

End-to-End Learning on Multimodal Knowledge Graphs

Xander Wilcke^{a,*}, Peter Bloem^a, Victor de Boer^a and Rein van 't Veer^b

^a *Faculty of Sciences, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands*

E-mails: w.x.wilcke@vu.nl, p.bloem@vu.nl, v.de.boer@vu.nl

^b *Geodan, Amsterdam, The Netherlands*

E-mail: rein.van.t.veer@geodan.nl

Abstract. Knowledge graphs enable data scientists to learn end-to-end on heterogeneous knowledge. However, most end-to-end models solely learn from the relational information encoded in graphs' structure: raw values, encoded as literal nodes, are either omitted completely or treated as regular nodes without consideration for their values. In either case we lose potentially relevant information which could have otherwise been exploited by our learning methods. We propose a multimodal message passing network which not only learns end-to-end from the structure of graphs, but also from their possibly diverse set of multimodal node features. Our model uses dedicated (neural) encoders to naturally learn embeddings for node features belonging to five different types of modalities, including numbers, texts, dates, images and geometries, which are projected into a joint representation space together with their relational information. We implement and demonstrate our model on node classification and link prediction for artificial and real-worlds datasets, and evaluate the effect that each modality has on the overall performance in an inverse ablation study. Our results indicate that end-to-end multimodal learning from any arbitrary knowledge graph is indeed possible, and that including multimodal information can significantly affect performance, but that much depends on the characteristics of the data.

Keywords: End-to-End Learning, Knowledge Graphs, Multimodal, Heterogeneous Knowledge, Message Passing Networks

1. Introduction

The recent adoption of knowledge graphs by multinationals such as Google and Facebook has made them interesting targets for various machine learning applications such as link prediction and node classification. Already, this interest has led to the development of message-passing models which enable data scientists to learn end-to-end¹ from any arbitrary graph. To do so, message-passing models propagate information over the edges of a graph, and can therefore be used to exploit the relational information encoded in a graph's structure to guide the learning process. The same approach has also been shown to work quite well on knowledge graphs, obtaining results that are comparable to dedicated models such as *RDF2Vec* [1] and Weisfeiler-Lehman kernels [2]. Nevertheless, by fo-

cus on a single modality—the graphs' structure—we are effectively throwing away a lot of other information that knowledge graphs tend to have, and which, if we were able to include it in the learning process, has the potential of improving the overall performance of our models.

Combining information from multiple modalities is a topic that is already well studied for information stored in *relational* form (for instance in relational database management systems). Here too, we often encounter *heterogeneous* knowledge, containing information from a wide variety of modalities (such as language, audio, or images). In [3], the case is made that

¹In the context of this paper, we define “end-to-end learning” as the use of machine learning models which operate directly on raw data, instead of relying on manually engineered features. In end-to-end learning, any information in the data can, in principle, be used by the model. See [3] for a more in-depth discussion.

*Corresponding Author, E-mail: w.x.wilcke@vu.nl.

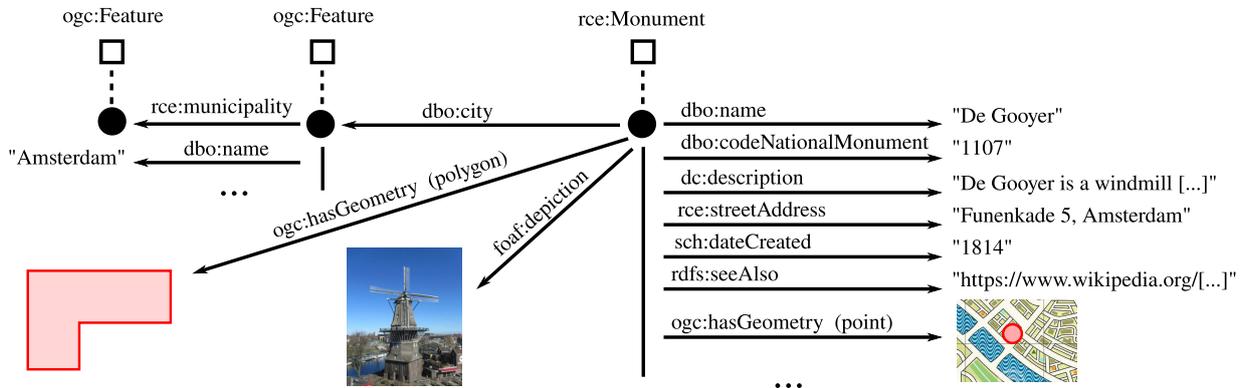


Fig. 1. A simplified and incomplete example from the Dutch Monuments Graph showing a single monument with several attributes of different modalities.

to truly learn *end-to-end* from a collection of heterogeneous, multimodal data, we must design machine learning models that can consume these data in as raw a form as possible, staying as close as we can to the original knowledge, and that we need to adopt a data model which can represent our data in a suitable format, for which the knowledge graph is a natural choice. In other words, even when our heterogeneous multimodal data is not initially represented as a knowledge graph, transforming it to this format is a natural first step in an end-to-end multimodal machine learning pipeline.

In this paper, we introduce and implement a multimodal message passing neural network, based on this principle, which can directly consume heterogeneous multimodal data, represented as knowledge graph, and which itself can learn to extract relevant information from each modality, based solely on the downstream task.

With the term *knowledge graph* we mean any labeled multidigraph that is built on top of the *Resource Description Framework* (RDF). We consider the relational information of such a graph, encoded in its structure, as a single modality. Other modalities that are commonly present in knowledge graphs are of numerical, textual, and temporal nature, such as various measurements, names, and dates, respectively, and, to a lesser degree, of visual, auditory, and spatial makeup. In a knowledge graph about monuments, for example, we might find that each monument has a detailed description, a registration number, a year in which it was built, a few pictures from different angles, and a set of coordinates (Figure 1). These and other attributes are encoded as raw values with corresponding datatype annotations, called *literals*, and tell us something about

the objects they are connected to, called *entities*. However, most of this information is lost when we reduce the literals to identifiers, as is currently common practice when we apply message passing networks to knowledge graphs.

By reducing literals to identifiers, we discard any information that is contained in their contents, retaining only the relational information encoded by their connections, and placing them on an equal footing with all other entities. This means that we are effectively feeding our models a subset of the original and complete knowledge, but also that we are depriving our models of the ability to compare inputs according to their modalities: measurements as numbers, descriptions as language, coordinates as geometries, etc. As a result, our models are unable to distinguish between literals that are closely together in the value space with those which are far apart. The name *Mary*, for example, would be seen as (dis)similar to *Maria* as it would be to *Bigglesworth*, as would the integer value 47 be to 42 and 6.626068×10^{-34} . Instead however, we want our models to use this information to guide their learning process.

By enabling our models to naturally ingest literal values, and by treating these values according to their modalities, tailoring their encodings to their specific characteristics, we stay much closer to the original and complete knowledge that is available to us. We believe that doing so enables our models to create better internal representations of the entities we are trying to learn over, potentially resulting in an increase in the overall performance of our models. By embedding this principle in the message passing framework, and by exploiting Semantic Web standards such as datatype annotations, we embrace the idea that this enables us to

learn end-to-end from any heterogeneous multimodal knowledge, as long as it is represented as a knowledge graph.

In this work, we propose a multimodal message passing model which incorporates the information from a diverse set of multimodal node features. Our model uses dedicated vectorization strategies and (neural) encoders to naturally learn embeddings for node features belonging to five different types of modalities, including images and geometries, which are projected into a joint representation space together with their relational information. We demonstrate our model on node classification and link prediction for both artificial and real-worlds knowledge graphs, and evaluate the effect that each modality has on the overall performance in an inverse ablation study. We also implement and publish our model as Python package capable of learning from any arbitrary knowledge graph out of the box, exploiting Semantic Web standards to automatically infer and incorporate multimodal information.

To summarize, the main contributions of this paper are:

1. A machine learning model, embedded in the message passing framework, which can learn end-to-end from a heterogeneous knowledge, encoded as a knowledge graph, and which can naturally ingest literal values according to their modalities.
2. An investigation of the potential usefulness of including information from multiple modalities, and the impact this has on the overall performance of our models.
3. An implementation of our model (named the MR-GCN), which can learn from any arbitrary knowledge graph, and which exploits Semantic-Web standards to automatically infer and incorporate multimodal information.

Our intent is emphatically *not* to show that our implementation achieves any kind of state-of-the-art, or even to measure its performance against related models. Rather, we aim to demonstrate that 1) by including as much of the original knowledge as possible, in as natural of a fashion as possible, we can, in certain cases, help our models obtain a better overall performance, and that 2) a model can be trained end-to-end on a heterogeneous knowledge graph such that it learns purely from the downstream task which patterns to extract from each modality.

2. Related Work

Machine learning from multimodal sources is a well-studied problem. A good introduction to this problem and its many perspectives is given by [4]. According to their taxonomy, our approach is one of *late fusion* by first encoding modalities using dedicated neural encoders, after which the resulting encodings are projected in a joint representation space. Different from most other research in this field we are not interested in *translation* (mapping one modality to another) nor in *alignment* (aligning the same subject over multiple modalities). Rather, information in a given modality is only ever used to learn node embeddings with the intent to improve the learning process by including as much of the original knowledge as possible.

2.1. Knowledge Graph Embeddings

Graph embedding techniques aim to represent graphs in a lower-dimensional space, making them more suitable to learn over. Numerous embedding techniques have been proposed over the years, and typically differ in which operations they apply between the node and edge embeddings, and which scoring function they use. Popular methods are those based on matrix factorization, random walks, translation models, and, more recently, deep neural networks [5]. Our approach falls in the latter group of methods, for its use of a message-passing network. A thorough overview of the different embedding methods can be found in one of the many recent survey papers, for example [5] and [6]. Here, we will limit ourselves to the graph embedding methods that consider multimodal information.

