

# INTERPRETING KNOWLEDGE GRAPH RELATION REPRESENTATION FROM WORD EMBEDDINGS

Carl Allen<sup>1\*</sup>, Ivana Balažević<sup>1\*</sup> & Timothy Hospedales<sup>1,2</sup>

<sup>1</sup> University of Edinburgh, UK    <sup>2</sup> Samsung AI Centre, Cambridge, UK  
 {carl.allen, ivana.balazevic, t.hospedales}@ed.ac.uk

## ABSTRACT

Many models learn representations of knowledge graph data by exploiting its low-rank latent structure, encoding known relations between entities and enabling unknown facts to be inferred. To predict whether a relation holds between entities, embeddings are typically compared in the latent space following a relation-specific mapping. Whilst their predictive performance has steadily improved, how such models capture the underlying latent structure of semantic information remains unexplained. Building on recent theoretical understanding of word embeddings, we categorise knowledge graph relations into three types and for each derive explicit requirements of their representations. We show that empirical properties of relation representations and the relative performance of leading knowledge graph representation methods are justified by our analysis.

## 1 INTRODUCTION

Knowledge graphs are large repositories of binary relations between words (or entities) in the form of (*subject, relation, object*) triples. Many models for representing entities and relations have been developed, so that known facts can be recalled and previously unknown facts can be inferred, a task known as *link prediction*. Recent link prediction models (e.g. Bordes et al., 2013; Trouillon et al., 2016; Balažević et al., 2019b) learn entity representations, or *embeddings*, of far lower dimensionality than the number of entities, by capturing latent structure in the data. Relations are typically represented as a mapping from the embedding of a subject entity to those of related object entities. Although the performance of link prediction models has steadily improved for nearly a decade, relatively little is understood of the low-rank latent structure that underpins them, which we address in this work. The outcomes of our analysis can be used to aid and direct future knowledge graph model design.

We start by drawing a parallel between the entity embeddings of knowledge graphs and context-free word embeddings, e.g. as learned by Word2Vec (W2V) (Mikolov et al., 2013a) and GloVe (Pennington et al., 2014). Our motivating premise is that the same latent word features (e.g. meaning(s), tense, grammatical type) give rise to the patterns found in different data sources, i.e. manifesting in word co-occurrence statistics and determining which words relate to which. Different embedding approaches may capture such structure in different ways, but if it is fundamentally the same, an understanding gained from one embedding task (e.g. word embedding) may benefit another (e.g. knowledge graph representation). Furthermore, the relatively limited but accurate data used in knowledge graph representation differs materially from the highly abundant but statistically noisy text data used for word embeddings. As such, theoretically reconciling the two embedding methods may lead to unified and improved embeddings learned jointly from both data sources.

Recent work (Allen & Hospedales, 2019; Allen et al., 2019) theoretically explains how semantic properties are encoded in word embeddings that (approximately) factorise a matrix of *pointwise mutual information* (PMI) from word co-occurrence statistics, as known for W2V (Levy & Goldberg, 2014). *Semantic* relationships between words, specifically similarity, relatedness, paraphrase and analogy, are proven to manifest as linear *geometric* relationships between rows of the PMI matrix (subject to known error terms), of which word embeddings can be considered low-rank projections. This explains, for example, the observations that similar words have similar embeddings and that embeddings of analogous word pairs share a common “vector offset” (e.g. Mikolov et al., 2013b).

---

\*Equal contribution

**Table 1:** Score functions of representative linear link prediction models.  $\mathbf{R} \in \mathbb{R}^{d_e \times d_e}$  and  $\mathbf{r} \in \mathbb{R}^{d_e}$  are the relation matrix and translation vector,  $\mathbf{W} \in \mathbb{R}^{d_e \times d_r \times d_e}$  is the core tensor and  $b_s, b_o \in \mathbb{R}$  are the entity biases.

Model		Linear Subcategory	Score Function
TransE	(Bordes et al., 2013)	additive	$-\ e_s + \mathbf{r} - e_o\ _2^2$
DistMult	(Yang et al., 2015)	multiplicative (diagonal)	$e_s^\top \mathbf{R} e_o$
TuckER	(Balažević et al., 2019b)	multiplicative	$\mathbf{W} \times_1 e_s \times_2 \mathbf{r} \times_3 e_o$
MuRE	(Balažević et al., 2019a)	multiplicative (diagonal) + additive	$-\ \mathbf{R}e_s + \mathbf{r} - e_o\ _2^2 + b_s + b_o$

We extend this insight to identify geometric relationships between PMI-based word embeddings that correspond to other relations, i.e. those of knowledge graphs. Such *relation conditions* define relation-specific mappings between entity embeddings (i.e. *relation representations*) and so provide a “blue-print” for knowledge graph representation models. Analysing the relation representations of leading knowledge graph representation models, we find that various properties, including their relative link prediction performance, accord with predictions based on these relation conditions, supporting the premise that a *common latent structure* is learned by word and knowledge graph embedding models, despite the significant differences between their training data and methodology.

In summary, the key contributions of this work are:

- to use recent understanding of PMI-based word embeddings to derive geometric attributes of a relation representation for it to map subject word embeddings to all related object word embeddings (*relation conditions*), which partition relations into three *types* (§3);
- to show that both per-relation ranking as well as classification performance of leading link prediction models corresponds to the model satisfying the appropriate relation conditions, i.e. how closely its relation representations match the geometric form derived theoretically (§4.1); and
- to show that properties of knowledge graph representation models fit predictions based on relation conditions, e.g. the strength of a relation’s *relatedness* aspect is reflected in the eigenvalues of its relation matrix (§4.2).

## 2 BACKGROUND

**Knowledge graph representation:** Recent knowledge graph models typically represent entities  $e_s, e_o$  as vectors  $e_s, e_o \in \mathbb{R}^{d_e}$ , and relations as transformations in the latent space from subject to object entity embedding, where the dimension  $d_e$  is far lower (e.g. 200) than the number of entities  $n_e$  (e.g.  $> 10^4$ ). Such models are distinguished by their *score function*, which defines (i) the form of the relation transformation, e.g. matrix multiplication and/or vector addition; and (ii) the measure of proximity between a transformed subject embedding and an object embedding, e.g. dot product or Euclidean distance. Score functions can be non-linear (e.g. Dettmers et al., 2018), or linear and sub-categorised as *additive*, *multiplicative* or both. We focus on linear models due to their simplicity and strong performance at link prediction (including state-of-the-art). Table 1 shows the score functions of competitive linear knowledge graph embedding models spanning the sub-categories: TransE (Bordes et al., 2013), DistMult (Yang et al., 2015), TuckER (Balažević et al., 2019b) and MuRE (Balažević et al., 2019a).

*Additive models* apply a relation-specific translation to a subject entity embedding and typically use Euclidean distance to evaluate proximity to object embeddings. A generic additive score function is given by  $\phi(e_s, r, e_o) = -\|e_s + \mathbf{r} - e_o\|_2^2 + b_s + b_o$ . A simple example is TransE, where  $b_s = b_o = 0$ .

*Multiplicative models* have the generic score function  $\phi(e_s, r, e_o) = e_s^\top \mathbf{R} e_o$ , i.e. a bilinear product of the entity embeddings and a relation-specific matrix  $\mathbf{R}$ . DistMult is a simple example with  $\mathbf{R}$  diagonal and so cannot model asymmetric relations (Trouillon et al., 2016). In TuckER, each relation-specific  $\mathbf{R} = \mathbf{W} \times_3 \mathbf{r}$  is a linear combination of  $d_r$  “prototype” relation matrices in a core tensor  $\mathbf{W} \in \mathbb{R}^{d_e \times d_r \times d_e}$  ( $\times_n$  denoting tensor product along mode  $n$ ), facilitating *multi-task learning* across relations.

Some models, e.g. MuRE, combine both multiplicative ( $\mathbf{R}$ ) and additive ( $\mathbf{r}$ ) components.

**Word embedding:** Algorithms such as Word2Vec (Mikolov et al., 2013a) and GloVe (Pennington et al., 2014) generate low-dimensional word embeddings that perform well on downstream tasks (Baroni et al., 2014). Such models predict the context words ( $c_j$ ) observed around a target word ( $w_i$ ) in a text corpus using shallow neural networks. Whilst recent language models (e.g. Devlin et al., 2018; Peters et al., 2018) achieve strong performance using *contextualised* word embeddings, we focus on “context-free” embeddings since knowledge graph entities have no obvious context and, importantly, they offer insight into embedding interpretability.

Levy & Goldberg (2014) show that, for a dictionary of  $n_e$  unique words and embedding dimension  $d_e \ll n_e$ , W2V’s loss function is minimised when its embeddings  $w_i, c_j$  form matrices  $W, C \in \mathbb{R}^{d_e \times n_e}$  that factorise a *pointwise mutual information* (PMI) matrix of word co-occurrence statistics (PMI( $w_i, c_j$ ) =  $\log \frac{P(w_i, c_j)}{P(w_i)P(c_j)}$ ), subject to a shift term. This result relates W2V to earlier count-based embeddings and specifically PMI, which has a history in linguistic analysis (Turney & Pantel, 2010). From its loss function, GloVe can be seen to perform a related factorisation.

