

2019

Augmentic Compositional Models for Knowledge Base Completion Using Gradient Representations

Matthias R. Lalisce

Johns Hopkins University, lalisce@jhu.edu

Paul Smolensky

Johns Hopkins University, smolensky@jhu.edu

Follow this and additional works at: <https://scholarworks.umass.edu/scil>



Part of the [Computational Linguistics Commons](#)

Recommended Citation

Lalisce, Matthias R. and Smolensky, Paul (2019) "Augmentic Compositional Models for Knowledge Base Completion Using Gradient Representations," *Proceedings of the Society for Computation in Linguistics*: Vol. 2 , Article 27.

DOI: <https://doi.org/10.7275/8et8-qd83>

Available at: <https://scholarworks.umass.edu/scil/vol2/iss1/27>

This Paper is brought to you for free and open access by ScholarWorks@UMass Amherst. It has been accepted for inclusion in Proceedings of the Society for Computation in Linguistics by an authorized editor of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

Augmenting Compositional Models for Knowledge Base Completion Using Gradient Representations

Matthias Lalissee

Dept of Cognitive Science
Johns Hopkins University
Baltimore, MD USA
lalissee@jhu.edu

Paul Smolensky

Dept of Cognitive Science
Johns Hopkins University
& Microsoft Research AI
Seattle, WA USA
smolensky@jhu.edu

Abstract

Neural models of Knowledge Base data have typically employed compositional representations of graph objects: entity and relation embeddings are systematically combined to evaluate the truth of a candidate Knowledge Base entry. Using a model inspired by Harmonic Grammar, we propose to tokenize triplet embeddings by subjecting them to a process of optimization with respect to learned well-formedness conditions on Knowledge Base triplets. The resulting model, known as Gradient Graphs, leads to sizable improvements when implemented as a companion to compositional models. Also, we show that the “supracompositional” triplet token embeddings it produces have interpretable properties that prove helpful in performing inference on the resulting triplet representations.

1 Introduction

As they are conventionally analyzed, representations of semantic or linguistic data are “compositional”: the meanings of complex representations are built up from the meanings of their constituent parts. This idea has motivated numerous models of graph data deployed in knowledge base completion (KBC), in which embeddings of entities and relations are combined into composite representations—pairs of entities in a particular relation with one another—that are built up systematically from the constituent parts. But what happens when the whole is *not* a simple function of the parts? A natural case arises in the interpretation of Noun-Noun compounds. The contrasting senses of *vampire cat* (*a-cat-that-is-a-vampire*) and *vampire stake* (*a-stake-used-to-kill-a-vampire*) has as much to do with the compatibility of the constituent nouns occurring in a given relation than with the meanings of the individual constituents.

Pursuing this line of thought, we propose **Gradient Graphs**, a neural network model for

KBC built on the principle that compositionally-obtained representations of semantic objects can be optimized to reflect context-specific aspects of the meanings of their constituents. The issue of context-conditioned, tokenized semantic representations has received little explicit attention in the KBC literature. However, precedents do exist. [Bordes et al. \(2011\)](#) model context-sensitive entity senses by embedding relations as pairs of matrices (R_{lhs}, R_{rhs}) that linearly transform entity embeddings into pairs of embeddings defined by the relation and the entities’ positions within it (the left-hand-side or right-hand-side). The distances of the resulting embeddings are then compared. [Socher et al. \(2013\)](#) cope with the context-sensitivity of relation meanings by learning a $k \times d \times d$ -dimensional tensor embeddings for each relation, letting their model represent polysemy by learning k versions of the relation represented in the k slices of its embedding tensor. The intuition underlying this approach is that, for instance, the relation **has part** has a different sense when applied to a biological organism than when predicated of a company. While the former has parts like organs and limbs, the latter has parts like subsidiaries and workers, which occupy very different parts of the semantic space. Each relational slice is then responsible for learning the compatibility of arguments within particular semantic subspaces.

In contrast to these other works, our approach is more radical in the sense that our context-sensitive representations of knowledge base entries are not just computed from the entries’ constituent elements (entity and relation embeddings), but are instead the result of a representation-optimization procedure that balances compositionally-derived representations with general knowledge about the characteristics of well-formed semantic structures. We show that this additional “supracompositional” processing, in addition to yielding sizable accuracy improvements over the compositional models we apply it to, leads to embeddings of entity tokens with interpretable characteristics.

1.1 Layout of the paper

Section 2 lays out the general framework, which is compatible with a variety of implementations. Section 3 presents two compositional embedding models proposed in the literature. We adapt these models to construct compositional embeddings, and in Section 4 report evaluations of Gradient versions of these models. Section 5 discusses the characteristics of the resulting semantic representations in greater detail, as well as their role in assisting inference. Section 6 concludes. Technical details about the model and the implementations are given in the Appendices.

2 Optimization of semantic tokens




Figure 1: Gradient Graph as a recurrent neural network. In addition to bias terms (omitted in the figure) and self-connections, hidden units are densely connected to one another via a layer of connections with symmetric (undirected) weights, and receive constant input weighted by λ from a single unit in the input layer. The composition function $f_{comp}(e_\ell, r, e_r)$, which differs between implementations, computes a **compositional embedding** x , which is fed into a **hidden layer** h of the network. The continuous-time dynamics of this network compute an internal representation \hat{h} of the input triplet that is optimal with respect to the Harmony (1)—a measure of the triplet’s semantic well-formedness.

The hypothesis underlying the approach we propose is that noncompositional effects in knowledge base data can be modeled by subjecting candidate facts to a process of optimization with respect to a set of learned semantic coherence conditions.

These semantic coherence conditions, encoded in a symmetric matrix, map out the covariance structure of the semantic space, indicating which semantic features are likely to co-occur with one another. The embedding of a given triplet is then the vector obtained by optimizing the semantic coherence of the the triplet embedding.

