

Geolocation of Cultural Heritage using Multi-View Knowledge Graph Embedding

Hebatallah A. Mohamed¹, Sebastiano Vascon^{1,2}, Feliks Hibraj², Stuart James³,
Diego Pilutti¹, Alessio Del Bue³, and Marcello Pelillo^{1,2}

¹ European Center for Living Technology (ECLT), Ca' Foscari University of Venice,
Italy

² Department of Environmental Sciences, Informatics and Statistics (DAIS), Ca'
Foscari University of Venice, Italy

³ Pattern Analysis and Computer Vision (PAVIS), Istituto Italiano di Tecnologia,
Genova, Italy

Abstract. Knowledge Graphs (KGs) have proven to be a reliable way of structuring data. They can provide a rich source of contextual information about cultural heritage collections. However, cultural heritage KGs are far from being complete. They are often missing important attributes such as geographical location, especially for sculptures and mobile or indoor entities such as paintings. In this paper, we first present a framework for ingesting knowledge about tangible cultural heritage entities from various data sources and their connected multi-hop knowledge into a geolocalized KG. Secondly, we propose a multi-view learning model for estimating the relative distance between a given pair of cultural heritage entities, based on the geographical as well as the knowledge connections of the entities.

Keywords: Cultural heritage · Geolocation · Knowledge graphs · Multi-view graph embedding.

1 Introduction

The term cultural heritage includes tangible heritage, which can be further specified in *i*) movable (such as paintings, sculptures, coins); *ii*) immovable (such as monuments, archaeological sites), and intangible heritage, such as traditions and performing arts [7]. Preparing cultural heritage collections for exploration by a wide range of users with different backgrounds requires to integrate heterogeneous data into modern information systems.

Knowledge graphs (KGs) have proven to be a reliable way of structuring data in wide range of domains, including the cultural heritage domain [8, 10]. KGs, such as Wikidata [17], are large directed network of real-world entities and relationships between them, where facts are represented as triplets in the form of (*head entity*, *relation*, *tail entity*). They enable to connect knowledge about cultural heritage collections, and enrich this knowledge with external data coming from heterogeneous sources.

A key challenge for KGs is the human annotation required, resulting in them remaining incomplete. Many general purpose KGs, such as Wikidata [17], are missing latitude and longitude coordinates for the entities representing tangible cultural heritage (especially the movable ones). In this paper, we present a method for estimating the relative distance between a given pair of cultural heritage entities in a KG. In particular, we propose a joint learning model to learn comprehensive and representative entity embeddings, based on entity geographical location and semantics. Moreover, we present a framework for ingesting city cultural heritage from multiple general purpose data sources into a geolocalized KG.

In brief, our main contributions are as follows:

- We present an ingestion tool and related system framework to create a geolocalized KG for cultural heritage exploration, by ingesting data from heterogeneous data sources based on city coordinates. The tool is based on Neo4j⁴ graph-database and is published as open source on GitHub⁵.
- We introduce new KG datasets for the geolocation prediction task, with a KG containing both spatial and non-spatial entities and relations.
- We propose a method that introduces a geospatial distance restriction to refine the embedding representations of geographic entities in a geolocalized KG, which fuses geospatial information and semantic information into a low-dimensional vector space. We then utilize this method for estimating the geographical distance between cultural heritage entities, where we outperform state-of-the-art methods.

2 Related Work

In this section, we discuss work related to cultural heritage KGs, and geolocation using KGs.

2.1 KGs for Cultural Heritage

There are some works that have proposed to use KGs to support cultural heritage exploration. For example, Pellegrino et. al. have proposed [8], a general-purpose approach to perform Question-Answering on KGs for querying data concerning cultural heritage. They have assessed the system performance on domain-independent KGs such as Wikidata [17] and DBpedia [1]. The authors of [10] have proposed a semantic graph-based recommender system of cultural heritage places, by modeling a graph representing the semantic relatedness of cultural heritage places using tags extracted from the descriptive content of the places from Wikipedia⁶.

⁴ <https://neo4j.com/>

⁵ <https://github.com/MEMEXProject/MEMEX-KG>

⁶ <https://www.wikipedia.org/>

In addition, there are some existing cultural heritage specific KGs, related to particular places or historical collections. For example, ArCO [2] is a KG that provides ontology to link people, events, and places about Italian artifacts and document collections. ArCO has been created from different data sources, including general purpose KGs such as Wikidata and DBpedia. Linked Stage Graph [14] which organizes and interconnects data about the Stuttgart State Theaters. ArDO [18] is an ontology to represent the dynamics of annotations of general archival resources. In our work, we propose a system framework for ingesting knowledge about any city cultural heritage using heterogeneous data coming from Wikidata [17], and any other geolocalized linked open data (LOD) such as Europeana [3], into a geolocalized KG.

2.2 Geolocation using KGs

There have been only few works that have focused on the task of geographical KG construction. The Guo et. al. have proposed GeoKG [4], a system framework for extracting geographical knowledge from geographic datasets and represent the knowledge by means of concepts and relations with the aid of GeoSPARQL⁷. However, this work has focused on specific semantic relationships such as *subclass_of* or *instance_of* to describe the subclass concepts of the geographical entities, such as *river* and *railway station*. In our work, we aim at constructing a geolocalized KG where tangible cultural heritage entities can be represented with richer semantic relations, such as *architectural_style*, *architect* and *exhibited_at*.