Recent work (Allen & Hospedales, 2019; Allen et al., 2019) shows how the semantic relationships of *similarity*, *relatedness*, *paraphrase* and *analogy* are encoded in PMI-based word embeddings by recognising such embeddings as low-rank projections of high dimensional rows of the PMI matrix, termed *PMI vectors*. Those semantic relationships are described in terms of *multiplicative* interactions between co-occurrence probabilities (subject to defined error terms), that correspond to *additive* interactions between (logarithmic) PMI statistics, and hence PMI vectors. Thus, under a sufficiently linear projection, those semantic relationships correspond to linear relationships between word embeddings. Note that although the relative geometry reflecting semantic relationships is preserved, the direct interpretability of dimensions, as in PMI vectors, is lost since the embedding matrices can be arbitrarily scaled/rotated if the other is inversely transformed. We state the relevant semantic relationships on which we build, denoting the set of unique dictionary words by  $\mathcal{E}$ :

- **Paraphrase:** word subsets  $\mathcal{W}, \mathcal{W}^* \subseteq \mathcal{E}$  are said to *paraphrase* if they induce similar distributions over nearby words, i.e.  $p(\mathcal{E}|\mathcal{W}) \approx p(\mathcal{E}|\mathcal{W}^*)$ , e.g.  $\{king\}$  paraphrases  $\{man, royal\}$ .
- **Analogy:** a common example of an *analogy* is “*woman* is to *queen* as *man* is to *king*” and can be defined as any set of word pairs  $\{(w_i, w_i^*)\}_{i \in \mathcal{I}}$  for which it is semantically meaningful to say “ $w_a$  is to  $w_a^*$  as  $w_b$  is to  $w_b^*$ ”  $\forall a, b \in \mathcal{I}$ .


Where one word subset paraphrases another, the sums of their embeddings are shown to be equal (subject to the independence of words within each set), e.g.  $w_{king} \approx w_{man} + w_{royal}$ . An interesting connection is established between the two semantic relationships: a set of word pairs  $\mathcal{A} = \{(w_a, w_a^*), (w_b, w_b^*)\}$  is an analogy if  $\{w_a, w_b^*\}$  paraphrases  $\{w_a^*, w_b\}$ , in which case the embeddings satisfy  $w_{a^*} - w_a \approx w_{b^*} - w_b$  (“vector offset”).

### 3 FROM ANALOGIES TO KNOWLEDGE GRAPH RELATIONS

Analogies from the field of word embeddings are our starting point for developing a theoretical basis for representing knowledge graph relations. The relevance of analogies stems from the observation that for an analogy to hold (see §2), its word pairs, e.g.  $\{(man, king), (woman, queen), (girl, princess)\}$ , must be *related* in the same way, comparably to subject-object entity pairs under a common knowledge graph relation. Our aim is to develop the understanding of PMI-based word embeddings (henceforth *word embeddings*), to identify the mathematical properties necessary for a relation representation to map subject word embeddings to all related object word embeddings.

Considering the paraphrasing word sets  $\{king\}$  and  $\{man, royal\}$  corresponding to the word embedding relationship  $w_{king} \approx w_{man} + w_{royal}$  (§2), *royal* can be interpreted as the semantic difference between *man* and *king*, fitting intuitively with the relationship  $w_{royal} \approx w_{king} - w_{man}$ . Fundamentally, this relationship holds because the difference between words that co-occur (i.e. occur more frequently than if independent) with *king* and those that co-occur with *man*, reflects those words that co-occur with *royal*. We refer to this difference in co-occurrence distribution as a “context shift”, from *man* (subject) to *king* (object). Allen & Hospedales (2019) effectively show that where multiple word pairs share a common context shift, they form an analogy whose embeddings satisfy the vector offset relationship. This result seems obvious where the context shift mirrors an identifiable word, the embedding of which is approximated by the common vector offset, e.g. *queen* and *woman* are related by the same context shift, i.e.  $w_{queen} \approx w_{woman} + w_{royal}$ , thus  $w_{queen} - w_{woman} \approx w_{king} - w_{man}$ . However, the same result holds, i.e. an analogy is formed with a common vector offset between embeddings, for an arbitrary (common) context shift that may reflect no particular word. Importantly, these context shift relations evidence a case in which it is known how a relation can be represented, i.e. by an additive vector (comparable to TransE) if entities are represented by word embeddings. More generally, this provides an interpretable foothold into relation representation.

Note that not all sets of word pairs considered analogies exhibit a clear context shift relation, e.g. in the analogy  $\{(car, engine), (bus, seats)\}$ , the difference between words co-occurring with *engine* and *car* is not expected to reflect the corresponding difference between *bus* and *seats*. This illustrates how



**Figure 1:** Relationships between PMI vectors (black rectangles) of subject/object words for different relation types. PMI vectors capture co-occurrence with every dictionary word: strong associations ( $\text{PMI} > 0$ ) are shaded (blue define the relation, grey are random other associations); red dash = *relatedness*; black dash = *context sets*. analogies are a loosely defined concept, e.g. their implicit relation may be semantic or syntactic, with several sub-categories of each (e.g. see Gladkova et al. (2016)). The same is readily observed for the relations of knowledge graphs. This likely explains the observed variability in “solving” analogies by use of vector offset (e.g. Köper et al., 2015; Karpinska et al., 2018; Gladkova et al., 2016) and suggests that further consideration is required to represent relations (or solve analogies) in general.

We have seen that the existence of a context shift relation between a subject and object word implies a (relation-specific) geometric relationship between word embeddings, thus the latter provides a *necessary condition for the relation to hold*. We refer to this as a “relation condition” and aim to identify relation conditions for other classes of relation. Once identified, relation conditions define a mapping from subject embeddings to all related object embeddings, by which related entities might be identified with a proximity measure (e.g. Euclidean distance or dot product). This is the precise aim of a knowledge graph representation model, but loss functions are typically developed heuristically. Given the existence of many representation models, we can verify identified relation conditions by contrasting the per-relation performance of various models with the extent to which their loss function reflects the appropriate relation conditions. Note that since relation conditions are necessary rather than sufficient, they do not guarantee a relation holds, i.e. false positives may arise.

Whilst we seek to establish relation conditions based on PMI word embeddings, the data used to train knowledge graph embeddings differs significantly to the text data used by word embeddings, and the relevance of conditions ultimately based on PMI statistics may seem questionable. However, where a knowledge graph representation model implements relation conditions and measures proximity between embeddings, the parameters of word embeddings necessarily provide a *potential* solution that minimises the loss function. Many equivalent solutions may exist due to symmetry as typical for neural network architectures. We now define relation types and identify their relation conditions (underlined); we then consider the completeness of this categorisation.

- **Similarity:** Semantically similar words induce similar distributions over the words they co-occur with. Thus their PMI vectors and word embeddings are similar (Fig 1a).
- **Relatedness:** The relatedness of two words can be considered in terms of the words  $\mathcal{S} \subseteq \mathcal{E}$  with which both co-occur similarly.  $\mathcal{S}$  defines the *nature* of relatedness, e.g. *milk* and *cheese* are related by  $\mathcal{S} = \{\text{dairy, breakfast, ...}\}$ ; and  $|\mathcal{S}|$  reflects the *strength* of relatedness. Since PMI vector components corresponding to  $\mathcal{S}$  are similar (Fig 1b), embeddings of  $\mathcal{S}$ -related words have similar components in the subspace  $\mathbb{V}_{\mathcal{S}}$  that spans the projected PMI vector dimensions corresponding to  $\mathcal{S}$ . The rank of  $\mathbb{V}_{\mathcal{S}}$  is thus anticipated to reflect relatedness strength. Relatedness can be seen as a weaker and more variable generalisation of similarity, its limiting case where  $\mathcal{S} = \mathcal{E}$ , hence  $\text{rank}(\mathbb{V}_{\mathcal{S}}) = d_e$ .
- **Context-shift:** As discussed above, words related by a common difference between their distributions of co-occurring words, defined as *context-shifts*, share a common vector offset between word embeddings. Context might be considered *added* (e.g. *man* to *king*), termed a **specialisation** (Fig 1c), *subtracted* (e.g. *king* to *man*) or both (Fig 1d). These relations are 1-to-1 (subject to synonyms) and include an aspect of *relatedness* due to the word associations in common. Note that, specialisations include hyponyms/hypernyms and context shifts include meronyms.
- **Generalised context-shift:** Context-shift relations generalise to 1-to-many, many-to-1 and many-to-many relations where the added/subtracted context may be from a (relation-specific) *context set* (Fig 1e), e.g. *any* city or *anything* bigger. The potential scope and size of context sets adds variability to these relations. The limiting case in which the context set is “small” reduces to a 1-to-1 context-shift (above) and the embedding difference is a known vector offset. In the limiting case of a “large” context set, the added/subtracted context is essentially unrestricted such that only the relatedness aspect of the relation, and thus a common subspace component of embeddings, is fixed.

**Categorisation completeness:** Taking intuition from Fig 1 and considering PMI vectors as *sets of word features*, these relation types can be interpreted as set operations: similarity as set equality; relatedness as subset equality; and context-shift as a relation-specific set difference. Since for any relation each feature must either remain unchanged (relatedness), change (context shift) or else be irrelevant, we conjecture that the above relation types give a complete partition of semantic relations.