Various approaches have explored using information from one or more additional modalities in machine learning models for knowledge graphs. In most cases, only a singly additional modality is included, always of numerical, textual, or visual nature [7]. This differs from our method, which also supports temporal and spatial literals. Our method also differs from most other approaches in that we address how information from different modalities can be 1) extracted from a graph, and 2) vectorized with minimal loss of information.

An early work described in [8] proposes an extension to the RESCAL [9] tensor factorization method which can also cope with textual attributes. This is done by introducing an additional tensor which is factorized together with the tensor holding the relational information. A similar separation is proposed by [10],

who generate a separate co-occurrence matrix for the relational and textual information, and which are then summed to produce the final embeddings. Both these methods scale well due to their use of basic matrix operations, whereas scalability remains a challenge for many message-passing models such as the one used in our approach.

In [11], the authors introduce a learnable function, called *LiteralE*, which replaces every entity embedding by a new embedding that is the fusion of the original entity embedding and its direct numerical attributes. The resulting vector representation can then be used in an arbitrary translation-based model. The fusion step is similar to our approach in that the embeddings of neighbouring nodes coalesce into the target entity, except that our model does this for every node (entity or literal), up to an arbitrary depth (determined by the number of layers in the message-passing network), and only after the modalities have been encoded according to their specific characteristics.

The authors of [12] propose an extension to *LiteralE* that incorporates textual features which they generate by performing entity resolution on (part of) the identifiers of entities and relations. The results are then mapped to integers and passed to *LiteralE* together with the corresponding entities.

A slightly different approach is proposed by [13], who perform a joint optimization of an existing translation model (*TransE* [14]) and a regression model specifically designed by the authors for numerical features. The work in [15] uses a similar approach, but for textual rather than numerical attributes and with a self-defined translation model instead of a regression model. Similar to our work, the authors use a CNN as encoder for textual attributes, but where our model employs a temporal CNN with one-hot encoded text as input, the authors here use a language-agnostic CNN with pretrained *word2vec* [16] embeddings as input.

Another extension to an arbitrary translation model is proposed in [17], who use a proven CNN architecture to learn image embeddings, which are then used in a self-defined translation model. For entities with more than one image attribute, the images embeddings are merged into one final embedding which is kept separate from the entity embedding to which they belong. Our model differs in that all neighbouring nodes, and not just images, coalesce into the corresponding entity embedding: separate image embeddings only exist prior to fusion.

Different from translation-based approaches is the work in [18], who propose using a dual network archi-

ture with a binary classifier to learn relational information and a regression model to learn numerical information. A joint optimization is used to train the model.

More modalities are considered by [19], who incorporate numerical and textual literals, as well as images. The numerical features are encoded using a feed-forward layer, which projects the values to a higher-dimensional space. For short strings, the authors employ a character-based GRU, whereas a language-aware CNN is used in combination with word sequences for longer strings. Finally, for images, the authors use the last hidden layer of a pretrained network on ImageNet [20]. The resulting embeddings are then paired with their corresponding entity embeddings (generated using a feed-forward network) and ultimately scored using *DistMult*. The use of dedicated neural encoders per modality is similar to our work, except for numerical features, which we feed directly to the message-passing network after normalization. Also similar is the use of different encoders for text of different lengths, but rather than have completely different models and input requirements, we employ three temporal CNNs of increasing size for short, medium, and long strings.

All the reviewed models are simple embedding models, based on basic matrix operations or on a score function applied to triples. By contrast, our approach includes a message passing layer, allowing multimodal information to be propagated through the graph, several hops and from *all* (direct and indirect) neighbours.

3. Preliminaries

Knowledge graphs and message passing neural networks are integral components of our research. We will here briefly introduce both concepts.

3.1. Knowledge Graphs

For the purposes of this paper we define a *knowledge graph* $G = (\mathcal{V}, \mathcal{E})$ over modalities $1, \dots, \mathcal{M}$ as a labeled multidigraph defined by a set of nodes $\mathcal{V} = \mathcal{I} \cup \bigcup \{\mathcal{L}^m | m \in \mathcal{M}\}$ and a set of directed edges \mathcal{E} , and with $n = |\mathcal{V}|$. Nodes belong to one of two categories: entities \mathcal{I} , which represent objects (monuments, people, concepts, etc.), and literals \mathcal{L}^m , which represent raw values in modality $m \in \mathcal{M}$ (numbers, strings, coordinates, etc.). We also define a set of rela-

tions \mathcal{R} , which contains the edge types that make up \mathcal{E} . Relations are also called *predicates*.

Information in G is encoded as triples \mathcal{T} of the form (h, r, t) , with head $h \in \mathcal{I}$, relation $r \in \mathcal{R}$, and tail $t \in \mathcal{I} \cup \mathcal{L}^1 \cup \dots \cup \mathcal{L}^m$. The combination of relations and literals are also called *attributes* or *node features*.

See Figure 1 for an example of knowledge graph with seven nodes, two of which are entities and the rest literals. All knowledge graphs in this paper are stored in the *Resource Description Framework* (RDF) format [21], but our model can be applied to any graph fitting the above definition.

3.2. Message Passing Neural Networks

A *message passing neural network* [22] is a graph neural network model that uses trainable functions to propagate node embeddings over the edges of the neural network. One simple approach to message passing is the graph convolutional neural network (GCN) [23]. The R-GCN [24], on which we build, is a straightforward extension to the knowledge graph setting.

Let \mathbf{H}^0 be a $n \times q$ matrix of q dimensional node embeddings for all n nodes in the graph. That is, the i -th row of \mathbf{H}^0 is an embedding for the i -th node in the graph². The R-GCN computes an updated $n \times l$ matrix \mathbf{H}^1 of l -dimensional node embeddings by the following computation (the *graph convolution*):

$$\mathbf{H}^1 = \sigma \left(\sum_{r \in \mathcal{R}} \mathbf{A}^r \mathbf{H}^0 \mathbf{W}^r \right) \quad (1)$$

Here, σ is an activation function like ReLU, applied element-wise. \mathbf{A}^r is the row-normalised adjacency matrix for the relation r and \mathbf{W}^r is a $q \times l$ matrix of learnable weights. This operation arrives at a new node embedding for a node by averaging the embeddings of all its neighbours, and linearly projecting to l dimensions by \mathbf{W}^r . The embeddings are then summed over all relations and a non-linearity σ is applied.

To allow information to propagate in both directions along an edge, all inverse relations are added to the predicate set. The identity relation is also added (for which $\mathbf{A}^r = \mathbf{I}$) so that the information in the current embedding can, in principle, be retained. To reduce overfitting, the weights \mathbf{W}^r can be derived from

²The standard R-GCN does not distinguish between literals and entities. Also, literals with the same value are collapsed into one node, therefore $n \leq |\mathcal{V}|$.

a smaller set of *basis weights* by linear combinations (see the original paper for details).

To use R-GCNs for entity classification with c classes, the standard approach is to start with one-hot vectors as initial node embeddings (that is, $\mathbf{H}^0 = \mathbf{I}$). These are transformed to h -dimensional node embeddings by a first R-GCN layer, which are transformed to c -dimensional node embeddings by a second R-GCN layer. The second layer has a row-wise softmax non-linearity, so that the final node embeddings can be read as class probabilities. The network is then trained by computing the cross-entropy loss for the known labels and backpropagating to update the weights. Using more than two layers of message passing does not commonly improve performance with current message passing models.

For link prediction, the R-GCNs can be viewed as encoder in a graph auto-encoder. In that role, the R-GCNs learn node embeddings that are used by a decoder to reconstruct the edges in the graph. As before, the standard approach for the R-GCNs is to have one or two layers, and to start with one-hot vectors as initial node embeddings. However, because we are now interested in the node embeddings themselves, the softmax on the end is replaced with an activation function like ReLU, applied element-wise. The decoder consists of a triple scoring function $s : \mathcal{V} \times \mathcal{R} \times \mathcal{V} \mapsto \mathbb{R}$, for which ideally holds that $s(h, r, t) > s(x, y, z)$ if (h, r, t) exists and (x, y, z) does not.

In this work, we use DistMult [25] for our decoder, which is known to perform well on link prediction tasks while keeping the number of parameters low [26]. DistMult uses the following bilinear scoring function:

$$s(\mathbf{y}_{v_i}, r, \mathbf{y}_{v_j}) = \mathbf{y}_{v_i}^T \text{diag}(\mathbf{R}_r) \mathbf{y}_{v_j} \quad (2)$$

Here, \mathbf{y}_{v_i} and \mathbf{y}_{v_j} are the output of the encoder for nodes $v_i, v_j \in \mathcal{V}$, and \mathbf{R}_r the embedding belonging to relation $r \in \mathcal{R}$. Both encoder and decoder are trained by minimizing the binary-cross entropy loss³ over the output of Equation 2 for both positive and negative samples (negative sampling) [24]. The set of negative samples \mathcal{T}^- can be obtained by randomly corrupting the head or tail of a portion ($\frac{1}{5}$) of the triples in \mathcal{T} .

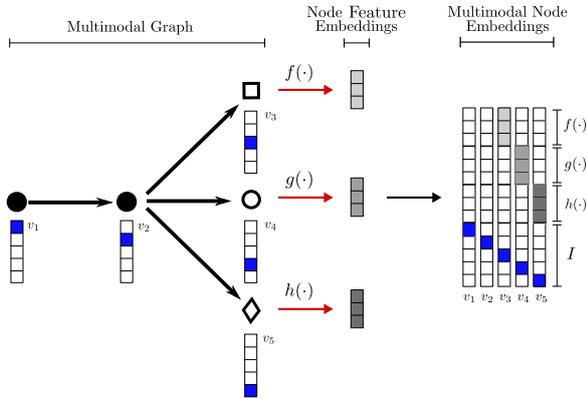


Fig. 2. Overview of how our model creates multimodal node embeddings for nodes v_1 to v_5 . Solid circles represent entities, whereas open shapes represent literals of different modalities. The nodes’ feature embeddings are learned using dedicated (neural) encoders (here f , g , and h), and concatenated to their identity vectors I to form multimodal node embeddings, which are fed to a message passing network.