We first lay out the model in abstract form, before introducing particular implementations. Let $x \in \mathbb{R}^d$ be a d -dimensional embedding of a knowledge base triplet (e_ℓ, r, e_r) obtained as some function f_{comp} —the **composition function**—of the embeddings of the left and right entities as well as the relation r . Section 3 provides several models for constructing the triplet embedding x . Also, let h be a d -dimensional vector giving the internal (“hidden”) state of the network. The *Harmony* of an internal state h of the network with respect to the triplet embedding x is

$$(1) \quad \mathcal{H}(h, x) = \frac{1}{2} [h^\top \mathbb{W} h + b^\top h - \lambda (h - x)^\top (h - x)]$$

where \mathbb{W} is a $d \times d$ weight matrix with $\mathbb{W} = \mathbb{W}^\top$ and b is a bias vector, both learned. (1) is composed of two terms: **Core Harmony**, a measure of the semantic coherence of the state vector h , and **Faithfulness**, a penalty incurred due to the state h ’s deviation from the compositional triplet embedding x . λ is a hyperparameter that controls the magnitude of the penalty incurred for straying from x .

$\mathcal{H}(h, x)$ may be rewritten as (2).

$$(2) \quad \mathcal{H}(h, x) = \frac{1}{2} [h^\top (\mathbb{W} - \lambda I) h + (b + 2\lambda x)^\top h - \lambda x^\top x]$$

If λ is greater than the largest eigenvalue of \mathbb{W} , then $V = \mathbb{W} - \lambda I$ is negative-definite, and $\mathcal{H}(h, x)$ has a unique global optimum $\hat{h} = \operatorname{argmax}_h \mathcal{H}(h, x)$ for each x . In closed form, this global optimum is

$$(3) \quad \mu(x) = \frac{1}{2} V^{-1} (b + 2\lambda x)$$

which depends only on the network parameters and on x . The expression $\mu(x)$ comes from observing that \hat{h} is the mean of a Gaussian distribution with inverse covariance matrix V , which implies that \hat{h} is the most probable state h of the network with respect to the probability distribution over the state space defined by $p(h|x) \propto \exp\{\mathcal{H}(h, x)\}$ (see Appendix A). We take the token embedding for a triplet x to be $\mu(x)$, which is the most semantically coherent triplet embedding given the compositional triplet x . In the limit as $\lambda \rightarrow \infty$, $\mu(x)$ is just x itself. Let $\lambda_{\mathbb{W}}$ denote the largest eigenvalue of \mathbb{W} ; then as $\lambda \rightarrow \lambda_{\mathbb{W}}$, $\mu(x)$ becomes independent of the triplet embedding x .

A Gradient Graph may be viewed as a neural network with weight matrix \mathbb{W} and bias vector $\frac{b}{2}$,

where the synaptic weights \mathbb{W} specify a feedback layer through which the values of the hidden state units affect one another. The construction is as follows. We stipulate that the hidden state of the network follows the gradient of Harmony over time:

$$(4) \quad \frac{dh}{dt} = \frac{\partial \mathcal{H}(x, h)}{\partial h}$$

Therefore,

$$\begin{aligned} \frac{dh_i}{dt} &= \frac{d}{dh_i} \frac{1}{2} \left[h^\top \mathbb{W} h + b^\top h - \lambda \|x - h\|^2 \right] \\ &= \frac{d}{dh_i} \frac{1}{2} \left[\sum_{jk} h_j \mathbb{W}_{jk} h_k + b_i h_i - \lambda (x_i - h_i)^2 \right] \\ &= \frac{1}{2} \left[\sum_j h_j \mathbb{W}_{ji} + \sum_k \mathbb{W}_{ik} h_k \right] + \frac{b_i}{2} + \lambda x_i - \lambda h_i \end{aligned}$$

The above specifies the connectivity of a network whose hidden units have the linear transfer function ($f(\text{input}) = \text{input}$), bias $\frac{b}{2}$ and external input x (weighted by λ). Each h_i also receives self-inhibitory input weighted by $-\lambda$, as well as inputs $\mathbb{W}_{ij} h_j$ from each h_j . The symmetry of \mathbb{W} implies that each term $\mathbb{W}_{ij} h_j = h_j \mathbb{W}_{ji}$ occurs twice, so that the factor of $\frac{1}{2}$ cancels. This connectivity structure is illustrated in Figure 1.

2.1 Relation to Harmonic Grammar

In addition to being globally optimal with respect to the Harmony function $\mathcal{H}(h, x)$ conditioned on a particular input x , $\mu(x)$ is the unique fixed point of this network’s state-evolution dynamics. It is interesting to note that such networks are the connectionist foundation for Harmonic Grammar (HG) and Optimality Theory (OT) in Linguistics (Smolensky and Legendre, 2006), where the dynamics of a neural network perform optimization over internal representations of an input structure. Appropriate output representations are then selected in accordance with well-formedness constraints encoded in the network parameters. There, the output representation balances Faithfulness to the input (an Underlying Form) and the network’s knowledge about the characteristics of well-formed structures in general.

Similarly, it is appealing to conceptualize the hidden layer of a Gradient Graph network as cleaning up a knowledge base triplet by subjecting it to semantic well-formedness conditions. The optimal triplet $\mu(x)$ is then the point to which the network converges in the limit of infinite computation time. However, our model differs from typical implementations of HG and OT in that the optimal structure $\mu(x)$ does not, in general, decompose into a unique combination of the input constituents (entity and relation embeddings). The resulting representations are in this sense gradient, rather than

being the product of a combination of discrete objects. Furthermore, Gradient Graphs are, to our knowledge, the first application of these ideas to the automatic learning of an appropriate semantic optimization function from a large amount of data.