**Table 2:** Categorisation of WN18RR relations.

Type	Relation	Examples ( <i>subject entity, object entity</i> )
R	verb_group	( <i>trim_down_VB_1, cut_VB_35</i> ), ( <i>hatch_VB_1, incubate_VB_2</i> )
	derivationally_related_form	( <i>lodge_VB_4, accommodation_NN_4</i> ), ( <i>question_NN_1, inquire_VB_1</i> )
	also_see	( <i>clean_JJ_1, tidy_JJ_1</i> ), ( <i>ram_VB_2, screw_VB_3</i> )
S	hypernym	( <i>land_reform_NN_1, reform_NN_1</i> ), ( <i>prickle-weed_NN_1, herbaceous_plant_NN_1</i> )
	instance_hypernym	( <i>yellowstone_river_NN_1, river_NN_1</i> ), ( <i>leipzig_NN_1, urban_center_NN_1</i> )
C	member_of_domain_usage	( <i>colloquialism_NN_1, figure_VB_5</i> ), ( <i>plural_form_NN_1, authority_NN_2</i> )
	member_of_domain_region	( <i>rome_NN_1, gladiator_NN_1</i> ), ( <i>usa_NN_1, multiple_voting_NN_1</i> )
	member_meronym	( <i>south_NN_2, sunshine_state_NN_1</i> ), ( <i>genus_carya_NN_1, pecan_tree_NN_1</i> )
	has_part	( <i>aircraft_NN_1, cabin_NN_3</i> ), ( <i>morocco_NN_1, atlas_mountains_NN_1</i> )
	synset_domain_topic_of	( <i>quark_NN_1, physics_NN_1</i> ), ( <i>harmonize_VB_3, music_NN_4</i> )

### 3.1 CATEGORISING REAL KNOWLEDGE GRAPH RELATIONS

Analysing the relations of popular knowledge graph datasets, we observe that they indeed imply (i) a relatedness aspect reflecting a common theme (e.g. both entities are animals or geographic terms); and (ii) contextual themes specific to the subject and/or object entities. Further, relations fall under a hierarchy of three *relation types*: highly related (**R**); generalised specialisation (**S**); and generalised context-shift (**C**). As above, “generalised” indicates that context differences are not restricted to be 1-1. From Fig 1, it can be seen that type R relations are a special case of S, which are a special case of C. Thus type C encompasses all considered relations. Whilst there are many ways to classify relations, e.g. by hierarchy, transitivity, the proposed relation conditions delineate relations by the required mathematical form (and complexity) of their representation. Table 2 shows a categorisation of the relations of the WN18RR dataset (Dettmers et al., 2018) comprising 11 relations and 40,943 entities.<sup>1</sup> An explanation for the category assignment is in Appx. A. Analysing the commonly used FB15k-237 dataset (Toutanova et al., 2015) reveals relations to be almost exclusively of type C, precluding a contrast of performance per relation type and hence that dataset is omitted from our analysis. Instead, we categorise a random subsample of 12 relations from the NELL-995 dataset (Xiong et al., 2017), containing 75,492 entities and 200 relations (see Tables 8 and 9 in Appx. B).

### 3.2 RELATIONS AS MAPPINGS BETWEEN EMBEDDINGS

Given the relation conditions of a relation type, we now consider mappings that satisfy them and thereby loss functions able to identify relations of each type, evaluating proximity between mapped entity embeddings by dot product or Euclidean distance. We then contrast our theoretically derived loss functions, specific to a relation type, with those of several knowledge graph models (Table 1) to predict identifiable properties and the relative performance of different knowledge graph models for each relation type.

**R:** Identifying  $\mathcal{S}$ -relatedness requires testing both entity embeddings  $e_s, e_o$  for a common subspace component  $\mathbb{V}_{\mathcal{S}}$ , which can be achieved by projecting both embeddings onto  $\mathbb{V}_{\mathcal{S}}$  and comparing their images. Projection requires multiplication by a matrix  $P_r \in \mathbb{R}^{d \times d}$  and cannot be achieved additively, except in the trivial limiting case of similarity ( $P_r = I$ ) when  $r \approx 0$  can be added.

Comparison by dot product gives  $(P_r e_s)^\top (P_r e_o) = e_s^\top P_r^\top P_r e_o = e_s^\top M_r e_o$  (for relation-specific symmetric  $M_r = P_r^\top P_r$ ). Euclidean distance gives  $\|P_r e_s - P_r e_o\|^2 = (e_s - e_o)^\top M_r (e_s - e_o) = \|P_r e_s\|^2 - 2e_s^\top M_r e_o + \|P_r e_o\|^2$ .

**S/C:** The relation conditions require testing for both  $\mathcal{S}$ -relatedness and relation-specific entity component(s) ( $v_r^s, v_r^o$ ). This is achieved by (i) multiplying both entity embeddings by a relation-specific projection matrix  $P_r$  that projects onto the subspace that spans the low-rank projection of dimensions corresponding to  $\mathcal{S}$ ,  $v_r^s$  and  $v_r^o$  (i.e. testing for  $\mathcal{S}$ -relatedness while preserving relation-specific entity components); and (ii) adding a relation-specific vector  $r = v_r^o - v_r^s$  to the transformed subject entity embeddings.

<sup>1</sup>We omit the relation “similar\_to” since its instances have no discernible structure, and only 3 occur in the test set, all of which are the inverse of a training example and trivial to predict.

Comparing the transformed entity embeddings by dot product equates to  $(P_r e_s + r)^\top P_r e_o$ ; and by Euclidean distance to  $\|P_r e_s + r - P_r e_o\|^2 = \|P_r e_s + r\|^2 - 2(P_r e_s + r)^\top P_r e_o + \|P_r e_o\|^2$  (cf MuRE:  $\|R e_s + r - e_o\|^2$ ).

Contrasting these theoretically derived loss functions with those of knowledge graph models (Table 1), we make the following predictions:

- P1:** The ability to learn the representation of a relation is expected to reflect:
- (a) the complexity of its type (R<S<C) independently of model choice; and
  - (b) whether relation conditions (e.g. additive/multiplicative interactions) are met by the model.
- P2:** Knowledge graph relation representations reflect the following type-specific properties:
- (a) relation matrices for relatedness (type R) relations are highly symmetric;
  - (b) offset vectors for relatedness relations have low norm; and
  - (c) as a proxy to the rank of  $\mathbb{V}_S$ , the eigenvalues of a relation matrix reflect relatedness strength.

To elaborate, our core prediction P1(b) anticipates that: (i) additive-only models (e.g. TransE) are not suited to identifying the relatedness aspect of relations, except in limiting cases of similarity, requiring a zero vector); (ii) multiplicative-only models (e.g. DistMult) should perform well on type R relations, but are not suited to identifying entity-specific features of type S/C (an asymmetric relation matrix in TuckER may help compensate); and (iii) the loss function of MuRE closely resembles that derived for type C relations, which generalise all others, and is thus expected to perform best overall.

## 4 EVIDENCE LINKING KNOWLEDGE GRAPH AND WORD EMBEDDINGS

We test whether the predictions P1 and P2, made on the basis of word embeddings, apply to knowledge graph relations by analysing the performance and properties of competitive knowledge graph models. We compare TransE, DistMult, TuckER and MuRE, which entail different forms of relation representation, on all WN18RR relations and a similar number of NELL-995 relations (spanning all relation types). All models have a comparable number of free parameters.

Since for TransE, the logistic sigmoid cannot be applied to the score function to give a probabilistic interpretation comparable to other models, for fair comparison we include MuRE<sub>I</sub>, a constrained variant of MuRE with  $R_s = R_o = I$ , as a proxy to TransE. Implementation details are included in Appx. D. For evaluation, we generate  $2n_e$  evaluation triples for each test triple ( $n_e = |\mathcal{E}|$  denoting the number of entities) by fixing the subject entity  $e_s$  and relation  $r$  and replacing the object entity  $e_o$  with each entity in turn and then keeping  $e_o$  and  $r$  fixed and varying  $e_s$ . Each model’s scores for the evaluation triples are ranked to give the standard metric Hits@10 (Bordes et al., 2013), i.e. the fraction of times a true triple appears in the top 10 ranked evaluation triples.

### 4.1 P1: JUSTIFYING THE RELATIVE PERFORMANCE OF KNOWLEDGE GRAPH MODELS

**Ranking performance:** Tables 3 and 4 report Hits@10 for each relation and include the relation type as well as known confounding influences: percentage of relation instances in the training and test sets (approximately equal), the actual number of instances in the test set (causing some results to be highly granular), Krackhardt hierarchy score (see Appx. E) (Krackhardt, 2014; Balažević et al., 2019a) and maximum and average shortest path between any two related nodes. A further confounding effect is dependence between relations: Lacroix et al. (2018) and Balažević et al. (2019b) independently show that constraining the rank of relation representations is beneficial for datasets with many relations due to *multi-task learning*, particularly when the number of instances per relation is low. This is expected to benefit TuckER on the NELL-995 dataset (200 relations).

As predicted in P1(a), all models tend to perform best at type R relations, with a clear performance gap to other relation types. Also, performance on type S relations appears higher in general than type C. In accordance with P1(b), additive-only models (TransE, MuRE<sub>I</sub>) perform worst on average since all relation types involve (multiplicative) relatedness. Best performance is achieved on type R relations, which can be represented by a small/zero additive vector. Multiplicative-only DistMult performs well, sometimes best, on type R relations, fitting expectation as it can represent those relations and has no inessential parameters, e.g. that may overfit to noise, which may explain instances where MuRE performs slightly worse. As expected, MuRE, performs best overall (particularly on WN18RR), and most strongly on S and C type relations, predicted to require both multiplicative and additive components. Comparable performance of TuckER on NELL-995 may be explained by its multi-task learning ability.