4. A Multimodal Message Passing Network

We introduce our model as an extension to message passing networks which can learn end-to-end from the structure of an arbitrary graph, and for which holds that $\mathbf{H}^0 = \mathbf{I}$. To do so, we let $f(\cdot)$, $g(\cdot)$, and $h(\cdot)$ be feature encoders that output feature embeddings of lengths ℓ_f , ℓ_g , and ℓ_h for nodes $v_i \in \mathcal{V}$. We define \mathbf{F} as the $n \times f$ matrix of multimodal feature embeddings with $f = \ell_f + \ell_g + \ell_h$, and concatenate \mathbf{F} to the identity matrix \mathbf{I} to form multimodal *node* embeddings:

$$\mathbf{H}^0 = [\mathbf{I} \mathbf{F}] \quad (3)$$

of size $n \times q$ (Fig. 2).

Embedding matrix \mathbf{H}^0 is fed together with \mathbf{A}^r to a message passing network, such as the R-GCN. Both encoders and network are trained end-to-end in unison by backpropagating the error signal from the network through the encoders all the way to the input.

4.1. Modality Encoders

We consider five different modalities which are commonly found in knowledge graphs. We forgo discussing relational information—the sixth modality—as that is already extensively discussed in related work on message passing networks. For numerical information, we use a straightforward one-to-one encoding and

let the message-passing layers handle it further. For all other modalities we use neural encoders: a feed-forward neural network for temporal information, and convolutional neural networks (CNN) for textual, spatial, and visual information. Each of these will be discussed next. We will also discuss the preceding vectorization process, which, if done poorly, can result in a loss of information.

In the following, we let \mathbf{e}_i^m be the embedding vector of node v_i for modality m . The concatenation of a node’s identity vector and all its feature embedding vectors \mathbf{e}_i^m for every $m \in \mathcal{M}$ equals the i -th row of \mathbf{H}^0 .

4.1.1. Numerical Information

Numerical information encompasses the set of real numbers \mathbb{R} , and corresponds to literal values with a datatype declaration of XSD:double, XSD:float, and XSD:decimal and any subtype thereof. For these, we can simply take the normalized values as their embeddings, and feed these directly to the message-passing layers. We also include values of the type XSD:boolean into this category, but separate their representations from those of real numbers to convey a difference in semantics.

More concretely, for all nodes $v_i \in \mathcal{V}$ holds that \mathbf{e}_i^{num} is the concatenation of their numerical and boolean components, encoded by functions f_{num} and f_{bool} , respectively. Here, $f_{num}(v_i) = v_i$ if v_i is a literal node with a value in \mathbb{R} . If v_i is a boolean instead, we let $f_{bool}(v_i)$ be 1.0 if v_i is true and -1.0 if v_i is false. In both cases, we represent missing or erroneous values with 0.0 (we assume a normalization between -1 and 1).

4.1.2. Temporal Information

Literal values with datatypes which follow the *Seven-property model*⁴ such as XSD:time, XSD:date and XSD:gMonth, are treated as temporal information. Different from numerical values, temporal values contain elements that are defined in a circular value space and which should be treated as such. For example, it is inaccurate to treat the months December and January as if they were 11 months apart, as would be implied by directly feeding the months’ number to our models. Instead, we can represent this as

$$f_{trig}(\phi, \psi) = \left[\sin\left(\frac{2\pi\phi}{\psi}\right), \cos\left(\frac{2\pi\phi}{\psi}\right) \right] \quad (4)$$

with ψ the number of elements in the value space (here 12), ϕ the integer representation of the element we want to encode, and f_{trig} a trigonometric function in

⁴A margin ranking loss is used in the original DistMult paper.

our encoder. This ensures that the representation of January is closer to that of December than it is to that of March.

We can use this representation for all other circular elements, such as hours ($\psi = 24$) and decades ($\psi = 10$). When dealing with years however, we represent smaller changes more granular than larger changes: years are split into centuries, decades, and (single) years fragments, with decades and years treated as circular elements but with centuries as numerical values (we limit our domain to years between -9999 and 9999).

Once vectorized, the vector representation \mathbf{v}_i is fed to a feed-forward neural network f_{emp} with input and output dimensions n_{in} and n_{out} , respectively, and for which holds that $n_{in} < n_{out}$, such that $\mathbf{e}_i^{temp} = f_{emp}(\mathbf{v}_i)$.

4.1.3. Textual Information

Vector representations for textual attributes with the datatype `XSD:string`, or any subtype thereof, and `XSD:anyURI` are created using a character-level encoding, as proposed in [27]. For this purpose, we let \mathbf{E}^s be a $|\Omega| \times |s|$ matrix representing string s using vocabulary Ω , such that $\mathbf{E}_{ij}^s = 1.0$ if $s_j = \Omega_i$, and 0.0 otherwise.

A character-level representation enables our models to be language agnostic and independent of controlled vocabularies (allowing it to cope with colloquialisms and identifiers for example), as well as provide some robustness to spelling errors. It also enables us to forgo the otherwise necessary stemming and lemmatization steps, which would remove information from the original text. The resulting embeddings are optimized by running them through a temporal CNN f_{char} with output dimension c , such that $\mathbf{e}_i^{textual} = f_{char}(\mathbf{E}^{v_i})$ for every node v_i with a textual value.

4.1.4. Visual Information

Images and other kinds of visual information (e.g. videos, which can be split in frames) can be included in a knowledge graph by either linking to them or by expressing them as binary string literals⁵ which are incorporated in the graph itself (as opposed to storing them elsewhere). In either case, we first have to obtain the raw image files by downloading and/or converting them.

Let im_i be the raw image file as linked to or encoded by node v_i . We can represent this image as a tensor \mathbf{E}^{im_i}

of size $channels \times width \times height$, which we can feed to a two-dimensional CNN f_{im} with output dimension c , such that $\mathbf{e}_i^{visual} = f_{im}(\mathbf{E}^{im_i})$ for the image associated with node v_i .

4.1.5. Spatial Information

Spatial information includes points, polygons, and any other spatial features that consist of one or more coordinates. These features can represent anything from real-life locations or areas to molecules or more abstract mathematical shapes. Literals with this type of information are commonly expressed using the *well-known text representation* (WKT) and carry the `OGC:wktLiteral` datatype declaration. The most elementary spatial feature is a coordinate (point geometry) in a d -dimensional space, expressed as `POINT($x_1 \dots x_d$)`, which can be combined to form more complex types such as lines and polygons.

We can use the vector representations proposed in [29] to represent spatial features. Let \mathbf{E}^{sf} be the $|\mathbf{x}| \times |sf|$ matrix representation for spatial feature sf consisting of $|sf|$ coordinates, and with \mathbf{x} the vector representation of one such coordinate. Vector \mathbf{x} holds all of the coordinate’s d points, followed by its other information (e.g. whether it is part of a polygon) encoded as binary values. For spatial features with more than one coordinate, we also need to separate their location from their shape to ensure that we capture both these components. To do so, we encode the location in \mathbb{R}^d by taking the mean of all coordinates that makeup the feature. To capture the shape, we compute the global mean of all spatial features in the graph, and subtract this from their coordinates to place their centre around the origin.

We feed the vector representations using a temporal CNN f_{sf} with output dimension c , such that $\mathbf{e}_i^{spatial} = f_{sf}(\mathbf{E}^{v_i})$ for all nodes v_i which express spatial features.

5. Implementation

We implement our model using the R-GCN as our main building block, onto which we stack our various encoders. We call this a multimodal R-GCN (MR-GCN). The R-GCN is a suitable choice for this purpose, as it can learn end-to-end on the structure of relational graphs, taking relation types into account. Our implementation is available as Python package⁶, and can be used with any arbitrary knowledge graph in RDF format.

⁴<https://www.w3.org/TR/xmlschema11-2>

⁵In [28], we advocate the use of KGBench’s `base64Image` for this purpose.

⁶Code available at <https://gitlab.com/wxwilcke/mrgcn>

In the simplest case, when we are only interested in learning from the graph’s structure or when no multimodal information is present in the graph, we let the initial node embedding matrix \mathbf{H}^0 be the nodes’ $n \times n$ identity matrix \mathbf{I} (i.e. $\mathbf{H}^0 = \mathbf{I}$). This reduces the MR-GCN to a plain R-GCN. To also include multimodal information in the learning process, we let \mathbf{F} be the $n \times f$ feature embedding matrix instead and concatenate this to \mathbf{H}^0 as in Equation 3 to form $\mathbf{H}^0 = [\mathbf{I} \ \mathbf{F}]$.

To accurately determine the most suitable encoder for each encountered literal, the MR-GCN exploits Semantic-Web standards to automatically infer this from the graph’s datatype annotations. Supported datatypes include many XSD classes, such as numbers, strings, and dates, as well as OGC’s `wktLiteral` for spatial information, and KGbench’s `base64Image` for binary-encoded images [28]. These modalities are assumed to be encoded directly in the graph, as opposed to reading them from separate files.

To cope with the increased complexity brought on by including node features we optimized the MR-GCN for sparse matrix operations by splitting up the computation of Equation 1 into the sum of the structural and feature component. For this, we once again split \mathbf{H}^0 into identity matrix $\mathbf{H}_I = \mathbf{I}$ and feature matrix $\mathbf{H}_F = \mathbf{F}$, and rewrite the computation as

$$\mathbf{H}^1 = \sigma \left(\sum_{r \in \mathcal{R}} \mathbf{A}^r \mathbf{H}_I \mathbf{W}_I^r + \mathbf{A}^r \mathbf{H}_F \mathbf{W}_F^r \right) \quad (5)$$

Here, \mathbf{W}_I^r and \mathbf{W}_F^r are the learnable weights for the structural and feature components, respectively. For layers $i > 0$ holds that $\mathbf{H}_F^i = \mathbf{H}^i$, and that $\mathbf{A}^r \mathbf{H}_I \mathbf{W}_I^r = 0$. Note that because $\mathbf{A}^r \mathbf{H}_I = \mathbf{A}^r$, we can omit this calculation when computing Equation 5, and thus also no longer need \mathbf{H}_I as input. Figure 3 illustrates this computation as matrix operations.