2.2 Comparison with translation-based approaches

Like a large class of *Translation*-based models (Bordes et al., 2011; Yoon et al., 2016; Lin et al., 2015; Ji et al., 2016), our inference procedure consists of the application of an affine transformation to an input x (Equation (3)), which is then scored using some regular operation. In our case, this scoring function is quadratic. A particular close cousin is the bilinear SEMANTIC MATCHING ENERGY (SME) method of Bordes et al. (2014), which learns a global third-order tensor \mathbb{W} that, when dotted along the third mode with a relation embedding r , yields a relation-specific matrix \mathbb{W}_r . Along with learned left and right bias vectors b_ℓ and b_r , this weight matrix is fed into the bilinear scoring function (5):

$$(5) \quad \text{score}_{\text{SME}}(e_\ell, r, e_r) = \frac{1}{2} (e_\ell^\top \mathbb{W}_r e_r + b_\ell^\top e_\ell + b_r^\top e_r)$$

Expanding out this expression, we get (6).

$$(6) \quad e_\ell^\top \mathbb{W}_r e_r + b_\ell^\top e_\ell + b_r^\top e_r + b_\ell^\top b_r$$

The relation-specific bilinear form $e_\ell^\top \mathbb{W}_r e_r$ is, like our global \mathbb{W} matrix, symmetric. The remaining terms, apart from the constant $b_\ell^\top b_r$, compute a pair of relation-specific bias vectors $b_\ell^\top \mathbb{W}_r$ and $b_r^\top \mathbb{W}_r$ applied to the pair of entity embeddings. The resulting Energy function used to score triplets has more than a passing similarity to our Harmony function (1) when $\lambda = \infty$ and, thus, no optimization takes place.

A distinctive characteristic of our approach in relation to these structurally similar models is that the transformation undergone by a Gradient Graph triplet is directly connected to the well-formedness criterion according to which triplets are evaluated in inference. As illustrated in the Discussion, our transformation of a compositional triplets using learned well-formedness criteria leads to two kinds of triplet embeddings: compositionally obtained *type* embeddings, and contextually optimized *token* embeddings. In qualitative and quantitative analyses of the learned representation, we see (1) that the space of compositionally obtained triplet embeddings has a reasonable structure, independently of the optimizing transformation, that is already sensitive to the context supplied by the relation, and (2) that semantic optimization improves these compositional representations in recognizable ways. Interestingly, im-

proving triplets with respect to the Harmony function does *not* uniformly place them in regions that are high-Harmony in a global sense. In fact, we find that whereas positive triplets end up close to other positive triplets, plausible but negative triplets tend to be detained in clusters with other negative instances (Tables 2 and 4).

3 Compositional and Gradient Models

Optimization with respect to \mathcal{H} can be implemented wherever we can construct a triplet embedding x . In our experiments, we apply Harmonic optimization of triplet representations to two compositional embedding models drawn from the knowledge base completion literature: DISTMULT and HOLE. Both models specify a scoring function for triplet embeddings obtained via operations applied to embeddings of the three triplet components—two entity vectors and a relation vector—with no additional learned components apart from these representations of the triplet constituents. We take the terms occurring in these scoring functions to be components of the *representation* of the triplet, specifying what information about the triplet elements is important to evaluating the triplet’s quality. Hence, we constructed Harmonic triplet embeddings according to the desideratum that every term occurring in the basic method’s scoring function should also appear in the triplet representation x in the Harmonic model. For instance, the score of a DISTMULT triplet is a sum of three-way products of the corresponding elements of the embeddings e_ℓ , r , and e_r . Setting the products $[e_\ell]_i[r]_i[e_r]_i$ to appear in our compositional triplet embeddings (as in Eqn (8)) satisfies this desideratum.

DISTMULT (Yang et al., 2015) is a baseline model for scoring knowledge base triplets using the scoring function (7):

$$(7) \quad score_{\text{DISTMULT}}(e_\ell, r, e_r) = e_\ell^\top \text{diag}(r) e_r$$

where e_ℓ, r, e_r are d -dimensional embeddings and $\text{diag}(r)$ is the $d \times d$ -dimensional matrix obtained by arranging the elements of r along the diagonal. Kadlec et al. (2017) have recently shown that DISTMULT can outperform many more complicated scoring functions when hyperparameters are properly optimized, making it a strong baseline comparison for the method we propose. In addition, DISTMULT often occurs as a subcomponent in state-of-the-art KBC models—e.g. (Schlichtkrull et al., 2017; Toutanova et al., 2015). From this starting-point, we construct HARMONIC DISTMULT (HDISTMULT) by setting the triplet embedding x to the elementwise multiplication of the relation and the pair of entity vectors:

$$(8) \quad x_{\text{HDM}} = e_\ell \odot r \odot e_r$$

where \odot denotes elementwise multiplication.

HOLOGRAPHIC EMBEDDINGS (HOLE) were introduced by Nickel et al. (2016) building on theoretical work by (Plate, 1995), as a means of constructing compressed tensor product representations of relational triplets. The method computes the score for a triplet (e_ℓ, r, e_r) from the similarity between a relation vector and the circular correlation $e_\ell \star e_r$ of the entity vectors and a relation vector:

$$(9) \quad score_{\text{HOLE}}(e_\ell, r, e_r) = r^\top (e_\ell \star e_r)$$

where the circular correlation of e_ℓ and e_r is computed as (10).

$$(10) \quad e_\ell \star e_r = \mathcal{F}^{-1} \left(\overline{\mathcal{F}(e_\ell)} \odot \mathcal{F}(e_r) \right)$$

\mathcal{F} and \mathcal{F}^{-1} denote the Fourier Transform and its inverse, and $\overline{\mathcal{F}(e_\ell)}$ is the complex conjugate of $\mathcal{F}(e_\ell)$.¹ Circular correlation is asymmetric ($e_\ell \star e_r \neq e_r \star e_\ell$)—allowing it to model asymmetric relations—and the result of the operation has the same dimensionality as the input vectors, while still carrying information about which pair of entities was bound together via correlation.