**Table 3: Hits@10 per relation on WN18RR.**

Relation Name	Type	%	#	Khs	Max/Avg Path	TransE	MuRE <sub>I</sub>	DistMult	TuckER	MuRE
verb_group	R	1%	39	0.00	-	-	0.87	0.95	<b>0.97</b>	<b>0.97</b>
derivationally_related_form	R	34%	1074	0.04	-	-	0.93	0.96	0.96	<b>0.97</b>
also_see	R	2%	56	0.24	44	15.2	0.59	<b>0.73</b>	0.67	<b>0.73</b>
instance_hypernym	S	4%	122	1.00	3	1.0	0.22	0.52	0.47	<b>0.54</b>
synset_domain_topic_of	C	4%	114	0.99	3	1.1	0.19	0.43	0.42	<b>0.53</b>
member_of_domain_usage	C	1%	24	1.00	2	1.0	0.42	0.42	0.48	<b>0.50</b>
member_of_domain_region	C	1%	26	1.00	2	1.0	0.35	0.40	0.40	<b>0.46</b>
member_meronym	C	8%	253	1.00	10	3.9	0.04	0.38	0.30	<b>0.39</b>
has_part	C	6%	172	1.00	13	2.2	0.04	0.31	0.28	<b>0.35</b>
hypernym	S	40%	1251	0.99	18	4.5	0.02	0.20	0.19	<b>0.28</b>
all		100%	3134				0.38	0.52	0.51	<b>0.57</b>

**Table 4: Hits@10 per relation on NELL-995.**

Relation Name	Type	%	#	Khs	Max/Avg Path	TransE	MuRE <sub>I</sub>	DistMult	TuckER	MuRE
teamploysagainsteam	R	2%	243	0.11	10	3.5	0.76	0.84	<b>0.90</b>	0.89
clothingtogowithclothing	R	1%	132	0.17	5	2.6	0.72	0.80	<b>0.88</b>	0.85
professionistypeofprofession	S	1%	143	0.38	7	2.5	0.37	0.55	0.62	<b>0.66</b>
animalistypeofanimal	S	1%	103	0.68	9	3.1	0.50	0.56	0.64	<b>0.68</b>
athleteplayssport	C	1%	113	1.00	1	1.0	0.54	0.58	0.58	<b>0.64</b>
chemististypeofchemical	S	1%	115	0.53	6	3.0	0.23	0.43	0.49	<b>0.60</b>
itemfoundinroom	C	2%	162	1.00	1	1.0	0.39	0.57	0.53	<b>0.59</b>
agentcollaborateswithagent	R	1%	119	0.51	14	4.7	0.44	0.58	<b>0.64</b>	0.61
bodypartcontainsbodypart	C	1%	103	0.60	7	3.2	0.30	0.38	0.54	<b>0.58</b>
atdate	C	10%	967	0.99	4	1.1	0.41	0.50	0.48	<b>0.52</b>
locationlocatedwithinlocation	C	2%	157	1.00	6	1.9	0.35	0.37	0.46	<b>0.48</b>
atlocation	C	1%	294	0.99	6	1.4	0.22	0.33	0.39	<b>0.44</b>
all		100%	20000				0.36	0.48	0.51	<b>0.52</b>

Other anomalous results also closely align with confounding factors. For example, all models perform poorly on the *hypernym* relation, despite it having a relative abundance of training data (40% of all instances), which may be explained by its *hierarchical* nature ( $Khs \approx 1$  and long paths). The same may explain the reduced performance on relations *also\_see* and *agentcollaborateswithagent*. As found previously (Balažević et al., 2019a), none of the models considered are well suited to modelling hierarchical structures. We also note that the percentage of training instances of a relation is not a dominant factor on performance, as would be expected if all relations could be equally represented.

**Classification performance:** We further evaluate whether P1 holds when comparing knowledge graph models by classification accuracy on WN18RR. Independent predictions of whether a given triple is true or false are not commonly evaluated, instead metrics such as mean reciprocal rank and Hits@ $k$  are reported that compare the prediction of a test triple against all evaluation triples. Not only is this computationally costly, the evaluation is flawed if an entity is related to  $l > k$  others ( $k$  is often 1 or 3). A correct prediction validly falling within the top  $l$  but not the top  $k$  would appear incorrect under the metric. Some recent works also note the importance of standalone predictions (Speranskaya et al., 2020; Pezeshkpour et al., 2020).

Since for each relation there are  $n_e^2$  possible entity-entity relationships, we sub-sample by computing predictions only for those  $(e_s, r, e_o)$  triples for which the  $e_s, r$  pairs appear in the test set. We split positive predictions ( $\sigma(\phi(e_s, r, e_o)) > 0.5$ ) between (i) *known truths* – training or test/validation instances; and (ii) *other*, the truth of which is not known. We then compute per-relation accuracy over the true training instances (“train”) and true test/validation instances (“test”); and the average number of “other” triples predicted true per  $e_s, r$  pair. Table 5 shows results for MuRE<sub>I</sub>, DistMult, TuckER and MuRE. All models achieve near perfect training accuracy. The additive-multiplicative MuRE gives best test set performance, followed (surprisingly) closely by MuRE<sub>I</sub>, with multiplicative models (DistMult and TuckER) performing poorly on all but type R relations in line with P1(b), with near-zero performance on most type S/C relations. Since the ground truth labels for “other” triples predicted to be true are not in the dataset, we analyse a sample of “other” true predictions for one relation of each type (see Appx. G). From this, we estimate that TuckER is relatively accurate but pessimistic ( $\sim 0.3$  correct of the 0.5 predictions  $\approx 60\%$ ), MuRE<sub>I</sub> is optimistic but inaccurate ( $\sim 2.3$  of  $7.5 \approx 31\%$ ), whereas MuRE is both optimistic and accurate ( $\sim 1.1$  of  $1.5 \approx 73\%$ ).

**Summary:** Our analysis identifies the best performing model per relation type as predicted by P1(b): multiplicative-only DistMult for type R, additive-multiplicative MuRE for types S/C; providing a basis for *dataset-dependent model selection*. The per-relation insight into where models perform

**Table 5:** Per relation prediction accuracy for MuRE<sub>I</sub> (M<sub>I</sub>), (D)istMult, (T)uckER and (M)uRE (WN18RR).

Relation Name	Type	# <i>train</i>	# <i>test</i>	Accuracy (train)				Accuracy (test)				# Other "True"			
				M <sub>I</sub>	D	T	M	M <sub>I</sub>	D	T	M	M <sub>I</sub>	D	T	M
verb_group	R	15	39	1.00	1.00	1.00	1.00	0.97	0.97	0.97	0.97	8.3	1.7	0.9	2.7
derivationally_related_form	R	1714	1127	1.00	1.00	1.00	1.00	0.96	0.94	0.95	0.95	8.8	0.5	0.6	1.7
also_see	R	95	61	1.00	1.00	1.00	1.00	0.64	0.64	0.61	0.59	7.9	1.6	0.9	1.9
instance_hyponym	S	52	122	1.00	1.00	1.00	1.00	0.32	0.32	0.23	0.43	1.3	0.4	0.3	0.9
member_of_domain_usage	C	545	43	0.98	1.00	1.00	1.00	0.02	0.00	0.02	0.00	1.5	0.6	0.0	0.3
member_of_domain_region	C	543	42	0.88	0.89	1.00	0.93	0.02	0.02	0.00	0.02	1.0	0.4	0.8	0.7
synset_domain_topic_of	C	13	115	1.00	1.00	1.00	1.00	0.42	0.10	0.14	0.47	0.7	0.6	0.1	0.2
member_meronym	C	1402	307	1.00	1.00	1.00	1.00	0.22	0.02	0.01	0.22	7.9	3.4	1.5	5.6
has_part	C	848	196	1.00	1.00	1.00	1.00	0.24	0.05	0.09	0.22	7.1	2.4	1.3	3.9
hyponym	S	57	1254	1.00	1.00	1.00	1.00	0.15	0.02	0.02	0.22	3.7	1.2	0.0	1.7
all		5284	3306	0.99	0.99	1.00	0.99	0.47	0.37	0.37	0.50	5.9	1.2	0.5	2.1

poorly, e.g. hierarchical or type C relations, can be used to aid and direct future model design. Analysis of the classification performance: (i) shows that MuRE is the most reliable fact prediction model; and (ii) emphasises the poorer ability of multiplicative-only models to represent S/C relations.


## 4.2 P2: PROPERTIES OF RELATION REPRESENTATION

**P2(a)-(b):** Table 6 shows the symmetry score ( $\in [-1, 1]$ ) indicating perfect anti-symmetry to symmetry; see Appx. F) for the relation matrix of TuckER and the norm of relation vectors of TransE, MuRE<sub>I</sub> and MuRE on the WN18RR dataset. As expected, type R relations have materially higher symmetry than both other relation types, fitting the prediction of how TuckER compensates for having no additive component. All additive models learn relation vectors of a noticeably lower norm for type R relations, which in the limiting case (similarity) require no additive component, than for types S or C.

**P2(c):** Fig 2 shows eigenvalue magnitudes (scaled relative to the largest and ordered) of relation-specific matrices  $\mathbf{R}$  of MuRE, labelled by relation type, as predicted to reflect the strength of a relation’s *relatedness* aspect. As expected, values are highest for type R relations. For relation types S and C the profiles are more varied, supporting the understanding that relatedness of such relations is highly variable, both in its nature (components of  $\mathcal{S}$ ) and strength (cardinality of  $\mathcal{S}$ ).

**Table 6:** Relation matrix symmetry score [-1.1] for TuckER; and relation vector norm for TransE, MuRE<sub>I</sub> and MuRE (WN18RR).