To support link prediction, the MR-GCN implements the DistMult [25] bilinear scoring function, shown in Equation 2. To reduce the number of parameters, we simulate relation embeddings $\text{diag}(\mathbf{R})$ by a $|\mathcal{R}| \times h$ matrix, with each row representing the diagonal of a theoretical relation embedding \mathbf{R}_r .

5.1. Neural Encoders

The MR-GCN implements neural encoders for all modalities listed in Section 4.1. For temporal information, we use a single layer fully connected feed-forward neural network of which the dimensions de-

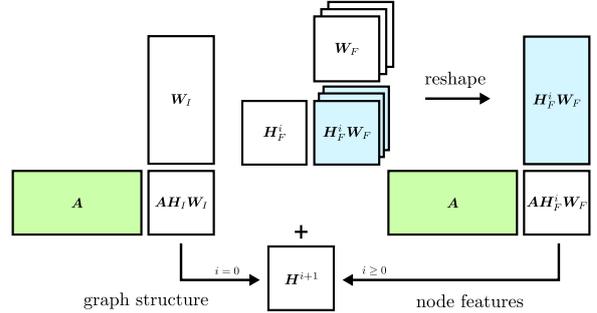


Fig. 3. Graphical depiction of our implementation of Equation 5, shown as matrix operations. The output of layer i , \mathbf{H}^{i+1} , is computed by summing the structure and node feature components. If $i > 0$, then $\mathbf{H}_F^i = \mathbf{H}^i$ and $\mathbf{A} \mathbf{H}_I \mathbf{W}_I = 0$.

Table 1

Configurations of the neural encoder for temporal information with h hidden nodes and output dimension n_{out} , listed per tested datatype. Note that $n_{in} = h$

Datatype	h	n_{out}
XSD:gYear	6	2
XSD:date	10	4
XSD:dateTime	14	6

pend on the datatype, as shown in Table 1. The three other neural encoders are all implemented using CNNs, each initiated using $\mathcal{N}(0, 1)$ and with an output dimension of 128.

For our visual encoder, we use the efficient MobileNet architecture from [30], which provides a good performance with relatively few parameters. For spatial information, we use a temporal CNN similar to that used in [29], which has 3 convolutional layers, each followed by ReLU, and 3 dense layers (Table 3). A similar setup is used for textual information, except that we use different architectures for short ($\ell < 20$), medium ($20 < \ell < 50$), and long ($\ell > 50$) strings, with ℓ denoting their length. The architecture for medium-length strings is listed in Table 2, whereas for long strings we double the number of filters to 128 and let the first dense layer have 1024 hidden nodes. For short strings, we omit the last convolutional and dense layer (layer 4 and 7), and reduce the number of hidden nodes in the first dense layer to 256.

The output of layer i from all encoders for all nodes in \mathcal{V} are concatenated to form \mathbf{H}_F^i , which is passed to Equation 5 together with \mathbf{A}^r .

Table 2

Configuration of the textual encoder for medium-length strings with 4 convolutional layers (top) and 3 dense layers (bottom). For pooling layers, $max(k/s)$ lists kernel size (k) and stride (s), or $max(\cdot)$ when it depends on the input sequence length.

Layer	Filters	Kernel	Padding	Pool
1	64	7	3	$max(2/2)$
2	64	7	3	$max(2/2)$
3	64	7	3	-
4	64	7	2	$max(\cdot)$

Layer	Dimensions
5	512
6	128
7	128

Table 3

Configuration of the spatial encoder with 3 convolutional layers (top) and 3 dense layers (bottom). For pooling layers, $max(k/s)$ lists kernel size (k) and stride (s), whereas $avg(\cdot)$ depends on the input sequence length.

layer	filters	kernel	padding	pool
1	16	5	2	$max(3/3)$
2	32	5	2	-
3	64	5	2	$avg(\cdot)$

layer	dimensions
4	512
5	128
6	128

6. Experiments

We evaluate the MR-GCN on node classification and link prediction while varying the modalities which are included in the learning process⁷. For this purpose, we compute the performance for each combination of structure and modality, as well as all modalities combined, and evaluate this against using only the relational information. To eliminate any confounding factors in real-world knowledge that might influence the results, we will first evaluate the MR-GCN on syn-

thetic knowledge (Section 6.1) before testing our implementation on real-world datasets (Section 6.2).

Another dimension that we vary is how much raw information is already implicitly encoded in the structure of a graph by having literals nodes with an indegree greater than one. This occurs when literals with the same value are coalesced into a single node, and is the standard approach to represent knowledge graphs in graph form. Encoding this information in a graph’s structure influences the potential gain in performance we can obtain by including node features in the learning process, possibly even masking it. Consider, for example, a classification problem in which a small range of literals perfectly separates our classes: when this information is already encoded in the structure there might be little to gain by enabling our models to compare these literals by their values, whereas doing so if this information is *not* encoded in the structure might yield a significant performance boost. In our experiments, we will use the term *split literals* to refer to the representation that keeps literals with the same value as separate nodes (i.e. indegree = 1), and use the term *merged literals* to refer to alternative representation in which literals with the same value are coalesced (i.e. indegree \geq 1).

For our node classification experiments we use an architecture similar to the plain R-GCN (Section 3.2). Concretely, we employ a two-layered MR-GCN with 32 hidden nodes, and with an element-wise ReLU activation function after the first layer. A row-wise softmax non-linearity is added to the second layer to output class probabilities. The network is trained by minimizing the cross-entropy loss in full batch mode with Adam for 400 epochs with an initial learning rate of 0.01.

For each configuration we report the mean classification accuracy and 95% confidence interval over 10 runs. To check the results on statistical significance, we use the Stuart-Maxwell marginal homogeneity test which tests whether two multi-class models have the same distribution of predictions [31, 32]. To obtain a single set of predictions per configuration for this purpose, we use a majority vote amongst the ordered output from all 10 runs.

Our link prediction experiments likewise use a graph auto-decoder architecture similar to the plain R-GCN (Section 3.2). More specific, we employ a single-layered MR-GCN with 200 hidden nodes, with

⁷Datasets available at <https://gitlab.com/wxwilcke/mmkg>

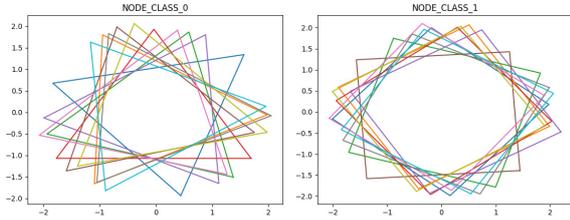


Fig. 4. Geometries belonging to 10 randomly-sampled entities per class from the SYNTH dataset. Apart from the number of points (which our model is agnostic to) the only difference between classes is the shape.

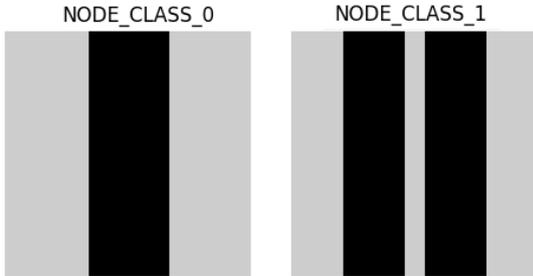


Fig. 5. Images belonging to entities per class from the SYNTH dataset, shown here without the noise normally present to ensure different string representations with a class.

an element-wise ReLU activation function at the end, and with DistMult as triple scoring function. We train the network by minimizing the binary cross-entropy loss in full batch mode with Adam for 1000 epochs with an initial learning rate of 0.01.

For each configuration we report the filtered mean reciprocal rank (MRR) and hits@ k with $k \in \{1, 3, 10\}$ over 5 runs, as well as the 95% confidence interval and statistical significance computed over the MRR⁸. To check for statistical significance, we use the computational-intensive randomised paired t-test [33], as suggested by [34], which tests whether two ordered sets of ranks have the same distribution of mean differences. Note that, with this method, the minimal achievable p-value depends on the size of the test set. As with classification, we obtain a single set of ranks per configuration by majority vote.

6.1. Evaluation on Synthetic Knowledge

We first evaluate the performance of the MR-GCN on synthetic data. These data serve as a controlled environment which enables us to eliminate any confounding factors in real-world data that would otherwise influence the results, ensuring that any observed difference can be confidently attributed to the addition or removal of a certain modality. For this purpose, we generated⁹ a synthetic knowledge graph (SYNTH) that contains strong multimodal signals, but which lacks relational information. General and modality-specific statistics are listed in Table 6 and 7, respectively.

The SYNTH dataset consists of 16,384 entities, all labeled, from two distinctly different classes, and connected by a random graph structure that is generated using the Watts–Strogatz algorithm. Each entity is provided with literals of different datatypes, encompassing all five modalities listed in Section 4.1. To ensure that the learning problem is both manageable and challenging, the literal values were drawn from two narrow and slightly overlapping distributions, with noise added where necessary. These distributions were generated with the corresponding modality in mind: numbers and years were drawn from Gaussian distributions, dates and times were sampled around specific months and hours, respectively, and strings were generated by combining a class-specific keyword with randomly sampled words from a dictionary. This principle is also shown in Figure 4 for geometries, which only differ in shape¹⁰ to force our model to capture this characteristic. Similarly in Figure 5 for images, which are unique per class and to an extent robust to transformations (e.g., scale, rotation, translation).

6.1.1. Node Classification Results

Table 4 reports the mean classification accuracy over 10 runs on SYNTH, together with its 95% confidence interval and corresponding p-values. We use *value_merged* [*value_split*] to express the performances in the merged and split configurations, respectively.