We construct HARMONIC HOLE (HHOLE) triplet embeddings via elementwise multiplication of relation vectors with the correlated pair of entity vectors:

$$(11) \quad x_{\text{HHOLE}} = r \odot (e_\ell \star e_r)$$

In both Harmonic models, the score for a candidate triplet (e_ℓ, r, e_r) with embedding x is calculated by taking the Harmony of its optimal instantiation, $\hat{h} = \mu(x)$, i.e.

$$(12) \quad score(x) = \mathcal{H}(\mu(x), x)$$

In the experiments, we train our networks using the log-softmax objective with negative sampling. For each positive training example (e_ℓ, r, e_r) with embedding x , we construct N negative examples $(\tilde{e}_\ell^n, \tilde{r}^n, \tilde{e}_r^n)$ obtained by deleting either the left or right entity of the true triplet and replacing it with a randomly sampled entity vector. Let \tilde{x}^n denote the embedding of the n th negatively sample triplet

¹The Fourier transform decomposes a function of time into its frequency components. In the context of holographic embeddings, its utility comes from the *Convolution Theorem*, which states that convolution in the time domain corresponds to elementwise multiplication in the frequency domain. This is useful in actual computations. The circular correlation—which consists of convolution with a time-reversed signal—can also be computed as a sum over off-diagonals of the tensor product of vectors, with time complexity $\mathcal{O}(d^2)$. In contrast, the Fast Fourier Transform (FFT) has time complexity $\mathcal{O}(n \log n)$ (Nickel et al., 2016).

$(\tilde{e}_\ell^n, \tilde{r}^n, \tilde{e}_r^n)$. The training objective is then to minimize (13):

$$(13) \quad \mathcal{L}_{\mathcal{H}}(e_\ell, r, e_r) = -\log \frac{\exp\{\mathcal{H}(\mu(x), x)\}}{\exp\{\mathcal{H}(\mu(x), x)\} + \sum_{n=1}^N \exp\{\mathcal{H}(\mu(\tilde{x}^n), \tilde{x}^n)\}}$$

This has the effect of increasing the Harmony of positive examples relative to negative samples. The learning rule is thus Harmony-maximizing: the network parameters maximize the well-formedness of the positive examples relative to negative samples.

4 Experiments

We evaluated Gradient Graphs using the standard WN18 and FB15K datasets (Bordes et al., 2013)—which are subsets of the WordNet (Miller, 1995) and Freebase (Bollacker et al., 2008) databases—in the Entity Reconstruction task. In Entity Reconstruction, the network ranks completions of triplets (\cdot, r, e_r) and (e_ℓ, r, \cdot) with deleted left and right entities. The model is successful if it ranks the true triplet above other candidate completions. We report results in the *filtered* evaluation setting (Bordes et al., 2013), in which a test triplet is only ranked against triplets that do not occur in the database. The rank of a test triplet is thus the rank of the *first* correct answer to the query. For both DISTMULT and HOLE, we report the originally reported results alongside results for our reimplementations, comparing these models with our Harmonic variants HDISTMULT and HHOLE with and without optimization of hidden layer representations. The Harmonic models with $\lambda = \infty$ have the Harmony function $\mathcal{H}(x, x)$, i.e. where the hidden representation is just the compositional embedding itself and the Faithfulness penalty in (1) is 0.

Our models used 256- to 512-dimensional embeddings and manually tuned values of the hyperparameter λ . In all models, entity and relation embeddings were normalized to $\|v\| = 1$. We do not regularize parameters, but instead set an upper bound $\lambda - \epsilon$ (ϵ a small constant) on the l_2 norm of the weight matrix \mathbb{W} , which helps constrain the spectral norm (maximum eigenvalue) of \mathbb{W} to remain lower than λ . This is equivalent to adopting a uniform prior on weight matrices lying within the n -ball with squared radius $\lambda - \epsilon$. Importantly, this procedure keeps the matrix $V = \mathbb{W} - \lambda I$ negative-definite—a necessary condition for the existence of a unique optimum for $\mathcal{H}(h, x)$.

Results from the experiments are reported in Table 1. Overall, we found that models using our quadratic scoring function (1) to perform best across the board. This effect was particularly seen in more stringent evaluation criteria—Hits@1 and

Hits@3—leading to, for instance—a 15% improvement in Hits@1 (accuracy) on Freebase between our DISTMULT reimplementations and quadratic HDISTMULT ($\lambda = \infty$). The bestDISTMULT models were those with high λ values; however, within-model comparison of HOLE shows dramatic improvements from including the optimization component—a 32% increase in FB15K accuracy between the results of (Nickel et al., 2016) and our HHOLE with a permissive λ -criterion of 1.0.

5 Discussion

In part, the appeal of our supracompositional representations stems from their ability to produce embeddings of *tokens* of semantic objects—that is, embeddings that take into account the context of a particular instance of a semantic type. Tokenized embeddings have proven useful in various settings. For instance, Dasigi et al. (2017) construct token embeddings by superposing learned vectors for WordNet senses in ratios determined by a probability distribution computed from the context. The resulting representation is a context-weighted sum of discrete senses drawn from a hand-crafted ontology. Closer to our approach, Belanger and Kakade (2015) model text as a linear dynamical system that generates texts through transitions of a continuous-state, discrete-time dynamical system across time. Estimates of the system’s most probable internal state can then be extracted as an embedding of the tokens, which prove useful in language modeling and other downstream tasks.