Relation	Type	Symmetry Score	Vector Norm		
			TuckER	TransE	MuRE <sub>I</sub>
verb_group	R	0.52	5.65	0.76	0.89
derivationally_related_form	R	0.54	2.98	0.45	0.69
also_see	R	0.50	7.20	0.97	0.97
instance_hyponym	S	0.13	18.26	2.98	1.88
member_of_domain_usage	C	0.10	11.24	3.18	1.88
member_of_domain_region	C	0.06	12.52	3.07	2.11
synset_domain_topic_of	C	0.12	23.29	2.65	1.52
member_meronym	C	0.12	4.97	1.91	1.97
has_part	C	0.13	6.44	1.69	1.25
hyponym	S	0.04	9.64	1.53	1.03

**Figure 2:** Eigenvalue magnitudes of relation-specific matrices  $\mathbf{R}$  for MuRE by relation type (WN18RR).

## 5 CONCLUSION

Many low-rank knowledge graph representation models have been developed, yet little is known of the latent structure they learn. We build on recent understanding of PMI-based word embeddings to theoretically establish a set of geometric properties of relation representations (relation conditions) required to map PMI-based word embeddings of subject entities to related object entities under knowledge graph relations. These conditions partition relations into three types and provide a basis to consider the loss functions of existing knowledge graph models. Models that satisfy the relation conditions of a particular type have a known set of model parameters that minimise the loss function, i.e. the parameters of PMI embeddings, together with potentially many equivalent solutions. We show that the better a model’s architecture satisfies a relation’s conditions, the better its performance at link prediction, evaluated under both rank-based metrics and accuracy. Overall, we generalise recent theoretical understanding of how particular semantic relations, e.g. similarity and analogy, are encoded between PMI-based word embeddings to the general relations of knowledge graphs. In doing so, we provide evidence in support of our initial premise: that common latent structure is exploited by both PMI-based word embeddings (e.g. W2V) and knowledge graph representation.



#### ACKNOWLEDGEMENTS

Carl Allen and Ivana Balažević were supported by the Centre for Doctoral Training in Data Science, funded by EPSRC (grant EP/L016427/1) and the University of Edinburgh.

#### REFERENCES

- Carl Allen and Timothy Hospedales. Analogies Explained: Towards Understanding Word Embeddings. In *International Conference on Machine Learning*, 2019.
- Carl Allen, Ivana Balažević, and Timothy Hospedales. What the Vec? Towards Probabilistically Grounded Embeddings. In *Advances in Neural Information Processing Systems*, 2019.
- Ivana Balažević, Carl Allen, and Timothy M Hospedales. Multi-relational Poincaré Graph Embeddings. In *Advances in Neural Information Processing Systems*, 2019a.
- Ivana Balažević, Carl Allen, and Timothy M Hospedales. TuckER: Tensor Factorization for Knowledge Graph Completion. In *Empirical Methods in Natural Language Processing*, 2019b.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. Don't Count, Predict! A Systematic Comparison of Context-Counting vs. Context-Predicting Semantic Vectors. In *Association for Computational Linguistics*, 2014.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems*, 2013.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2D Knowledge Graph Embeddings. In *Association for the Advancement of Artificial Intelligence*, 2018.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*, 2018.
- Anna Gladkova, Aleksandr Drozd, and Satoshi Matsuoka. Analogy-based Detection of Morphological and Semantic Relations With Word Embeddings: What Works and What Doesn't. In *NAACL Student Research Workshop*, 2016.
- Lawrence J Hubert and Frank B Baker. Evaluating the Symmetry of a Proximity Matrix. *Quality & Quantity*, 13(1):77–84, 1979.
- Marzena Karpinska, Bofang Li, Anna Rogers, and Aleksandr Drozd. Subcharacter Information in Japanese Embeddings: When Is It Worth It? In *Workshop on the Relevance of Linguistic Structure in Neural Architectures for NLP*, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.
- Maximilian Köper, Christian Scheible, and Sabine Schulte im Walde. Multilingual reliability and "semantic" structure of continuous word spaces. In *International Conference on Computational Semantics*, 2015.
- David Krackhardt. Graph Theoretical Dimensions of Informal Organizations. In *Computational Organization Theory*. Psychology Press, 2014.
- Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. Canonical Tensor Decomposition for Knowledge Base Completion. In *International Conference on Machine Learning*, 2018.
- Omer Levy and Yoav Goldberg. Neural Word embedding as Implicit Matrix Factorization. In *Advances in Neural Information Processing Systems*, 2014.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*, 2013a.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. Linguistic regularities in continuous space word representations. In *North American Chapter of the Association for Computational Linguistics*, 2013b.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing*, 2014.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep Contextualized Word Representations. In *North American Chapter of the Association for Computational Linguistics*, 2018.
- Pouya Pezeshkpour, Yifan Tian, and Sameer Singh. Revisiting Evaluation of Knowledge Base Completion Models. In *Automated Knowledge Base Construction*, 2020.
- Marina Speranskaya, Martin Schmitt, and Benjamin Roth. Ranking vs. Classifying: Measuring Knowledge Base Completion Quality. In *Automated Knowledge Base Construction*, 2020.
- Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoifung Poon, Pallavi Choudhury, and Michael Gamon. Representing Text for Joint Embedding of Text and Knowledge Bases. In *Empirical Methods in Natural Language Processing*, 2015.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex Embeddings for Simple Link Prediction. In *International Conference on Machine Learning*, 2016.
- Peter D Turney and Patrick Pantel. From Frequency to Meaning: Vector Space Models of Semantics. *Journal of Artificial Intelligence Research*, 37:141–188, 2010.
- Wenhan Xiong, Thien Hoang, and William Yang Wang. DeepPath: A Reinforcement Learning Method for Knowledge Graph Reasoning. In *Empirical Methods in Natural Language Processing*, 2017.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *International Conference on Learning Representations*, 2015.

## A CATEGORISING WN18RR RELATIONS

Table 7 describes how each WN18RR relation was assigned to its respective category.

**Table 7:** Explanation for the WN18RR relation category assignment.

Type	Relation	Relatedness	Subject Specifics	Object Specifics
R	verb_group	both verbs; similar in meaning	-	-
	derivationally_related_form	different syntactic categories; semantically related	-	-
	also_see	semantically similar	-	-
S	hypernym	semantically similar	instance of the object	-
	instance_hyponym	semantically similar	instance of the object	-
C	member_of_domain_usage	loosely semantically related (word usage features)	usage descriptor	broad feature set
	member_of_domain_region	loosely semantically related (regional features)	region descriptor	broad feature set
	member_meronym	semantically related	collection of objects	part of the subject
	has_part	semantically related	collection of objects	part of the subject
	synset_domain_topic_of	semantically related	broad feature set	domain descriptor

## B CATEGORISING NELL-995 RELATIONS

Categorisation of NELL-995 relations and the explanation for the category assignment of are shown in Tables 8 and 9 respectively.

**Table 8:** Categorisation of NELL-995 relations.

Type	Relation	Examples ( <i>subject entity, object entity</i> )
R	teamploysagainstteam	<i>(rangers, mariners), (phillies, tampa_bay_rays)</i>
	clothingtogowithclothing	<i>(shirts, trousers), (shoes, black_shirt)</i>
	agentcollaborateswithagent	<i>(white_stripes, jack_white), (barack_obama, hillary_clinton)</i>
S	professionistypeofprofession	<i>(trial_lawyers, attorneys), (engineers, experts)</i>
	animalistypeofanimal	<i>(cats, small_animals), (chickens, livestock)</i>
	chemicalistypeofchemical	<i>(moisture, gas), (oxide, materials)</i>
C	athleteplayssport	<i>(joe_smith, baseball), (chris_cooley, football)</i>
	itemfoundinroom	<i>(bed, den), (refrigerator, kitchen_area)</i>
	bodypartcontainsbodypart	<i>(system002, eyes), (blood, left_ventricle)</i>
	atdate	<i>(scotland, n2009), (wto, n2003)</i>
	locationlocatedwithinlocation	<i>(medellin, colombia), (jackson, wyoming)</i>
atlocation	<i>(ogunquin, maine), (palmer_lake, colorado)</i>	

**Table 9:** Explanation for the NELL-995 relation category assignment.

Type	Relation	Relatedness	Subject Specifics	Object Specifics
R	teamploysagainstteam	both sport teams	-	-
	clothingtogowithclothing	both items of clothing that go together	-	-
	agentcollaborateswithagent	both people or companies; related industries	-	-
S	professionistypeofprofession	semantically related (both profession types)	instance of the object	-
	animalistypeofanimal	semantically related (both animals)	instance of the object	-
	chemicalistypeofchemical	semantically related (both chemicals)	instance of the object	-
C	athleteplayssport	semantically related (sports features)	athlete descriptor	sport descriptor
	itemfoundinroom	semantically related (room/furniture features)	item descriptor	room descriptor
	bodypartcontainsbodypart	emantically related (specific body part features)	collection of objects	part of the subject
	atdate	loosely semantically related (start date features)	broad feature set	date descriptor
	locationlocatedwithinlocation	semantically related (geographical features)	part of the subject	collection of objects
atlocation	semantically related (geographical features)	part of the subject	collection of objects	

## C SPLITTING THE NELL-995 DATASET

The test set of NELL-995 created by Xiong et al. (2017) contains only 10 out of 200 relations present in the training set. To ensure a fair representation of all training set relations in the validation and test sets, we create new validation and test set splits by combining the initial validation and test sets with the training set and randomly selecting 10,000 triples each from the combined dataset.

## D IMPLEMENTATION DETAILS

All algorithms are re-implemented in PyTorch with the Adam optimizer (Kingma & Ba, 2015) that minimises binary cross-entropy loss, using hyper-parameters that work well for all models (learning rate: 0.001, batch size: 128, number of negative samples: 50). Entity and relation embedding dimensionality is set to  $d_e = d_r = 200$  for all models except TUCKER, for which  $d_r = 30$  (Balažević et al., 2019b).