Overall, the results indicate that, for all modalities and literal configurations, including node features considerably increases the performance over that of the

⁸As the hits@ k is derived from the MRR, no new information is gained by also computing the confidence interval and statistical significance of the former.

⁹Code available at <https://gitlab.com/wxwilcke/graphsynth>

¹⁰The neural encoders in our model are agnostic to the number of points.

baseline (structure only). When all node features are taken into account, this performance increase raises the accuracy from near random (0.616 [0.495]) to near perfect (0.995 [0.996]). All reported performance gains are statistically significant, with as highest p-value 5.21×10^{-04} .

When comparing the performance gain per modality it is evident that this differs widely between modalities: including just textual or spatial information increases the performance to a near perfect accuracy of 0.995 [0.996] and 0.957 [0.949], respectively, whereas including only visual information just provides a slight (although still significant) gain to an accuracy of 0.642 [0.556]. The remaining two modalities—numerical and temporal information—lie in between these two extremes and provide a moderate performance boost with an accuracy of 0.744 [0.785] and 0.763 [0.625], respectively. When all modalities are included, the performance gain is roughly equal to that of the best single modality.

The differences between the merged and split literal configurations indicate that, despite our best efforts, information from the node features has leaked into the structure. In the split configuration, the baseline performance is, as expected, near random with an accuracy equalling that of a majority class classifier (0.495). However, in the merged configuration the performance is roughly one-tenth higher than expected (0.616), indicating that some literals have an indegree greater than one. Judging from the differences between modalities, these literals most likely express temporal or visual information, which drop with roughly the same amount when moving from merged to split configuration.

6.1.2. Link Prediction Results

Table 5 reports the mean MRR and hits@ k over 5 runs on SYNTH, together with its 95% confidence interval and corresponding p-values. We use the same *value_merged* [*value_split*] notation as before to express the performances in the merged and split configurations, respectively.

Overall, the results show that, for most modalities, including their information considerably improves the performance when compared to the baseline (structure only). In all cases, these differences are statistically significant. When information from all modalities is included, the performance also increases noticeably, irrespective of literal configuration, from 0.045 [0.038] to 0.069 [0.057]. However, rather than performing roughly the same as the best performing single

modality (0.084 [0.068] for numerical information), including all modalities yields a performance that is slightly lower. This contrasts with our classification results.

Similar to the classification results there is considerable variation between the performances per modality: including just numerical information yields a large boost in performance, both for the merged and split literal configuration, whereas including textual or spatial information results in a drop in performance to an MRR of 0.030 [0.035] and 0.034 [0.031], respectively. Also similar is the limited influence of including visual information, although a slight but significant gain to an MRR of 0.050 is still visible in the merged literal configuration.

A final observation is that there exists a difference in performance on the baseline of 0.007 between the split and merged literal configurations, supporting our previous supposition that some information from the literals is encoded in the graph’s structure. As before, this effect seems most evident with temporal and visual information, both of which drop considerably in performance from 0.073 to 0.048 and from 0.050 to 0.028, respectively, when changing from merged to split literals.

6.1.3. Discussion

Our results indicate that, in the most ideal setting, including node features in the learning process improves the performance most or all of the times, depending on the task. This is most clear for node classification, which obtains a significance performance boost irrespective of the modality we include. With link prediction the results are less clear cut, although most modalities seem to have a positive effect on the overall performance. However, since a perfect score is practically unobtainable in this setting, it is difficult to gauge how much these effects actually matter or whether we can achieve the same by simply running the baseline for a higher number of epoch. Similarly, the drop in performance for some modalities might just as well be caused by the increased difficulty of the learning task. Some support for this supposition might be found with the drop in performance when either textual or spatial information is included, both of which require a relatively large number of parameters but still result in a near perfect score in node classification. Another possible reason is that this dataset, which is optimized for classification, lacks properties that make it an ideal testbed for link prediction.

Table 4: Entity classification results for SYNTH in accuracy, averaged over 10 runs and with 95% confidence interval, for both merged and split literals configuration. *Structure* uses only the relational information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. All p-values are in relation to using only relational information.

	<i>merged literals</i>		<i>split literals</i>	
	accuracy	p-value	accuracy	p-value
Majority Class	0.503	-	0.503	-
Structure	0.616 (± 0.003)	-	0.497 (± 0.005)	-
Structure + Features	0.996 (± 0.000)	2.33×10^{-20}	0.995 (± 0.000)	5.09×10^{-33}
Structure + Numerical	0.744 (± 0.011)	4.12×10^{-03}	0.785 (± 0.012)	3.01×10^{-29}
Structure + Temporal	0.763 (± 0.019)	1.80×10^{-14}	0.625 (± 0.012)	1.78×10^{-14}
Structure + Textual	0.995 (± 0.000)	2.39×10^{-19}	0.996 (± 0.000)	3.97×10^{-34}
Structure + Visual	0.642 (± 0.063)	5.21×10^{-04}	0.556 (± 0.044)	3.58×10^{-53}
Structure + Spatial	0.957 (± 0.002)	2.33×10^{-20}	0.949 (± 0.001)	1.22×10^{-30}

Table 5: Link prediction results for SYNTH, averaged over 5 runs and with 95% confidence interval, for both merged and split literals configuration. Listed are mean reciprocal rank (MRR) and hits@k with $k \in \{1, 3, 10\}$. *Structure* uses only the relational information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. All p-values are in relation to using only relational information.

	<i>merged literals</i>				<i>split literals</i>					
	MRR	H@1	H@3	H@10	p-value	MRR	H@1	H@3	H@10	p-value
Structure	0.045 (± 0.001)	0.041	0.048	0.050	-	0.038 (± 0.000)	0.032	0.045	0.046	-
Structure + Features	0.069 (± 0.009)	0.065	0.072	0.074	2.50×10^{-05}	0.057 (± 0.003)	0.053	0.060	0.063	2.50×10^{-05}
Structure + Numerical	0.084 (± 0.001)	0.081	0.085	0.088	2.50×10^{-05}	0.068 (± 0.000)	0.064	0.071	0.075	2.50×10^{-05}
Structure + Temporal	0.073 (± 0.001)	0.070	0.074	0.078	2.50×10^{-05}	0.048 (± 0.001)	0.043	0.056	0.060	2.50×10^{-05}
Structure + Textual	0.030 (± 0.003)	0.023	0.036	0.040	2.50×10^{-05}	0.035 (± 0.000)	0.024	0.044	0.045	2.50×10^{-05}
Structure + Visual	0.050 (± 0.002)	0.044	0.053	0.063	2.50×10^{-05}	0.028 (± 0.002)	0.026	0.029	0.034	2.50×10^{-05}
Structure + Spatial	0.034 (± 0.001)	0.028	0.038	0.041	2.50×10^{-05}	0.031 (± 0.000)	0.022	0.040	0.041	2.50×10^{-05}

Despite the aforementioned differences between tasks, we would expect to see that each modalities affects the performance roughly similar, especially with classification since literals from each modality carry a strong positive signal. As our classification results show that this is not the case, any difference in performance in this task must have originated in the MR-GCN and/or the dataset. For numerical and temporal information the precise cause is unclear and more elaborate testing is needed to determine whether the less-than-perfect performance stems from our encoders, or their implementation, or whether the fault lies with an imperfect data generation process. In contrast, since we use the proven MobileNet architecture for our visual encoder, it is likely that our image generation process is to blame for the lackluster performance when visual information is included.

When all modalities are included in the learning process, the overall performance approaches or equals that of the best performing single modality. This suggests that the message-passing network largely succeeds in learning, by itself, which information to include and which to ignore. This effect is again more profound in our classification results, for which including all modalities yield near perfect accuracy, but is still visible in the link prediction setting. As before, this difference between tasks may stem from the focus of the dataset on classification, resulting in less clear signals when used for link prediction.

6.2. Evaluation on Real-World Knowledge

Whereas previously we evaluated the MR-GCN on synthetic knowledge, we here evaluate our implementation on real-world knowledge graphs from various domains and with different (combinations of) modalities.

6.2.1. Node Classification

We evaluate the MR-GCN on five real-world knowledge graphs on node classification. General and modality-specific statistics about each of these are listed in Table 6 and 7, respectively. A short description of each dataset is given next.

AIFB+ The AIFB dataset is a benchmark knowledge graph about scientific publications from a research group, and about the people working there [1]. This is the smallest of the datasets in our experiments, and lacks the datatype annotations needed to accurately determine the literals'

modalities. These annotations were added by us, creating AIFB+.

MUTAG MUTAG is a benchmark dataset about molecules, the atoms they consist of, and any mutagenic properties that they might have [1]. This dataset only contains a single additional modality, encoded by numerical literals.

BGS The BGS dataset contains information about geological measurements in Great Britain, and includes rock composition and age [1]. Also present is spatial information, in the form of point locations and polygons.

AM+ The Amsterdam Museum dataset (AM) is a benchmark knowledge graph which contains information about the collection of a museum in The Netherlands [1]. We use the AM+ version from [28] in our experiments, which has been extended with datatype annotations and images, and which has a much higher number of labeled samples.

DMG The Dutch Monument Graph (DMG) is a benchmark dataset for multimodal entity classification [28]. The DMG includes information from all five modalities listed in Section 4.1 (in addition to relational information), with a strong emphasis on spatial information. The example given in Figure 1 is from this dataset.

Results Table 8 and 9 list the results of our classification experiments for merged and split literal configurations, respectively, and report the mean classification accuracy over 10 runs on the test sets, together with its 95% confidence interval. Corresponding p-values are available in Appendix A. We once again use the *value_merged* [*value_split*] notation to express the performances in the merged and split configurations, respectively.