In our framework, *types* correspond to static entity and relation embeddings that are the input to f_{comp} , and the triplet embeddings resulting from their combination. *Token* triplet embeddings are produced by optimization of the hidden layer of a GG. To understand the effect of optimizing the hidden layer of a GGRAPH both on its learned representations and on its performance in inference, we used the best-performing trained HHOLE model to produce token embeddings of database triplets in order to inspect their semantic neighborhoods. For a given compositional triplet embedding $f_{comp}(e_\ell, r, e_r) \equiv x$, we first computed the optimized triplet representation $\mu(x) \equiv \hat{h}$ using Equation (3). Treating \hat{h} as the contextually optimal (token) embedding of the triplet (e_ℓ, r, e_r) , we then examined the semantic neighborhood by computing the 5 closest optimized embeddings in the context of the same relation. Table 4 shows the semantic neighborhoods of compositional triplets x and optimized triplets \hat{h} for different possible completions of a number of queries. Rows 1 and 2 display completions of the query $(\cdot, \text{office_position_or_title}, \text{US_President})$, and Rows 3 and 4 consider the neighborhood of the entity embedding of *Bob Dylan* in the context of

Model	FB15K						WN18					
	λ	Rank		Hits@			λ	Rank		Hits@		
		MR	MRR	1	3	10		MR	MRR	1	3	10
DISTMULT	-	-	.350	-	-	.577	-	-	.830	-	-	.942
ENSEMBLE DM [†]	-	36	.837	.797	-	.904	-	457	.790	.784	-	.950
DISTMULT*	-	28	.710	.605	.792	.876	-	220	.825	.714	.938	.950
HDISTMULT	∞	23	.806	.751	.845	.898	∞	164	.841	.740	.943	.955
HDISTMULT	50.0	23	.742	.661	.799	.881	3.0	184	.831	.732	.931	.945
HOLE	-	-	.524	.402	.613	.739	-	-	.938	.930	.945	.949
HOLE*	-	39	.409	.289	.464	.647	-	205	.916	.893	.936	.946
HHOLE	∞	32	.682	.575	.763	.850	∞	293	.919	.903	.934	.942
HHOLE	1.0	21	.796	.727	.848	.901	2.0	183	.939	.931	.945	.951

Table 1: Results on FB15K and WN18. The results from the original DISTMULT and HOLE models are drawn from (Yang et al., 2015) and (Nickel et al., 2016). Our reimplementations* of DISTMULT and HOLE differ in numerous details from those in the original papers (see Appendix B for technical details). ENSEMBLE DISTMULT[†] refers to the hyperparameter-optimized Ensemble (product of experts) reimplementation of DistMult proposed by Kadlec et al. (2017). For each model, we report Mean Rank (MR) and Mean Reciprocal Rank (MRR), as well as Hits@N for $N \in \{1, 3, 10\}$. Hits@N denotes the fraction of test instances in which the true triplet completion had rank less than or equal to N. The best results within each category (DISTMULT and HOLE) are marked in **bold**, and the best results overall are additionally **underlined**.

queries about his profession ($e_\ell = \text{Bob.Dylan}, r = \text{has_profession}$) while varying the profession e_r . This illustrates how the representation of *Bob Dylan* varies across his different professional guises.

The table illustrates the utility of token embeddings in inference. Token embeddings of *George W. Bush* and *Barack Obama* in the context of a query about their having held the office of U.S. president are in semantic neighborhoods with a greater density of true instances of U.S. Presidents than their type embeddings. The negative examples *John McCain* and *Hilary Rodham Clinton* have type embeddings that are close to actual presidents. This is sensible since, for instance, Hilary Rodham Clinton is married to Bill Clinton—one of her nearest neighbors. But both of the negative examples’ token embeddings have neighborhoods that are mostly cleared of actual presidents—despite having type embedding neighborhoods that are relatively dense with presidents.

Turning to the second half of the table, we note that *Bob Dylan*’s type embedding is already in a neighborhood dense with singer-songwriters. It is appropriate, then, that this neighborhood undergoes no change apart from minor re-ranking when the triplet (*Bob.Dylan,has_profession,singer-songwriter*) is optimized. For more difficult cases, however, where *Dylan* is not a prototypical example, the semantic neighborhoods undergo dramatic reconfiguration. For instance, optimizing the triplet (*Bob.Dylan,has_profession,disc_jockey*) correctly places *Dylan* in the neighborhood of other DJs, despite the implausibility of this association in the neighborhood of his type embedding, which contains no DJs. This places him in the token neighborhood of *Moby*, who is otherwise

quite unlike *Bob Dylan* except in respect of their common career as DJs.

Combined with our finding that optimization yields the most dramatic improvements in the more stringent evaluation criteria (Hits@1 and Hits@3), this suggests that our optimization procedure is particularly helpful in arbitrating between difficult cases. This qualitative observation about the neighborhoods of compositional and supracompositional triplets can be quantified. Using the triplet classification dataset introduced by Socher et al. (2013), which contains an equal number of positive and negative triplets, we find (Table 2) that positive triplets, on average, end up in supracompositional neighborhoods that are more dense in positive examples than their compositional counterparts. On the other hands, negative triplets suffer a decrease in the number of positive triplets in the neighborhoods of their supracompositional embeddings.