## E KRACKHARDT HIERARCHY SCORE

The Krackhardt hierarchy score measures the proportion of node pairs  $(x, y)$  where there exists a directed path  $x \rightarrow y$ , but not  $y \rightarrow x$ ; and it takes a value of one for all directed acyclic graphs, and zero for cycles and cliques (Krackhardt, 2014; Balažević et al., 2019a).

Let  $M \in \mathbb{R}^{n \times n}$  be the binary *reachability matrix* of a directed graph  $\mathcal{G}$  with  $n$  nodes, with  $M_{i,j} = 1$  if there exists a directed path from node  $i$  to node  $j$  and 0 otherwise. The Krackhardt hierarchy score of  $\mathcal{G}$  is defined as:

$$\text{Khs}_{\mathcal{G}} = \frac{\sum_{i=1}^n \sum_{j=1}^n \mathbb{1}(M_{i,j} == 1 \wedge M_{j,i} == 0)}{\sum_{i=1}^n \sum_{j=1}^n \mathbb{1}(M_{i,j} == 1)}. \quad (1)$$

## F SYMMETRY SCORE

The symmetry score  $\in [-1, 1]$  (Hubert & Baker, 1979) for a relation matrix  $R \in \mathbb{R}^{d_e \times d_e}$  is defined as:

$$s = \frac{\sum \sum_{i \neq j} R_{ij} R_{ji} - \frac{(\sum \sum_{i \neq j} R_{ij})^2}{d_e(d_e-1)}}{\sum \sum_{i \neq j} R_{ij}^2 - \frac{(\sum \sum_{i \neq j} R_{ij})^2}{d_e(d_e-1)}}, \quad (2)$$

where 1 indicates a symmetric and -1 an anti-symmetric matrix.

## G “OTHER” PREDICTED FACTS

Tables 10 to 13 shows a sample of the unknown triples (i.e. those formed using the WN18RR entities and relations, but not present in the dataset) for the *derivationally\_related form* (R), *instance\_hyponym* (S) and *synset\_domain\_topic\_of* (C) relations at a range of probabilities ( $\sigma(\phi(e_s, r, e_o)) \approx \{0.4, 0.6, 0.8, 1\}$ ), as predicted by each model. True triples are indicated in bold; instances where a model predicts an entity is related to itself are indicated in blue.

Table 10: "Other" facts as predicted by MuRE<sub>r</sub>.

Relation (Type)	$\sigma(\phi(e_s, r, e_o)) \approx 0.4$	$\sigma(\phi(e_s, r, e_o)) \approx 0.6$	$\sigma(\phi(e_s, r, e_o)) \approx 0.8$	$\sigma(\phi(e_s, r, e_o)) \approx 1$
derivationally-related_form (R)	<p>equalizer_NN_2, set_off_VB_5)</p> <p>(constellation_NN_2, satellite_NN_3)</p> <p>(shrink_VB_3, subtraction_NN_2)</p> <p>(continue_VB_10, proceed_VB_1)</p> <p>(support_VB_6, defend_VB_5)</p> <p>(shutter_NN_1, fill_up_VB_3)</p> <p>(yawing_NN_1, patellar_reflex_NN_1)</p> <p>(yaw_NN_1, spiral_VB_1)</p> <p>(stratum_NN_2, social_group_NN_1)</p> <p>(duel_VB_1, scuffle_NN_3)</p>	<p>(extrapolation_NN_1, maths_NN_1)</p> <p>(spread_VB_5, circularize_VB_3)</p> <p>(flaunt_NN_1, showing_NN_2)</p> <p>(extrapolate_VB_3, synthesis_NN_3)</p> <p>(strategist_NN_1, machination_NN_1)</p> <p>(crush_VB_4, grind_VB_2)</p> <p>(spike_VB_5, steady_VB_2)</p> <p>(licking_NN_1, vanquish_VB_1)</p> <p>(synthetical_JJ_1, synthesizer_NN_2)</p> <p>(realization_NN_2, embodiment_NN_3)</p>	<p>(sewer_NN_2, stitcher_NN_1)</p> <p>(land_VB_1, vegetable_oil_NN_1)</p> <p>(snuggle_NN_1, draw_close_VB_3)</p> <p>(train_VB_3, training_NN_1)</p> <p>(scratch_VB_3, skin_sensation_NN_1)</p> <p>(scheme_NN_5, schematization_NN_1)</p> <p>(ordain_VB_3, vest_VB_1)</p> <p>(lie_VB_1, front_end_NN_1)</p> <p>(tread_NN_1, step_NN_9)</p> <p>(register_NN_3, file_away_VB_1)</p>	<p>(trail_VB_2, trail_VB_2)</p> <p>(worship_VB_1, worship_VB_1)</p> <p>(steer_VB_1, steer_VB_1)</p> <p>(sort_out_VB_1, sort_out_VB_1)</p> <p>(make_full_VB_1, make_full_VB_1)</p> <p>(utilize_VB_1, utilize_VB_1)</p> <p>(geology_NN_1, geology_NN_1)</p> <p>(zoology_NN_2, zoology_NN_2)</p> <p>(uranology_NN_1, uranology_NN_1)</p> <p>(travel_VB_1, travel_VB_1)</p>
instance-hypernym (S)	<p>(thomas_aquinas_NN_1, martyr_NN_2)</p> <p>(volcano_islands_NN_1, volcano_NN_2)</p> <p>(cape_horn_NN_1, urban_center_NN_1)</p> <p>(bergen_NN_1, national_capital_NN_1)</p> <p>(marshall_NN_2, generalship_NN_1)</p> <p>(nansen_NN_1, venturer_NN_2)</p> <p>(wisconsin_NN_2, state_capital_NN_1)</p> <p>(prussia_NN_1, stockade_NN_2)</p> <p>(de_mille_NN_1, dancing_master_NN_1)</p> <p>(aegean_sea_NN_1, island_NN_1)</p>	<p>(taiwan_NN_1, asian_nation_NN_1)</p> <p>(st_gregory_of_n_n_NN_1, canonization_NN_1)</p> <p>(st_gregory_of_n_n_NN_1, saint_VB_2)</p> <p>(mccormick_NN_1, find_VB_8)</p> <p>(st_gregory_1_NN_1, bishop_NN_1)</p> <p>(richard_buckminster_f_NN_1, technological_JJ_2)</p> <p>(thomas_aquinas_NN_1, archbishop_NN_1)</p> <p>(marshall_NN_2, general_officer_NN_1)</p> <p>(newman_NN_2, primateship_NN_1)</p> <p>(thomas_the_apostle_NN_1, sanctify_VB_1)</p>	<p>(prophets_NN_1, gospels_NN_1)</p> <p>(malcolm_x_NN_1, passive_resister_NN_1)</p> <p>(taiwan_NN_1, national_capital_NN_1)</p> <p>(truth_NN_5, abolitionism_NN_1)</p> <p>(thomas_aquinas_NN_1, saint_VB_2)</p> <p>(central_america_NN_1, s_american_nation_NN_1)</p> <p>(de_mille_NN_1, dance_VB_1)</p> <p>(st_gregory_1_NN_1, apostle_NN_3)</p> <p>(fertile_crescent_NN_1, asian_nation_NN_1)</p> <p>(robert_owen_NN_1, industry_NN_1)</p>	<p>(helsinki_NN_1, urban_center_NN_1)</p> <p>(mannheim_NN_1, stockade_NN_2)</p> <p>(nippon_NN_1, nippon_NN_1)</p> <p>(victor_hugo_NN_1, novel_NN_1)</p> <p>(regionomachus_NN_1, uranology_NN_1)</p> <p>(prophets_NN_1, book_NN_10)</p> <p>(thomas_aquinas_NN_1, church_father_NN_1)</p> <p>(woody_guthrie_NN_1, minstrel_VB_1)</p> <p>(central_america_NN_1, c_american_nation_NN_1)</p> <p>(aegean_sea_NN_1, island_NN_1)</p>
synset_domain-topic_of (C)	<p>(write_VB_8, tape_VB_3)</p> <p>(introvert_NN_1, scientific_discipline_NN_1)</p> <p>(libel_NN_1, slur_NN_2)</p> <p>(evangelizing_NN_1, law_NN_1)</p> <p>(temple_NN_4, place_of_worship_NN_1)</p> <p>(trial_impression_NN_1, proof_VB_1)</p> <p>(friend_of_the_court_NN_1, war_machine_NN_1)</p> <p>(multiv_analysis_NN_1, applied_math_NN_1)</p> <p>(sell_VB_1, transaction_NN_1)</p> <p>(draw_VB_6, represent_VB_9)</p>	<p>(draw_VB_6, creative_person_NN_1)</p> <p>(suborder_NN_1, taxonomic_group_NN_1)</p> <p>(draw_VB_6, draw_VB_6)</p> <p>(first_sacker_NN_1, ballgame_NN_2)</p> <p>(alchemize_VB_1, modify_VB_3)</p> <p>(sermon_NN_1, sermon_NN_1)</p> <p>(saint_VB_2, catholic_church_NN_1)</p> <p>(male_JJ_1, masculine_JJ_2)</p> <p>(fire_VB_3, zoology_NN_2)</p> <p>(sell_VB_1, sell_VB_1)</p>	<p>(libel_NN_1, sully_VB_3)</p> <p>(relationship_NN_4, relationship_NN_4)</p> <p>(etymologizing_NN_1, linguistics_NN_1)</p> <p>(turn_VB_12, cultivation_NN_2)</p> <p>(brynhild_NN_1, mythologize_VB_2)</p> <p>(brynhild_NN_1, myth_NN_1)</p> <p>(assist_NN_2, am_football_game_NN_1)</p> <p>(mitzvah_NN_2, human_activity_NN_1)</p> <p>(drive_NN_12, drive_VB_8)</p> <p>(relationship_NN_4, biology_NN_1)</p>	<p>(libel_NN_1, disparagement_NN_1)</p> <p>(roll_on_roll_off_NN_1, transport_NN_1)</p> <p>(prance_VB_4, equestrian_sport_NN_1)</p> <p>(libel_NN_1, troaducement_NN_1)</p> <p>(sell_VB_1, selling_NN_1)</p> <p>(troi_VB_2, ride_horseback_VB_1)</p> <p>(prance_VB_4, ride_horseback_VB_1)</p> <p>(gallop_VB_1, ride_horseback_VB_1)</p> <p>(brynhild_NN_1, mythology_NN_2)</p> <p>(drive_NN_12, badminton_NN_1)</p>

Table 11: ‘‘Other’’ facts as predicted by DistMult.

