 2012; Bond and Foster, 2013; Lafourcade, 2007).

Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments. We discussed fast motif-counting algorithms with Polo Chau and Oded Green, and received early feedback from Jordan Boyd-Graber, Erica Briscoe, Martin Hyatt, Bryan Leslie Lee, Martha Palmer, and Oren Tsur. This research was funded by the Defense Threat Research Agency under award HDTRA1-15-1-0019.

References

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with

subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.

Francis Bond and Ryan Foster. 2013. Linking and extending an open multilingual wordnet. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1352–1362.

Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2014. A semantic matching energy function for learning with multi-relational data. *Machine Learning*, 94(2):233–259.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2787–2795. Curran Associates, Inc.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12:2493–2537.

Tim Dettmers, Minervini Pasquale, Stenertorp Pontus, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the 32th AAAI Conference on Artificial Intelligence*.

Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.

Christiane Fellbaum and Piek Vossen. 2012. Challenges for a multilingual wordnet. *Language Resources and Evaluation*, 46(2):313–326.

Fan Guo, Steve Hanneke, Wenjie Fu, and Eric P Xing. 2007. Recovering temporally rewiring networks: A model-based approach. In *Proceedings of the 24th international conference on Machine learning*, pages 321–328. ACM.

Kelvin Guu, John Miller, and Percy Liang. 2015. Traversing knowledge graphs in vector space. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 318–327, Lisbon, Portugal. Association for Computational Linguistics.

Paul W. Holland and Samuel Leinhardt. 1981. An exponential family of probability distributions for directed graphs. *Journal of the American Statistical Association*, 76(373):33–50.

Mathieu Lafourcade. 2007. Making people play for lexical acquisition with the jeuxdemots prototype. In *SNLP’07: 7th international symposium on natural language processing*, page 7.

- Zhengdong Lu, Berkant Savas, Wei Tang, and Inderjit S. Dhillon. 2010. [Supervised link prediction using multiple sources](#). In *2010 IEEE International Conference on Data Mining*, pages 923–928.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of International Conference on Learning Representations*.
- Pasquale Minervini, Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. 2017. Adversarial sets for regularising neural link predictors. *arXiv preprint arXiv:1707.07596*.
- Martina Morris, Mark S. Handcock, and David R. Hunter. 2008. Specification of exponential-family random graph models: terms and computational aspects. *Journal of statistical software*, 24(4):1548.
- Arvind Neelakantan, Benjamin Roth, and Andrew McCallum. 2015. Compositional vector space models for knowledge base inference. In *2015 aaai spring symposium series*.
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, Kevin Duh, Manaal Faruqui, Cynthia Gan, Dan Garrette, Yangfeng Ji, Lingpeng Kong, Adhiguna Kuncoro, Gaurav Kumar, Chaitanya Malaviya, Paul Michel, Yusuke Oda, Matthew Richardson, Naomi Saphra, Swabha Swayamdipta, and Pengcheng Yin. 2017. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*.
- Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. 2018. [A novel embedding model for knowledge base completion based on convolutional neural network](#). 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 327–333, New Orleans, Louisiana. Association for Computational Linguistics.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In *ICML*, volume 11, pages 809–816.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543.
- Yuval Pinter, Robert Guthrie, and Jacob Eisenstein. 2017. Mimicking word embeddings using subword rnns. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 102–112.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M. Marlin. 2013. [Relation extraction with matrix factorization and universal schemas](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 74–84, Atlanta, Georgia. Association for Computational Linguistics.
- Sascha Rothe and Hinrich Schütze. 2015. [AutoExtend: Extending Word Embeddings to Embeddings for Synsets and Lexemes](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1793–1803, Beijing, China. Association for Computational Linguistics.
- Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2017. Modeling relational data with graph convolutional networks. *stat*, 1050:17.
- Christian S. Schmid and Bruce A. Desmarais. 2017. Exponential random graph models with big networks: Maximum pseudolikelihood estimation and the parametric bootstrap. *arXiv preprint arXiv:1708.02598*.
- Baoxu Shi and Tim Weninger. 2017. [Proje: Embedding projection for knowledge graph completion](#). In *AAAI*, volume 17, pages 1236–1242.
- Mariano Sigman and Guillermo A. Cecchi. 2002. Global organization of the wordnet lexicon. *Proceedings of the National Academy of Sciences*, 99(3):1742–1747.
- Tom AB Snijders. 2002. Markov chain monte carlo estimation of exponential random graph models. *Journal of Social Structure*, 3(2):1–40.
- Richard Socher, Danqi Chen, Christopher D. Manning, and Andrew Y. Ng. 2013. Reasoning with neural tensor networks for knowledge base completion. In *Advances in Neural Information Processing Systems* 26.
- Mark Steyvers and Joshua B. Tenenbaum. 2005. The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive science*, 29(1):41–78.
- David Strauss and Michael Ikeda. 1990. Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association*, 85(409):204–212.
- Ben Taskar, Carlos Guestrin, and Daphne Koller. 2004. [Max-margin markov networks](#). In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 25–32. MIT Press.
- Kristina Toutanova and Danqi Chen. 2015. [Observed versus latent features for knowledge base and text inference](#). In *Proceedings of the 3rd Workshop on*

Continuous Vector Space Models and their Compositionality, pages 57–66, Beijing, China. Association for Computational Linguistics.

Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoi-fung Poon, Pallavi Choudhury, and Michael Gamon. 2015. Representing text for joint embedding of text and knowledge bases. In *EMNLP*, volume 15, pages 1499–1509.

Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*.