To further quantify the role of semantic optimization in inference, we correlated the difference between the Harmony (score) of input triplets pre- and post-optimization with the change in its rank on the FB15K dataset. The change in Harmony is computed as $\Delta\mathcal{H} = \mathcal{H}(\mu(x), x) - \mathcal{H}(x, x)$, i.e. the difference between the Harmony of the token embedding and the Harmony of the type embedding. This comparison is model-internal—it does not compare models trained to do token inference with models trained for type inference. However, it serves as a useful index of the performance gains attributable to the optimization procedure. If optimizing a triplet representation indeed improves its relative position among all candidate triplets, we expect changes in Harmony to be negatively correlated with the change in rank of positive

	Δ density	t statistic	p
POS	0.241	$t = 99.7$	$p \ll 10^{-10}$
NEG	-0.059	$t = -62.4$	$p \ll 10^{-10}$

Table 2: Change in neighborhood (top-5 closest neighbors) density of true triplets (Δ density) for positive and negative triplets drawn from the triplet classification dataset introduced by Socher et al. (2013), which is derived from the FB15K test set and consists of 59,071 positive triplets and the same number of negative triplets. This resulted in $N = 118,142$ queries for both positive and negative examples (two for each triplet, querying both the left and right entity). After computing each triplet’s neighborhood, we counted the number of triplet neighbors that were in fact in the training, validation, or test sets of FB15K, yielding a measure of the concentration of true and false examples in the neighborhood of both type and token triplet embeddings.

triplets. Consistent with this, we find that optimization leads to significant improvements in raw rank in our best trained HHOLE model (Spearman’s $\rho = -0.0157, p < 10^{-6}$, Figure 5). When considering the change in Mean Reciprocal Rank, a more standard evaluation metric, we find that $\Delta\mathcal{H}$ is positively associated with improvements in MRR ($\rho = .1370, p \ll 10^{-10}$),² particularly when triplets whose ranks do not change at all are omitted ($\rho = .3746, p \ll 10^{-10}$). In other words, when semantic optimization makes a difference, it does so for the better.

For HDISTMULT, $\Delta\mathcal{H}$ is significantly associated with increases in the rank of true triplets ($\rho = 0.1226, p \ll 10^{-10}$), a result consistent with our finding that this class of models disprefers low settings of λ . This illustrates the importance of choices of representational format for embeddings of semantic data. Our optimization procedure can only operate over information that is contained in its compositional input. Hence, choices about how to combine the learned features of entities and relations—i.e. about the manner of composition—are central to our framework.

5.1 Desiderata of a composition function

What factors affect the success of semantic optimization in combination with a particular composition scheme? We suspect that multiplicative interactions across embedding components—which are present in HHOLE and absent in HDISTMULT—are essential for our optimization procedure to contribute helpfully to inference. Both

² Δ MRR is computed as $\text{MRR}(\mu(x)) - \text{MRR}(x)$, i.e. the difference between the Mean Reciprocal Rank of the supracompositional and the compositional triplets.

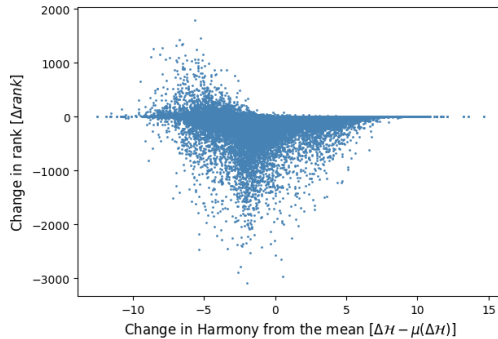


Figure 2: Effect of optimization on the rank of FB15K validation set triplets ($N = 100,000$; 50,000 triplets with two queries per triplet) from the best-performing HHOLE model ($d = 512, \lambda = 1.0$). The horizontal axis is a triplet’s change in Harmony pre- and post-optimization ($\Delta\mathcal{H} \equiv \mathcal{H}(\mu(x), x) - \mathcal{H}(x, x)$) minus the mean change in Harmony for all triplets ($\mu(\Delta\mathcal{H})$). This is plotted against $\text{rank}_q(\mu(x)) - \text{rank}_q(x)$, the triplet’s change in rank due to optimization for query q . A negative correlation indicates reductions in rank (improvements) associated with increasing optimization of triplet representations.

DISTMULT and HOLE are special cases of contracted Tensor Product Representations (TPRs), obtained by summing over (HOLE) or discarding (DISTMULT) terms from the three-way tensor product $e_\ell \otimes r \otimes e_r$.³ In particular, DISTMULT retains only multiplicative interactions *within* components, omitting terms with non-matching indices. This fact appears to be crucial. In a follow-up experiment, we implemented a series of full TPR models trained on FB15K, using the composition operation (14):

$$(14) \quad x_{\text{HTPR}} = e_\ell \otimes r \otimes e_r$$

Such models are necessarily small in size due to the rapid growth of dimensionality for TPRs as a function of the dimensionality of entity and relation embeddings. Consequently, their performance is also poor in comparison to our other implementations. However, the trend matched that which we observed within the HHOLE class: models including the optimization procedure consistently outperformed those with $\lambda = \infty$ (see Table 3). From this, we conclude that other embedding-based KBC models incorporating cross-component multiplicative interactions are likely to see improvements from the addition of a semantic optimization step prior to scoring.

³See (Nickel et al., 2016) for discussion of holographic embeddings as compressed tensor products.

λ	MR	MRR	H@1	H@3	H@10
∞	.150	.278	.192	.305	.447
1.0	.134	.295	.204	.326	.471

Table 3: Performance of HTPR models with and without optimization (controlled by λ). For both models, entities were 5-dimensional and relations 20-dimensional. This trend held across other hyperparameter settings.