Overall, our classification results show that the effects of including node features in the learning process are considerable, influencing the performance both positively and negatively, and that these effects vary greatly between datasets and modalities: including temporal information, for example, has a (slight) positive effect on the performance on AIFB+, from an accuracy of 0.933 [0.883] to that of 0.939 [0.894], but including the same form of information with DMG results in a noticeably performance drop from 0.717 [0.450] to 0.695 [0.400]. Similar effects are observable for other modalities. Moreover, including all modalities does not necessarily result in a higher accuracy, irrespective of dataset and literal configuration: only on

Table 6

Datasets used in our experiments. The AIFB+, MUTAG, and SYNTH datasets are used in both classification and link prediction, DMG, AM+, and BGS only for classification, and ML100k+ and YAGO3-10+ only for link prediction. Literals with the same value are counted as the same node in the merged count, whereas they are counted separately in the split count.

Dataset		AIFB+	MUTAG	YAGO3-10+	SYNTH	ML100k+	DMG	BGS	AM+
Relations		46	24	44	42	13	60	104	33
Entities		2,835	22,540	50,639	16,386	56,204	148,127	103,055	1,026,150
Literals	merged	5,468	1,104	20,797	112,319	32,055	195,468	230,790	127,520
	split	8,705	11,185	32,448	132,790	115,495	488,745	386,254	799,660
Facts	total	29,219	74,567	167,848	181,942	227,399	777,124	916,345	2,521,035
	train	21,175	54,547	127,802	141,899	187,393	-	-	-
	test	4,022	10,010	20,023	20,023	20,003	-	-	-
	valid	4,022	10,010	20,023	20,023	20,003	-	-	-
Classes		4	2	-	2	-	5	2	8
Labeled	total	176	340	-	16,384	-	8,399	146	73,423
	train	112	218	-	10,484	-	5,394	94	33,423
	test	36	68	-	3,278	-	2,001	29	20,000
	valid	28	54	-	2,622	-	1,001	23	20,000

Table 7

Distribution of datatypes in the datasets. Numerical information includes all subsets of real numbers, as well as booleans, whereas date, years, and other similar types are listed under temporal information. Textual information includes strings and its subsets, as well as raw URIs (e.g. links). Images and geometries are listed under visual and spatial information, respectively.

Dataset	AIFB+	MUTAG	YAGO3-10+	SYNTH	ML100k+	DMG	BGS	AM+
Numerical	115	11,185	-	29,565	55,058	17,205	12,332	160,959
Temporal	1,227	-	12,447	44,207	55,661	1,800	13	202,304
Textual	7,363	-	10,001	29,540	3,200	398,938	279,940	376,150
Visual	-	-	10,000	14,758	1,576	46,108	-	58,855
Spatial	-	-	-	14,720	-	20,866	73,870	-
Other	-	-	-	-	-	-	20,098	-

AM+, do we observe an increase when learning on all modalities, from an accuracy of 0.751 [0.578] to that of 0.760 [0.598].

Looking at the differences in baseline performance between the merged and split configurations, it is evident that all datasets express some information from the literals in their structure. This is particularly clear in the case of DMG, which drops considerably in performance from 0.717 to 0.450 when we keep literals with the same values as separate nodes. However, this effect does enable us to observe that including textual and spatial information significantly improves the accuracy on DMG to 0.518 and 0.511, respectively. Similar on AM+ for textual information, which improves the performance in the split literal configuration from 0.578 to 0.606. In both cases, the added value is masked when part of this information is encoded in

the structure. In contrast, the baseline performance on BGS stays roughly the same (0.845 [0.849]), suggesting that only few literals share a value.

Finally, our tests indicate that only the results on DMG and AM+ are statistically significant. This is most likely the result of the large number of labeled samples in the test sets of these datasets. Note that the difference of 0.001 on DMG between the performance of the baseline and that of including all features in the split literal configuration is still statistically significant because the Stuart-Maxwell test compares individual predictions rather than accuracies.

6.2.2. Link Prediction

We evaluate the MR-GCN for link prediction on four multimodal real-world datasets. Two of these—AIFB+ and MUTAG— were also used in our node classification experiments, whereas the remaining two are

Table 8

Entity classification results in accuracy, averaged over 10 runs and with 95% confidence interval, with merged literal configuration. *Structure* uses only the relation information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. Corresponding p-values are reported in Table 12. Statistically significant results are annotated with †.

Dataset	AIFB+	MUTAG	DMG	BGS	AM+
Majority Class	0.415	0.621	0.478	0.637	0.300
Structure	0.933 (± 0.013)	0.689 (± 0.024)	0.717 (± 0.001)	0.845 (± 0.010)	0.751 (± 0.004)
Structure + Features	0.908 (± 0.011)	0.658 (± 0.001)	0.475 (± 0.028) [†]	0.748 (± 0.054)	0.760 (± 0.013) [†]
Structure + Numerical	0.939 (± 0.011)	0.664 (± 0.015)	0.678 (± 0.006) [†]	0.828 (± 0.000)	0.756 (± 0.006) [†]
Structure + Temporal	0.947 (± 0.001)	-	0.695 (± 0.001) [†]	0.845 (± 0.010)	0.765 (± 0.004) [†]
Structure + Textual	0.903 (± 0.001)	-	0.538 (± 0.012) [†]	0.853 (± 0.010)	0.713 (± 0.013) [†]
Structure + Visual	-	-	0.466 (± 0.028) [†]	-	0.764 (± 0.011) [†]
Structure + Spatial	-	-	0.741 (± 0.003) [†]	0.807 (± 0.045)	-

Table 9

Entity classification results in accuracy, averaged over 10 runs and with 95% confidence interval, with split literal configuration. *Structure* uses only the relation information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. Corresponding p-values are reported in Table 13. Statistically significant results are annotated with †.

Dataset	AIFB+	MUTAG	DMG	BGS	AM+
Majority Class	0.415	0.621	0.478	0.637	0.300
Structure	0.883 (± 0.017)	0.662 (± 0.000)	0.450 (± 0.004)	0.849 (± 0.010)	0.578 (± 0.004)
Structure + Features	0.865 (± 0.001)	0.653 (± 0.012)	0.451 (± 0.021) [†]	0.829 (± 0.019)	0.598 (± 0.018) [†]
Structure + Numerical	0.869 (± 0.011)	0.655 (± 0.004)	0.369 (± 0.011) [†]	0.827 (± 0.008)	0.560 (± 0.004) [†]
Structure + Temporal	0.894 (± 0.001)	-	0.400 (± 0.002) [†]	0.841 (± 0.010)	0.515 (± 0.005) [†]
Structure + Textual	0.861 (± 0.011)	-	0.518 (± 0.025) [†]	0.852 (± 0.010)	0.606 (± 0.012) [†]
Structure + Visual	-	-	0.468 (± 0.031) [†]	-	0.594 (± 0.004) [†]
Structure + Spatial	-	-	0.511 (± 0.003) [†]	0.826 (± 0.012)	-

exclusively used for link prediction. The DMG and AM+ datasets are not used here, since their relative large number of facts would translate to exorbitant long training durations. We also abstain from testing the MR-GCN on standard link prediction benchmark datasets, such as FB15k-237 and WN18RR, as these lack node features.

General and modality-specific statistics about each of the datasets are listed in Table 6 and 7, respectively. All training, testing, and validation splits are stratified on predicate. A short description of two datasets that are exclusively used for link prediction is given next. Because of the added complexity accompanying link prediction, both datasets were subsampled to still allow for GPU acceleration.

ML100k+ MovieLens-100k is a well-known benchmark dataset about users, movies, and ratings given to these movies by the users, and contains various information that includes, amongst others, the genders and ages of users, and the release

dates and titles of movies [35]. We use a subset of the version introduced in [19], which extends the original dataset with movie posters. This subset was generated by selecting the 500 users with the highest rating count, together with all information to which they are linked.

YAGO-10+ A popular link prediction benchmark dataset is the YAGO knowledge graph. Emphasizing general knowledge, the dataset contains various information about people, cities, countries, movies, and organizations [36]. Similar as with ML100k+, we use a subset of the version introduced in [19], which enriches the original graph with images, texts, and dates. The subset was generated by taking the intersection of all entities with images, texts, and dates, together with all information to which they are linked.

Results Table 10 and 11 reports the mean MRR and its 95% confidence interval over 5 runs on the tests sets. Corresponding p-values and hits@k statis-

tics are available in Appendix A. As before, we use the *value_merged* [*value_split*] notation to express the performances in the merged and split configurations, respectively.

Overall, our results indicate that, for link prediction on real-world knowledge, including node features can have a profound effect on the performance, and that this effect can be both positive and negative. For MUTAG, this effect results in a considerable performance boost from an MRR of 0.162 [0.135] to that of 0.225 [0.202], whereas, for the three remaining datasets, this effect results in a moderate drop in performance (e.g. AIFB+, from 0.252 [0.215] to 0.215 [0.161]) to a considerable drop (e.g. YAGO3-10+, from 0.053 [0.050] to 0.025 [0.021]). These results are statistically significant for all datasets and configurations, except for AIFB+ which, when numerical information is included, achieves roughly the same performance as the baseline. A quick glance at Table 7 shows that AIFB+ only contains few numerical literals, suggesting that this result is a poor indicator of the effect that including numerical information has on the overall performance and can best be ignored.

Similar to our classification results, there appears to exist no discernible pattern in the performances amongst modalities. Instead, here too, the results for individual modalities vary much between datasets. For MUTAG, for example, adding numerical information results in a moderate performance boost from 0.162 [0.135] to 0.192 [0.140], whereas, for ML100k+, including this form of information results in a decrease in performance from 0.124 [0.028] to 0.042 [0.004]. Also similar is that, when including information from all modalities, the overall performance seems to roughly equal the average performance of all separate modalities combined.

The differences in baseline performance between the merged and split configurations shows that all datasets have some information from the literals encoded in their structure. This is most evident for ML100k+, which drops from 0.124 to 0.028 when this information is lost. In contrast, the drop in performance on YAGO3-10+ is only minor (± 0.003), indicating that only few literals have an indegree greater than one. Irrespective, for all datasets and configuration, the performance in the split configuration is the same or worse than that in the merged setting.