Relation (Type)	$\sigma(\phi(e_s, r, e_o)) \approx 0.4$	$\sigma(\phi(e_s, r, e_o)) \approx 0.6$	$\sigma(\phi(e_s, r, e_o)) \approx 0.8$	$\sigma(\phi(e_s, r, e_o)) \approx 1$
derivationally-related_form (R)	(sag_VB_3, undercover_work_NN_1) (print_VB_4, publisher_NN_2) (crier_NN_3, pichman_NN_2) (play_VB_26, turn_NN_10) (count_VB_4, recite_VB_2) (witness_NN_2, imbue_VB_3) (sea_mew_NN_1, larus_NN_1) (alkali_NN_2, acidify_VB_2) (see_VB_17, understand_VB_2) (stun_VB_1, hedging_NN_2)	(dish_NN_2, stew_NN_2) (expose_VB_3, show_NN_1) (system_NN_9, orderliness_NN_1) (spread_NN_4, stew_VB_1) (take_down_VB_2, put_VB_2) (wrestle_VB_4, wrestle_NN_1) (autotr_organism_NN_1, epiphytic_JJ_1) (duel_VB_1, slugfest_NN_1) (vocal_NN_2, rock_star_NN_1) (smelling_NN_1, scent_VB_1)	(shrink_NN_1, pedology_NN_1) (finish_VB_6, finishing_NN_2) (play_VB_26, playing_NN_3) (centralization_NN_1, unite_VB_6) (existence_NN_1, living_NN_3) (mouth_VB_3, sassing_NN_1) (constellation_NN_2, star_NN_1) (print_VB_4, publishing_house_NN_1) (puzzle_VB_2, secret_NN_3) (uranology_NN_1, ur_NN_1)	(calliterate_VB_1, versifier_NN_1) (geology_NN_1, structural_JJ_5) (reset_VB_1, amputation_NN_2) (nutrition_NN_3, man_NN_4) (saint_NN_3, sanctify_VB_1) (right_fielder_NN_1, leftfield_NN_1) (list_VB_4, slope_NN_2) (lieutenancy_NN_1, captain_NN_1) (tread_NN_1, step_VB_7) (exentiation_NN_1, enucleate_VB_2)
instance-hyponym (S)	(wisconsin_NN_2, urban_center_NN_1) (marshal_NN_2, lieutenant_general_NN_1) (abidjan_NN_1, cote_d_ivoire_NN_1) (world_war_i_NN_1, urban_center_NN_1) (sr_paul_NN_2, evangelist_NN_2) (deep_south_NN_1, urban_center_NN_1) (nupise_NN_1, urban_center_NN_1) (ticino_NN_1, urban_center_NN_1) (aegean_sea_NN_1, aegean_island_NN_1) (coopens_NN_1, war_of_am_ind_NN_1)	(mississippi_river_NN_1, american_state_NN_1) (r_e_byrd_NN_1, commissioned_officer_NN_1) (kobenhavn_NN_1, urban_center_NN_1) (the_gambia_NN_1, africa_NN_1) (tirich_mir_NN_1, urban_center_NN_1) (r_e_byrd_NN_1, military_advisor_NN_1) (r_e_byrd_NN_1, aide-de-camp_NN_1) (tampa_boy_NN_1, urban_center_NN_1) (tidewater_region_NN_1, south_NN_1) (r_e_byrd_NN_1, executive_officer_NN_1)	(deep_south_NN_1, south_NN_1) (capital_of_gambia_NN_1, urban_center_NN_1) (south_west_africa_NN_1, africa_NN_1) (brandenburg_NN_1, urban_center_NN_1) (sierra_lewada_NN_1, urban_center_NN_1) (malcolm_x_NN_1, emancipationist_NN_1) (north_platte_river_NN_1, urban_center_NN_1) (oslo_NN_1, urban_center_NN_1) (sare_river_NN_1, urban_center_NN_1) (transylvanian_alps_NN_1, urban_center_NN_1)	(helsinki_NN_1, urban_center_NN_1) (the_nazarene_NN_1, save_VB_7) (irish_capital_NN_1, urban_center_NN_1) (r_e_byrd_NN_1, inspector_general_NN_1) (r_e_byrd_NN_1, chief_of_staff_NN_1) (central_america_NN_1, c_am_nation_NN_1) (malcolm_x_NN_1, environmentalist_NN_1) (the_nazarene_NN_1, christian_JJ_1) (thomas_aquinas_NN_1, church_father_NN_1) (the_nazarene_NN_1, el_nino_NN_2)
synset_domain-topic_of (C)	(limitation_NN_4, trammel_VB_2) (light_colonel_NN_1, colonel_NN_1) (nurse_VB_1, nursing_NN_1) (sermon_NN_1, prophesy_VB_2) (libel_NN_1, practice_of_law_NN_1) (slugger_NN_1, baseball_player_NN_1) (rna_NN_1, chemistry_NN_1) (metrify_VB_1, versify_VB_1) (trial_impression_NN_1, publish_VB_1) (turn_VB_12, plowman_NN_1)	(roll_on_roll-off_NN_1, transport_NN_1) (hizb_ut_tahrir_NN_1, asia_NN_1) (slugger_NN_1, softball_game_NN_1) (sermon_NN_1, sermonize_VB_1) (draw_VB_6, drawer_NN_3) (turn_VB_12, plow_NN_1) (assist_NN_2, softball_game_NN_1) (council_NN_2, assembly_NN_4) (throughput_NN_1, turnout_NN_4) (cream_VB_1, cream_NN_2)	(etymologizing_NN_1, explanation_NN_1) (ferry_VB_3, travel_VB_1) (public_prosecutor_NN_1, prosecute_VB_2) (alchemize_VB_1, modify_VB_3) (libel_NN_1, libel_VB_1) (turn_VB_12, till_VB_1) (hit_NN_1, hit_VB_1) (fire_VB_3, flaming_NN_1) (ring_NN_4, chemical_chain_NN_1) (libidinal_energy_NN_1, charge_NN_9)	(flat_JJ_5, matte_NN_2) (etymologizing_NN_1, derive_VB_3) (hole_out_VB_1, hole_NN_3) (relationship_NN_4, relation_NN_1) (drive_NN_12, badminton_NN_1) (etymologizing_NN_1, etymologize_VB_2) (matrix_algebra_NN_1, diagonalization_NN_1) (cabinetwork_NN_2, woodworking_NN_1) (cabinetwork_NN_2, bottom_VB_1) (cabinetwork_NN_2, upholster_VB_1)

Table 12: “Other” facts as predicted by TUCKER.