6 Conclusion

In this paper, we proposed Gradient Graphs, a general method for augmenting compositional representations of Knowledge Graphs with a post-composition procedure that optimizes the well-formedness of triplet embeddings, highlighting the model’s connection to Harmonic Grammar and Optimality Theory. The resulting model shows marked improvements over the compositional models it is implemented alongside, and also produces triplet token embeddings with properties that prove useful for inference about knowledge base entities. In future work, we intend to explore the utility of semantically-optimized token embeddings in other linguistic settings.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., and Zheng, X. (2016). Tensorflow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI’16, pages 265–283, Berkeley, CA, USA. USENIX Association.
- Belanger, D. and Kakade, S. (2015). A linear dynamical system model for text. In *ICML 32*, volume 37.
- Bollacker, K., Evans, C., Paritosh, P., Sturge, T., and Taylor, J. (2008). Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’08, pages 1247–1250, New York, NY, USA. ACM.
- Bordes, A., Glorot, X., Weston, J., and Bengio, Y. (2014). A semantic matching energy function for learning with multi-relational data. *Machine Learning*, 94(2):233–259.
- Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. (2013). Translating embeddings for modeling multi-relational data. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 26*, pages 2787–2795. Curran Associates, Inc.
- Bordes, A., Weston, J., Collobert, R., and Bengio, Y. (2011). Learning structured embeddings of knowledge bases. In *AAAI*.
- Dasigi, P., Ammar, W., Dyer, C., and Hovy, E. (2017). Ontology-aware token embeddings for prepositional phrase attachment. In *ACL55*, pages 2089–2098.
- Ji, G., Liu, K., He, S., and Zhao, J. (2016). Knowledge graph completion with adaptive sparse transfer matrix. In *AAAI-16*, pages 985–991.
- Kadlec, R., Bajgar, O., and Kleindienst, J. (2017). Knowledge base completion: Baselines strike back. In *2nd Workshop on Representation Learning for NLP*, pages 69–74.
- Kingma, D. and Ba, J. (2015). Adam: A method for stochastic optimization. In *ICLR*.
- Lin, Y., Liu, Z., Sun, M., Liu, Y., and Zhu, X. (2015). Learning entity and relation embeddings for knowledge graph completion. In *AAAI*, pages 2181–2187.
- Ling, X. and Weld, D. (2012). Fine-grained entity recognition. In *AAAI 26*.
- Miller, G. A. (1995). Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41.
- Nickel, M., Rosasco, L., and Poggio, T. (2016). Holographic embeddings of knowledge graphs. In *AAAI*.
- Nickel, M., Tresp, V., and Kriegel, H.-P. (2011). A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on Machine Learning*.
- Plate, T. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, 6:623–641.
- Schlichtkrull, M., Kipf, T. N., Bloem, P., Berg, R. v. d., Titov, I., and Welling, M. (2017). Modeling relational data with graph convolutional networks. *arXiv preprint arXiv:1703.06103*.
- Smolensky, P. and Legendre, G. (2006). *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar*, volume 1: Cognitive Architecture. The MIT Press.
- Socher, R., Chen, D., Manning, C. D., and Ng, A. (2013). Reasoning with neural tensor networks for knowledge base completion. In *NIPS*.
- Toutanova, K., Chen, D., Pantel, P., Poon, H., Choudhury, P., and Gamon, M. (2015). Representing text for joint embedding of text and knowledge bases. In *EMNLP. ACL – Association for Computational Linguistics*.

US Presidents

George W. Bush			Barack Obama		
n	x (compositional)	\hat{h} (optimized)	n	x (compositional)	\hat{h} (optimized)
1	George H. W. Bush	George H. W. Bush	1	Hillary Rodham Clinton	George W. Bush
2	Bill Clinton	Bill Clinton	2	Al Gore	Bill Clinton
3	Jimmy Carter	Jimmy Carter	3	George W. Bush	John F. Kennedy
4	John F. Kennedy	Ronald Reagan	4	Bill Clinton	Ronald Reagan
5	Ronald Reagan	Barack Obama	5	John F. Kennedy	George H. W. Bush

John McCain			Al Gore		
n	x (compositional)	\hat{h} (optimized)	n	x (compositional)	\hat{h} (optimized)
1	John Kerry	John Kerry	1	Barack Obama	Condoleezza Rice
2	Hillary Rodham Clinton	Colin Powell	2	George W. Bush	John C. Calhoun
3	Colin Powell	Nancy Pelosi	3	Colin Powell	Colin Powell
4	Richard Nixon	Joe Biden	4	Condoleezza Rice	Hillary Rodham Clinton
5	Herbert Hoover	Dick Cheney	5	John F. Kennedy	John Kerry

Guises of Bob Dylan

Singer-Songwriter			Screenwriter		
n	x (compositional)	\hat{h} (optimized)	n	x (compositional)	\hat{h} (optimized)
1	Eric Clapton	Bonnie Raitt	1	John Lennon	John Lennon
2	Bonnie Raitt	Eric Clapton	2	Jimi Hendrix	Barbara Streisand
3	Van Morrison	Van Morrison	3	Barbara Streisand	Eric Idle
4	B.B. King	B.B. King	4	Eric Clapton	Nick Cave
5	Bob Seger	Bob Seger	5	Eddie Vedder	Alan Bergman

Disc Jockey			Writer		
n	x (compositional)	\hat{h} (optimized)	n	x (compositional)	\hat{h} (optimized)
1	Tom Petty	Steven Van Zandt	1	John Lennon	Alanis Morissette
2	Warren Zevon	Erykah Badu	2	Alanis Morissette	John Lennon
3	Willie Nelson	Alice Cooper	3	Paul McCartney	Leonard Cohen
4	John Mayer	John Mayer	4	Tina Turner	Leonard Bernstein
5	Steve Earle	Moby	5	Dolly Parton	Prince

Table 4: Semantic neighborhoods of type (pre-) and token (post-optimization) triplets output by the best-performing HHOLE model ($d = 512, \lambda = 1.0$). *US Presidents*: Effect of optimization on the semantic neighborhoods of entity embeddings in the context of the query ($\cdot, \text{office_title}, \text{US_President}$). *Guises of Bob Dylan*: Effect of optimization on the semantic neighborhood of Bob.Dylan in the context of four queries about his profession: Bob Dylan as **singer-songwriter**, **screenwriter**, **disc.jockey**, and **writer**. **Bob.Dylan** is a positive instance of each of these professions in FB15K. For each entity, we retrieved the 5 closest (Euclidian Distance) compositional triplet embeddings, as well as the five closest triplets, among all candidate triplets, when all these candidates are optimized. Triplet completions that in fact occur in FB15K are marked in **bold**. Human-readable entity names were retrieved from a mapping between Freebase machine IDs and names of Wikipedia articles built by [Ling and Weld \(2012\)](#). See main text for discussion of the results.