6.2.3. Discussion

Our results on real-world knowledge show that, overall, the effects of including node features in the learn-

ing process vary widely: for some datasets, including information from a certain modality results in a slight to considerable performance boost, whereas for other datasets that same modality does little or even results in a performance drop. This suggests that the potential gain of including node features strongly depends on the characteristics of the data and on the strength of the signals provided by the modalities. Moreover, when all modalities are included, our results show that the overall performance stays behind that of the best performing single modality. This could suggest that the message-passing model has difficulties ignoring the negative signals, or that the positive signals lack sufficient strength in many real-world datasets for the message-passing model to overcome this.

Comparing the results on AIFB+ and MUTAG from our node classification and link prediction experiments shows that the effect of including a modality on the performance differs between tasks. On AIFB+, for example, incorporating temporal information results in a slight performance gain in the classification setting, whereas the opposite is true in the link prediction setting. Similar on MUTAG for numerical information, which provides a considerable gain or drop in performance depending on which problem we are trying to solve. These results suggest that the influence of certain modalities on one task does not necessarily carry over to other tasks. A similar observation was made for our results on artificial knowledge. However, since, here, none of the classification results on either dataset is statistically significant, it remain unclear whether the differences between tasks really matter, or whether they stem from instabilities caused by the small test sets.

7. Discussion

Our results show that including node features from various modalities can have a profound effect on the overall performance of our models. However, the direction and magnitude of this effect differs depending on which dataset we use, what modalities we include, and even which tasks we perform.

When learning on on artificial knowledge, our results indicate that including multimodal information can significantly improve performance, and that the underlying message-passing model is capable of learning, by itself, which features to including and which to ignore. This contrasts with our results on real-world knowledge, which show that including node features

Table 10

Mean reciprocal rank (filtered), averaged over 5 runs and with 95% confidence interval, with merged literal configuration. *Structure* uses only the relation information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. Corresponding hits@ k and p-values are reported in Appendix A. Statistically significant results are annotated with †.

Dataset	AIFB+	MUTAG	YAGO3-10+	ML100k+
Structure	0.252 (± 0.006)	0.162 (± 0.008)	0.053 (± 0.002)	0.124 (± 0.014)
Structure + Features	0.215 (± 0.004)†	0.225 (± 0.006)†	0.025 (± 0.001)†	0.066 (± 0.010)†
Structure + Numerical	0.254 (± 0.004)	0.192 (± 0.006)†	-	0.042 (± 0.006)†
Structure + Temporal	0.237 (± 0.004)†	-	0.042 (± 0.001)†	0.111 (± 0.012)†
Structure + Textual	0.213 (± 0.005)†	-	0.021 (± 0.002)†	0.125 (± 0.010)†
Structure + Visual	-	-	0.024 (± 0.001)†	0.101 (± 0.014)†
Structure + Spatial	-	-	-	-

Table 11

Mean reciprocal rank (filtered), averaged over 5 runs and with 95% confidence interval, with split literal configuration. *Structure* uses only the relation information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. Corresponding hits@ k and p-values are reported in Appendix A. Statistically significant results are annotated with †.

Dataset	AIFB+	MUTAG	YAGO3-10+	ML100k+
Structure	0.215 (± 0.004)	0.135 (± 0.009)	0.050 (± 0.001)	0.028 (± 0.008)
Structure + Features	0.161 (± 0.003)†	0.202 (± 0.009)†	0.021 (± 0.001)†	0.003 (± 0.001)†
Structure + Numerical	0.214 (± 0.006)	0.140 (± 0.007)†	-	0.004 (± 0.001)†
Structure + Temporal	0.205 (± 0.003)†	-	0.043 (± 0.002)†	0.002 (± 0.000)†
Structure + Textual	0.154 (± 0.006)†	-	0.022 (± 0.003)†	0.019 (± 0.001)†
Structure + Visual	-	-	0.022 (± 0.002)†	0.003 (± 0.001)†
Structure + Spatial	-	-	-	-

can have very different effects depending on which dataset we use and what modalities we include. Moreover, the same message-passing model seemed unable to overcome the negative influence of some of the modalities, sometimes even resulting in an overall worse performance with node features than without. This difference between artificial and real-world knowledge might have been caused by our decision to abstain from hyperparameter optimization. However, since the same hyperparameters were effective on artificial knowledge, this is unlikely to produce such a large difference. Similar for our choices of (neural) encoders, which were unchanged between experiments. Instead, it is more likely that our chosen message-passing model has difficulties coping with negative signals and/or noise. This would explain why weak, but still positive, signals such as the visual information in SYNTH pose no problem, whereas the negative signals in some of the real-world datasets drag the overall performance down considerably.

A comparison of results between the merged and split literal configurations shows that the potential performance gain from including node features is influ-

enced by how much information from these features is already encoded in the structure of a graph. In some cases, our results show that including the same information can have little effect in the merged setting while providing a considerable performance boost in the split configuration. This suggests that much of this information is already stored as relational information, and that we gain little by also feeding the raw values to our model. This is not necessarily a problem if, by nevertheless including this information, the performance does not decrease either. However, our results show that, for some datasets and modalities, including node features results in a drop in performance. This might be caused by the added complexity that makes the problem more difficult to solve. Reducing the number of model parameters might be a first step to alleviate this problem (See also Section 8.1).

Finally, we observed that only half the datasets used in our classification experiments—SYNTH, AM+, and DMG—produced statistically significant results. The datasets in question have a considerably higher number of labeled instances, allowing for a more precise evaluation of the results. To accurately establish

which model architectures performs well in this setting we need more datasets with similarly sized test sets. However, the observed difference in statistical significance between datasets with few and many labeled instances does suggest that the Stuart-Maxwell test is suitable to compare classification results with. Similarly, in our link prediction experiments, we observed only a single result that lacked statistical significance. A quick inspection suggested that this was justified, since the dataset—AIFB+—contained only few features of the modality being tested. This suggests that the randomised paired t-test is suitable to validate link prediction results with. Since most literature in this field forgoes with statistical testing, we hope that these results encourage others to use these or similar tests for machine learning experiments on knowledge graphs.

8. Conclusion

In this work, we have proposed an end-to-end multimodal message passing model for multimodal knowledge graphs. By embedding our model in the message passing framework, and by treating literals as first-class citizen, we embrace the idea that this enables data scientists to learn end-to-end from any heterogeneous multimodal knowledge, as long as it is represented as a knowledge graph. To test our hypothesis, we have implemented our model and evaluated its performance for both node classification and link prediction on a large number of artificial and real-world knowledge graphs from various domains and with different degrees of multimodality.

Our results indicate that, overall, including information from other modalities can have a considerable effect on the performance of our models, but that the direction and magnitude of this effect strongly depends on the characteristics of the knowledge. In the most ideal situation, when the dataset contains little noise and strong positive signals, incorporating node features has the potential to significantly improve performance. When faced with real-world knowledge, however, our results show that this effect can vary considerably between datasets, modalities, and even tasks.

Despite the mixed results on real-world knowledge, we believe that this work supports our hypothesis that by enabling our models to naturally ingest literal values, and by treating these values according to their modalities, tailoring their encodings to their specific characteristics, we stay much closer to the original and complete knowledge that is available to us, potentially

resulting in an increase in the overall performance of our models.

Learning end-to-end on heterogeneous knowledge has a lot of promise which we have only scratched the surface of. A model that learns in a purely data-driven way to use information from different modalities, and to integrate such information along known relations, has the potential to allow practitioners a much greater degree of hands-free machine learning on multimodal heterogeneous knowledge.

8.1. Limitations and future work

Our aim has currently been to demonstrate that we can train a multimodal message passing model end-to-end which can exploit the information contained in a graph’s literals and naturally combine this with its relational counterpart, rather than to established that our implementation reaches state-of-the-art performance, or even to measure its performance relative to other published models. We therefore performed little hyperparameter tuning in our experiments, ensuring that any observable difference in performance could be confidently attributed to the inclusion or exclusion of information from a certain modality, rather than have been caused by a particular hyperparameter setting.

To properly establish which type of model architecture performs best in multimodal settings, and whether message passing models provide an advantage over more shallow embedding models without message passing, we require more extensive, high-quality, standard benchmark datasets with well-defined semantics (i.e. datatype and/or relation range declarations) and a large number of labeled instances. Recently, some datasets have seen the light which are suitable for this purpose (e.g. [28]). However, to perform more precise evaluations and more accurate models comparisons, we need even more datasets from a wide range of domains and with a large number of different modalities. Nevertheless, to determine precisely what kind of knowledge is most fitting for this form of learning we are likely to require an iterative process where each generation of models provides inspiration for the next generation of benchmark datasets and vice versa.

In other work, currently under submission, we explore techniques to reduce the overall complexity of a multimodal model by reducing the number of parameters by merging some of the weight matrices. Our main motivation for this is the necessity of full batch learning with many message passing networks—a known limitation—which makes it challenging to learn from

large graphs; a problem which becomes even more evident as we start adding multimodal node features. Future work will also investigate the other side of the spectrum by using a separate set of learnable weights per relation, as opposed to sharing weights amongst literals of the same modality. While this adds some additional complexity, it allows a more natural encoding of a graph in our model by capturing the semantics per relation. To illustrate this, compare learning a single set of weights for age and height, both of which are numeric, against learning a separate set of weights for each.

Lastly, a promising direction of research is the use of pretrained encoders. In our experiments, we show that the encoders receive enough of a signal from the downstream network to learn a useful embedding, but this signal is complicated by the message passing head of the network, and the limited amount of data. Using a modality-specific, pretrained encoder, such as GPT-2 for language data [37] or Inception-v4 for image data [38], may provide us with good general-purpose feature at the start of training, which can then be fine-tuned to the specifics of the domain.

Acknowledgments

We express our gratitude to Lucas van Berkel for his insightful comments. This project is supported by the NWO Startimpuls programme (VWData - 400.17.605) and by the Amsterdam Academic Alliance Data Science (AAA-DS) Program Award to the UvA and VU Universities.