Relation (Type)	$\sigma(\phi(e_s, r, e_o)) \approx 0.4$	$\sigma(\phi(e_s, r, e_o)) \approx 0.6$	$\sigma(\phi(e_s, r, e_o)) \approx 0.8$	$\sigma(\phi(e_s, r, e_o)) \approx 1$
derivationally-related_form (R)	( <i>tympanist</i> _NN_1, <i>gong</i> _NN_2) ( <i>indication</i> _NN_1, <i>signalize</i> _VB_2) ( <i>turn</i> _over_VB_3, <i>rotation</i> _NN_3) ( <i>date</i> _VB_5, <i>geological_dating</i> _NN_1) ( <i>set</i> _VB_23, <i>emblem</i> _NN_2) ( <i>tyro</i> _NN_1, <i>start</i> _VB_5) ( <i>identification</i> _NN_1, <i>name</i> _VB_5) ( <i>stabber</i> _NN_1, <i>thrust</i> _VB_5) ( <i>justification</i> _NN_1, <i>apology</i> _NN_2) ( <i>manufacture</i> _VB_1, <i>prevarication</i> _NN_1)	( <i>mash</i> _NN_2, <i>mill</i> _VB_2) ( <i>walk</i> _VB_9, <i>immer_frame</i> _NN_1) ( <i>use</i> _VB_5, <i>utility</i> _NN_2) ( <i>musical_instrument</i> _NN_1, <i>write</i> _VB_6) ( <i>lining</i> _NN_3, <i>wrap_up</i> _VB_1) ( <i>scrap</i> _VB_2, <i>struggle</i> _NN_2) ( <i>tape</i> _VB_3, <i>tape_recorder</i> _NN_1) ( <i>vindicate</i> _VB_2, <i>justification</i> _NN_2) ( <i>leaching</i> _NN_1, <i>percolate</i> _VB_3) ( <i>synchronize</i> _VB_2, <i>synchroscope</i> _NN_1)	( <i>take_chances</i> _VB_1, <i>venture</i> _NN_1) ( <i>shutter</i> _NN_1, <i>fill_up</i> _VB_3) ( <i>exit</i> _NN_3, <i>leave</i> _VB_1) ( <i>trembler</i> _NN_1, <i>vibrate</i> _VB_1) ( <i>motivator</i> _NN_1, <i>trip</i> _VB_4) ( <i>support</i> _VB_6, <i>indorsement</i> _NN_1) ( <i>federate</i> _VB_2, <i>confederation</i> _NN_1) ( <i>take_over</i> _VB_6, <i>return</i> _NN_7) ( <i>wait_on</i> _VB_1, <i>supporter</i> _NN_3) ( <i>denote</i> _VB_3, <i>promulgation</i> _NN_1)	( <i>venturer</i> _NN_2, <i>venturer</i> _NN_2) ( <i>dynamitist</i> _NN_1, <i>dynamitist</i> _NN_1) ( <i>love</i> _VB_3, <i>lover</i> _NN_2) ( <i>snuggle</i> _NN_1, <i>squeeze</i> _VB_8) ( <i>departed</i> _NN_1, <i>die</i> _VB_2) ( <i>position</i> _VB_1, <i>placement</i> _NN_1) ( <i>repentant</i> _JJ_1, <i>repentant</i> _JJ_1) ( <i>thead</i> _NN_1, <i>step</i> _VB_7) ( <i>stockist</i> _NN_1, <i>stockist</i> _NN_1) ( <i>philanthropist</i> _NN_1, <i>philanthropist</i> _NN_1)
instance_hyponym (S)	( <i>deep_south</i> _NN_1, <i>south</i> _NN_1) ( <i>st_paul</i> _NN_2, <i>organist</i> _NN_1) ( <i>helsinki</i> _NN_1, <i>urban_center</i> _NN_1) ( <i>malcolm_x</i> _NN_1, <i>emancipationist</i> _NN_1) ( <i>thomas_the_apollo</i> _NN_1, <i>church_father</i> _NN_1) ( <i>st_gregory_of_n</i> _NN_1, <i>sermonizer</i> _NN_1) ( <i>robert_owen</i> _NN_1, <i>movie_maker</i> _NN_1) ( <i>the_resh</i> _NN_1, <i>monk</i> _NN_1) ( <i>st_paul</i> _NN_2, <i>philosopher</i> _NN_1) ( <i>ibn-roshd</i> _NN_1, <i>pedagogue</i> _NN_1)	( <i>thomas_aquinas</i> _NN_1, <i>bishop</i> _NN_1) ( <i>irish_capital</i> _NN_1, <i>urban_center</i> _NN_1) ( <i>thomas_the_apollo</i> _NN_1, <i>apostle</i> _NN_2) ( <i>st_paul</i> _NN_2, <i>apostle</i> _NN_3) ( <i>mcormick</i> _NN_1, <i>painter</i> _NN_1) ( <i>thomas_the_apollo</i> _NN_1, <i>trogodyte</i> _NN_1) ( <i>mcormick</i> _NN_1, <i>electrical_engineer</i> _NN_1) ( <i>mississippi_river</i> _NN_1, <i>american_state</i> _NN_1)	( <i>cowpens</i> _NN_1, <i>siege</i> _NN_1) ( <i>mcormick</i> _NN_1, <i>arms_manufacturer</i> _NN_1) ( <i>thomas_the_apollo</i> _NN_1, <i>evangelist</i> _NN_2) ( <i>mcormick</i> _NN_1, <i>technologist</i> _NN_1) ( <i>st_gregory_i</i> _NN_1, <i>church_father</i> _NN_1)	( <i>r_e_byrd</i> _NN_1, <i>siege</i> _NN_1) ( <i>shaw</i> _NN_3, <i>women's</i> _rightist_NN_1) ( <i>aegean_sea</i> _NN_1, <i>aegean_island</i> _NN_1) ( <i>thomas_aquinas</i> _NN_1, <i>church_father</i> _NN_1)
synset_domain_topic_of (C)	( <i>roll_on_roll_off</i> _NN_1, <i>motorcar</i> _NN_1) ( <i>libel</i> _NN_1, <i>legislature</i> _NN_1) ( <i>roll_on_roll_off</i> _NN_1, <i>passenger_vehicle</i> _NN_1)	( <i>drive</i> _NN_12, <i>badminton</i> _NN_1)		

Table 13: "Other" facts as predicted by MuRE.

Relation (Type)	$\sigma(\phi(e_s, r, e_o)) \approx 0.4$	$\sigma(\phi(e_s, r, e_o)) \approx 0.6$	$\sigma(\phi(e_s, r, e_o)) \approx 0.8$	$\sigma(\phi(e_s, r, e_o)) \approx 1$
derivationally-related_form (R)	(surround_VB_1, wall_NN_1) (unpleasant_JJ_1, unpalatableness_NN_1) (love_VB_3, enjoyment_NN_2) (magnitude_NN_1, tail_JJ_1) (testify_VB_2, information_NN_1) (connect_VB_6, converging_NN_1) (connect_VB_6, connexion_NN_4) (operate_VB_4, psypop_NN_1) (market_VB_1, trade_NN_4) (operate_VB_4, mission_NN_2)	(word_picture_NN_1, sketch_VB_2) (develop_VB_10, adjustment_NN_4) (gloss_VB_3, commentary_NN_1) (violate_VB_2, violation_NN_3) (suffocate_VB_1, strangler_tree_NN_1) (number_VB_3, point_NN_12) (develop_VB_10, organic_process_NN_1) (plication_NN_1, twist_VB_4) (split_up_VB_3, separation_NN_5) (plication_NN_1, wrinkle_VB_2)	(smelling_NN_1, wind_VB_4) (try_out_VB_1, somatic_cell_nuclear_transplantation_NN_1) (lighting_NN_4, set_on_fire_VB_1) (symphtalagus_NN_1, one-half_NN_1) (just_JJ_3, validity_NN_1) (reprove_VB_1, talking_to_NN_1) (sustain_VB_5, beam_NN_2) (spring_NN_6, hurdle_VB_1) (spark_NN_1, scintillate_VB_1) (utility_NN_2, functional_JJ_1)	(spoliation_NN_2, sack_VB_1) (desire_NN_2, hope_VB_2) (shuffle_VB_3, whine_NN_1) (nasalization_NN_1, sound_out_VB_1) (wear_NN_1, step_VB_7) (unreliableness_NN_1, arbitrary_JJ_1) (travesty_NN_2, travesty_NN_2) (spark_NN_1, sparkle_VB_1) (stockist_NN_1, stockist_NN_1)
instance-hyponym (S)	(malcolm_x_NN_1, hipster_NN_1) (the_nazarene_NN_1, juddism_NN_2) (old_line_state_NN_1, river_NN_1) (r_e_hyrd_NN_1, commissioned_officer_NN_1) (south_korea_NN_1, peninsula_NN_1) (st_gregory_of_n_NN_1, vicar_of_christ_NN_1) (nippon_NN_1, italian_region_NN_1) (robert_owen_NN_1, tycoon_NN_1) (mandalay_NN_1, national_capital_NN_1) (nan_ling_NN_1, urban_center_NN_1)	(central_america_NN_1, central_america_NN_1) (st_gregory_i_NN_1, church_father_NN_1) (south_korea_NN_1, african_nation_NN_1) (malcolm_x_NN_1, passive_resister_NN_1) (malcolm_x_NN_1, birth-control_reformer_NN_1) (los_angeles_NN_1, port_NN_1) (great_lakes_NN_1, canadian_province_NN_1) (transylvanian_alps_NN_1, urban_center_NN_1) (gettysburg_NN_2, steege_NN_1) (wisconsin_NN_2, geographical_region_NN_1)	(theresa_NN_1, monk_NN_1) (nippon_NN_1, european_nation_NN_1) (great_lakes_NN_1, river_NN_1) (r_e_hyrd_NN_1, noncommissioned_officer_NN_1) (world_war_i_NN_1, pitched_battle_NN_1) (irish_capital_NN_1, urban_center_NN_1) (volcano_islands_NN_1, urban_center_NN_1) (nippon_NN_1, american_state_NN_1) (helsinki_NN_1, urban_center_NN_1) (capital_of_gambia_NN_1, urban_center_NN_1)	
synset_domain-topic_of (C)	(libel_NN_1, criminal_law_NN_1) (brynithid_NN_1, mythology_NN_2) (slugger_NN_1, sport_NN_1) (sell_VB_1, law_NN_1) (semitic_deity_NN_1, mythology_NN_1) (nuclear_deterrence_NN_1, law_NN_1) (reception_NN_5, baseball_game_NN_1) (photosynthesis_NN_1, chemistry_NN_1) (isolde_NN_1, parable_NN_1) (assist_NN_2, court_game_NN_1)	(write_VB_8, transcription_NN_5) (temple_NN_4, muslimism_NN_2) (assist_NN_2, hockey_NN_1) (relationship_NN_4, biology_NN_1) (oposite_NN_3, western_church_NN_1) (assist_NN_2, sport_NN_1) (trot_VB_2, equestrian_sport_NN_1) (rma_NN_1, chemistry_NN_1) (assist_NN_2, soccer_NN_1) (assist_NN_2, football_game_NN_1)	(assist_NN_2, am_football_game_NN_1) (drive_NN_12, court_game_NN_1) (sell_VB_1, offense_NN_3) (slugger_NN_1, softball_game_NN_1) (drive_NN_12, badminton_NN_1)	