Yang, B., Yih, W., He, X., Gao, J., and Deng, L. (2015). Embedding entities and relations for learning and inference in knowledge bases. In *ICLR*.

Yoon, H., Song, H., Park, S., and Park, S. (2016). A translation-based knowledge graph embedding preserving logical property of relations. In *Proceedings of NAACL-HLT*, pages 907–916.

7 Appendix A: Model details

(15) **Claim:** $\mu(x) = \frac{1}{2}(\mathbb{W} - \lambda I)^{-1}(b + 2\lambda x)$ is the unique global optimum for $\mathcal{H}(h, x)$ for any fixed x .

We define the Harmony of hidden state h with respect to triplet embedding x as in (1):

$$\begin{aligned} \mathcal{H}(h, x) &\equiv \frac{1}{2} [h^\top \mathbb{W} h + b^\top h - \lambda(h - x)^\top (h - x)] \\ &= \frac{1}{2} [h^\top (\mathbb{W} - \lambda I) h + (b + 2\lambda x)^\top h - \lambda x^\top x] \\ &\equiv \frac{1}{2} [h^\top V h + m(x)^\top h - \lambda x^\top x] \end{aligned}$$

Completing the square yields:

$$\begin{aligned} \mathcal{H}(h, x) &= \frac{1}{2} \left[\left(h - \frac{1}{2} V^{-1} m(x) \right)^\top V \left(h - \frac{1}{2} V^{-1} m(x) \right) \right] \\ &\quad + \frac{1}{2} \left[-\lambda x^\top x - \frac{1}{4} m(x)^\top V^{-1} m(x) \right] \\ &\equiv \frac{1}{2} \left[(h - \mu(x))^\top V (h - \mu(x)) \right] + \ell(x) \end{aligned}$$

This is valid because $V = \mathbb{W} - \lambda I$ is symmetric. $\ell(x)$ does not depend on h , so it is sufficient to optimize $\frac{1}{2} \left[(h - \mu(x))^\top V (h - \mu(x)) \right]$. Setting $\frac{\partial \mathcal{H}(h, x)}{\partial h} = 0$ yields $2V(h - \mu(x)) = 0$; $\therefore h = \mu(x)$. Since V is negative-definite, this point is a *maximum*.

The truth of the claim may also be perceived from observing that $\mathcal{H}(h, x)$ defines a Gaussian distribution over the hidden state variable h with mean $\mu(x)$ and inverse covariance matrix $\Sigma^{-1} \equiv V$. The optimality of $\mu(x)$ then follows from the unimodality of Gaussians.

The training objective (13) may be justified by the following considerations. We take the compositional triplet data to be generated by hidden states of the gradient graph network, and maximize the log probability of the training data using the maximum a posteriori point estimate of the hidden state h . The “complete data” are then $\mathcal{D} = \{\langle \hat{h}, x \rangle\} = \{\langle \mu(x), x \rangle\}$. For fixed x , $\mathcal{H}(h, x)$ models the conditional distribution $p(h|x)$, with

$$(16) \quad p(h|x) = \frac{\exp\{\mathcal{H}(h, x)\}}{Z(x)}$$

where $Z(x) = \int_{h'} \exp\{\mathcal{H}(h', x)\} dh' = |2\pi V^{-1}|^{\frac{1}{2}} \exp\{\ell(x)\}$ is the partition function

conditioned on x . Let $\chi_q = \{x'\}$ be the set of candidate triplet embeddings consistent with a given query q . Choosing the discrete distribution $p(x) = \frac{\exp\{\ell(x)\}}{\sum_{x' \in \chi_q} \exp\{\ell(x')\}}$ over triplet embeddings as the prior probability of the embedding x , we have that:

$$p(\mu(x), x|q) \propto \frac{\exp\{\mathcal{H}(\mu(x), x)\}}{|2\pi V^{-1}|^{\frac{1}{2}} \sum_{x' \in \chi_q} \exp\{\ell(x')\}}$$

For given parameters, the denominator is constant. So, renormalizing over the discrete triplets χ_q gives:

$$p(\mu(x), x|q) = \frac{\exp\{\mathcal{H}(\mu(x), x)\}}{\sum_{x' \in \chi_q} \exp\{\mathcal{H}(\mu(x), x)\}}$$

Approximating the discrete distribution over all of χ_q with a negative sample yields the objective (13).

8 Appendix B: Implementation details

In initial experiments, we searched through a number of candidate models. These included two Harmonic variants of the RESCAL model (Nickel et al., 2011), as well as models that constructed x as a simple concatenation of entity and relation vectors, as well as three-way tensor products of these vectors. These initial experiments led us to focus on DISTMULT and HOLE as the best-performing candidates. Our Harmonic models and reimplementations of the DISTMULT and HOLE baselines were written in TensorFlow (Abadi et al., 2016) and estimated using the Adam optimizer (Kingma and Ba, 2015). With the exception of the HOLE reimplementation, we uniformly used the log-softmax loss (13), which performed best in initial experiments. In contrast, Yang et al. (2015) use a margin-based ranking loss that is linear in the margin between the scores of positive and negative examples up to a threshold, and Nickel et al. (2016) use the pairwise linear margin loss applied to the scores squashed by the logistic function. For HOLE, we used the *linear* margin loss, which provided by far the best performance in the experiments. For each model, we trained until performance on the validation set decrease, then chose the best-performing embedding size from among $d \in \{256, 512\}$. Batch size (512), negative sampling rate (500), and learning rate (0.001) were kept constant across models. We note in passing that regions of the hyperparameter space for DISTMULT explored by Kadlec et al. (2017) were inaccessible to us for technical reasons. For the Harmonic models, we manually tuned the λ hyperparameter.