Appendix A. Detailed Results

The following tables list more detailed results from our experiments. Tables 12 and 13 list the statistical significance for our classification results for merged and split literal configurations, respectively. For our link prediction experiments, tables 14, 15, 16, and 17 list the hits@ k and the statistical significance for AIFB+, MUTAG, YAGO3-10+, and ML100k+ respectively.

References

- [1] P. Ristoski, G.K.D. De Vries and H. Paulheim, A collection of benchmark datasets for systematic evaluations of machine learning on the semantic web, in: *International Semantic Web Conference*, Springer, 2016, pp. 186–194.
- [2] N. Shervashidze, P. Schweitzer, E.J.v. Leeuwen, K. Mehlhorn and K.M. Borgwardt, Weisfeiler-lehman graph kernels, *Journal of Machine Learning Research* **12**(Sep) (2011), 2539–2561.
- [3] X. Wilcke, P. Bloem and V. de Boer, The Knowledge Graph as the Default Data Model for Learning on Heterogeneous Knowledge, *Data Science* **1**(1–2) (2017), 39–57, DOI:10.3233/DS-170007.
- [4] T. Baltrušaitis, C. Ahuja and L.-P. Morency, Multimodal machine learning: A survey and taxonomy, *IEEE transactions on pattern analysis and machine intelligence* **41**(2) (2018), 423–443.
- [5] H. Cai, V.W. Zheng and K.C.-C. Chang, A Comprehensive Survey of Graph Embedding: Problems, Techniques and Applications, *arXiv:1709.07604 [cs]* (2018), arXiv: 1709.07604. <http://arxiv.org/abs/1709.07604>.
- [6] Q. Wang, Z. Mao, B. Wang and L. Guo, Knowledge graph embedding: A survey of approaches and applications, *IEEE Transactions on Knowledge and Data Engineering* **29**(12) (2017), 2724–2743.
- [7] G.A. Gesese, R. Biswas, M. Alam and H. Sack, A Survey on Knowledge Graph Embeddings with Literals: Which model links better Literal-ly?, *arXiv preprint arXiv:1910.12507* (2019), Under review at SWJ.
- [8] M. Nickel, V. Tresp and H.-P. Kriegel, Factorizing yago: scalable machine learning for linked data, in: *Proceedings of the 21st international conference on World Wide Web*, 2012, pp. 271–280.
- [9] M. Nickel, V. Tresp and H.-P. Kriegel, A three-way model for collective learning on multi-relational data., in: *Icml*, Vol. 11, 2011, pp. 809–816.
- [10] M. Cochez, M. Garofalo, J. Lenßen and M.A. Pellegrino, A first experiment on including text literals in KGloVe, *arXiv preprint arXiv:1807.11761* (2018).
- [11] A. Kristiadi, M.A. Khan, D. Lukovnikov, J. Lehmann and A. Fischer, Incorporating literals into knowledge graph embeddings, *arXiv preprint arXiv:1802.00934* (2018).
- [12] G. de Assis Costa and J.M.P. de Oliveira, Towards Exploring Literals to Enrich Data Linking in Knowledge Graphs, in: *2018 IEEE First International Conference on Artificial Intelligence and Knowledge Engineering (AIKE)*, IEEE, 2018, pp. 110–114.
- [13] Y. Wu and Z. Wang, Knowledge graph embedding with numeric attributes of entities, in: *Proceedings of The Third Workshop on Representation Learning for NLP*, 2018, pp. 132–136.
- [14] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston and O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: *Advances in neural information processing systems*, 2013, pp. 2787–2795.
- [15] R. Xie, Z. Liu, J. Jia, H. Luan and M. Sun, Representation learning of knowledge graphs with entity descriptions, in: *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [16] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado and J. Dean, Distributed representations of words and phrases and their compositionality, in: *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [17] R. Xie, Z. Liu, H. Luan and M. Sun, Image-embodied knowledge representation learning, *arXiv preprint arXiv:1609.07028* (2016).

Table 12: Statistical significance for entity classification results, averaged over 10 runs, for merged literal configuration. *Structure* uses only the relational information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. All p-values are in relation to using only relational information.

	AIFB+	MUTAG	DMG	BGS	AM+
Structure + Features	5.72×10^{-01}	8.33×10^{-02}	1.96×10^{-105}	3.17×10^{-01}	5.12×10^{-22}
Structure + Numerical	1.00×10^{-00}	8.57×10^{-02}	6.62×10^{-29}	1.00×10^{-00}	9.52×10^{-03}
Structure + Temporal	1.00×10^{-00}	-	3.21×10^{-28}	1.00×10^{-00}	6.25×10^{-112}
Structure + Textual	5.72×10^{-01}	-	1.99×10^{-108}	3.17×10^{-01}	6.58×10^{-87}
Structure + Visual	-	-	1.44×10^{-49}	-	7.91×10^{-104}
Structure + Spatial	-	-	1.26×10^{-02}	1.00×10^{-00}	-

Table 13: Statistical significance for entity classification results, averaged over 10 runs, for split literal configuration. *Structure* uses only the relational information whereas *Structure + Features* also includes information from all supported modalities. The rest provides a breakdown per modality. All p-values are in relation to using only relational information.

	AIFB+	MUTAG	DMG	BGS	AM+
Structure + Features	6.40×10^{-01}	1.00×10^{-00}	2.26×10^{-159}	1.56×10^{-01}	8.26×10^{-54}
Structure + Numerical	1.00×10^{-00}	1.00×10^{-00}	6.14×10^{-87}	3.17×10^{-01}	6.72×10^{-48}
Structure + Temporal	8.01×10^{-01}	-	3.82×10^{-36}	3.17×10^{-01}	9.60×10^{-135}
Structure + Textual	6.16×10^{-01}	-	1.07×10^{-144}	1.00×10^{-00}	5.12×10^{-58}
Structure + Visual	-	-	4.31×10^{-89}	-	1.40×10^{-44}
Structure + Spatial	-	-	6.01×10^{-18}	8.33×10^{-02}	-

- [18] Y. Tay, L.A. Tuan, M.C. Phan and S.C. Hui, Multi-task neural network for non-discrete attribute prediction in knowledge graphs, in: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 2017, pp. 1029–1038.
- [19] P. Pezeshkpour, L. Chen and S. Singh, Embedding Multimodal Relational Data for Knowledge Base Completion, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018, pp. 3208–3218.
- [20] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li and L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 2009, pp. 248–255.
- [21] O. Lassila, R.R. Swick et al., Resource description framework (RDF) model and syntax specification (1998).
- [22] J. Gilmer, S.S. Schoenholz, P.F. Riley, O. Vinyals and G.E. Dahl, Neural message passing for quantum chemistry, in: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, JMLR.org, 2017, pp. 1263–1272.
- [23] T.N. Kipf and M. Welling, Semi-supervised classification with graph convolutional networks, *arXiv preprint arXiv:1609.02907* (2016).
- [24] M. Schlichtkrull, T.N. Kipf, P. Bloem, R. Van Den Berg, I. Titov and M. Welling, Modeling relational data with graph convolutional networks, in: *European Semantic Web Conference*, Springer, 2018, pp. 593–607.
- [25] B. Yang, W.-t. Yih, X. He, J. Gao and L. Deng, Embedding entities and relations for learning and inference in knowledge bases, *arXiv preprint arXiv:1412.6575* (2014).
- [26] D. Ruffinelli, S. Broscheit and R. Gemulla, You CAN teach an old dog new tricks! on training knowledge graph embeddings, in: *International Conference on Learning Representations*, 2019.
- [27] X. Zhang, J. Zhao and Y. LeCun, Character-level convolutional networks for text classification, in: *Advances in neural information processing systems*, 2015, pp. 649–657.
- [28] P. Bloem, L. van Berkel, W. Wilcke and V. de Boer, kgbench: A Collection of Datasets for Multimodal and Relational Learning on Heterogeneous Knowledge, in: *Proceedings of the European Semantic Web Conference 2021*, 2021, Under Submission.
- [29] R. van't Veer, P. Bloem and E. Folmer, Deep Learning for Classification Tasks on Geospatial Vector Polygons, *arXiv preprint arXiv:1806.03857* (2018).
- [30] A.G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto and H. Adam, Mobilenets: Efficient convolutional neural networks for mobile vision applications, *arXiv preprint arXiv:1704.04861* (2017).
- [31] A.E. Maxwell, Comparing the Classification of Subjects by Two Independent Judges, *The British Journal of Psychiatry* **116**(535) (1970), 651–655. doi:10.1192/bjp.116.535.651. <http://www.cambridge.org/core/journals/the-british-journal-of-psychiatry/article/comparing-the-classification-of-subjects-by-two-independent-judges/B2F3A3380DDFCBA7B98C157F93FD1A0E>.
- [32] A. Stuart, A Test for Homogeneity of the Marginal Distributions in a Two-Way Classification, *Biometrika* **42**(3/4) (1955), 412–416. doi:10.2307/2333387. <https://www.jstor.org/stable/2333387>.
- [33] P.R. Cohen, *Empirical methods for artificial intelligence*, Vol. 139, MIT press Cambridge, MA, 1995.
- [34] A. Yeh, More accurate tests for the statistical significance of result differences (2000).
- [35] F.M. Harper and J.A. Konstan, The movielens datasets: History and context, *Acm transactions on interactive intelligent systems (tiis)* **5**(4) (2015), 1–19.
- [36] F.M. Suchanek, G. Kasneci and G. Weikum, Yago: a core of semantic knowledge, in: *Proceedings of the 16th international conference on World Wide Web*, 2007, pp. 697–706.
- [37] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei and I. Sutskever, Language Models are Unsupervised Multitask Learners (2019).
- [38] C. Szegedy, S. Ioffe, V. Vanhoucke and A.A. Alemi, Inception-v4, inception-resnet and the impact of residual connections on learning, in: *Thirty-first AAAI conference on artificial intelligence*, 2017.