Moreover, there are works that focused on geographical KG completion and location prediction. For example, the authors of [11] have proposed a translational embedding model that has the ability to capture and embed the geospatial distance restriction with the semantic information of the geographical KG into a vector space. Then, the optimized model outputs the refined representations of geo-entities and geo-relations, which improves the completion performance on the sparse geographical KG. However, this work does not consider the existence of non geo-entity and non geo-relations in the KG. In our work, we aim at predicting the distance between two entities in a KG that contains both geographical and non-geographical entities and relations.

3 Data Curation for Geolocalized Knowledge Graph

In this section, we discuss details about our developed ingestion tool, and the related system framework for ingesting data about a given city and the cultural heritage it contains into a geolocalized KG.

⁷ <http://www.opengis.net/ont/geosparql>

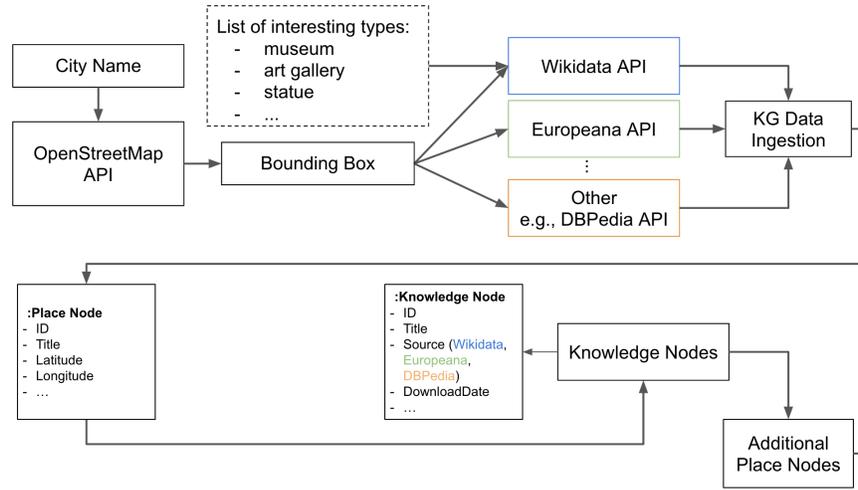


Fig. 1. Information extraction workflow.

3.1 Data Sources

In this work, Wikidata⁸ and Europeana⁹ have been selected for enriching our geolocalized KG, however other general purpose KGs and LOD can be incorporated. Wikidata is a public crowd-sourced knowledge base, with facts and figures about tens of millions of items (93,207,191). The data are offered freely, with no restriction on the reuse and modification (Creative Commons Zero). Wikidata provides its data openly in an easy to use standard Resource Description Framework (RDF). It provides much of its data in hundreds of languages. Already, large amounts of data about cultural assets are being shared with Wikidata by formal partnerships with Galleries, Libraries, Archives and Museums (GLAMS).

On the other hand, Europeana aims to facilitate the usage of digitized cultural heritage resources from and about Europe [5]. It seeks to enable users to access content in all European languages via the Europeana collections portal and allow applications to use cultural heritage metadata via its open APIs. Europeana holds metadata from over 3,700 providers [3], mostly GLAMS. However, it uses a strictly federated ontology limiting the diversity of the meta-data.

3.2 Information Extraction Workflow

We propose a localized approach that grows a KG based on association to known landmarks. We define two types of KG nodes: *Place* and *Knowledge*. The former represents tangible data while the latter represents intangible data. As shown in Figure 1, first, we use OpenStreetMap (OSM)¹⁰ to get the bounding box of the

⁸ <https://www.wikidata.org/>

⁹ <https://www.europeana.eu/>

¹⁰ <https://www.openstreetmap.org/>

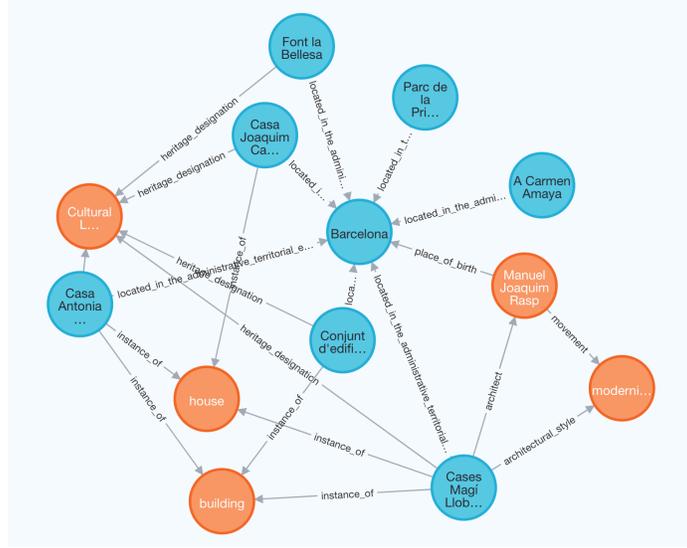


Fig. 2. A subset of our geolocated KG (orange nodes are *Knowledge* and cyan ones are *Places*).

target city. Then coordinates are used to query Wikidata to retrieve geo-entities within the selected scope, using Wikidata *coordinate_location* (*P625*) attribute. Moreover, a list of interesting entity types can be predefined by the user, to allow the ingestion tool to focus on specific cultural heritage related entities. We use the Wikidata *instance_of* (*P31*) and *subclass_of* (*P279*) attributes to select entities with specific types. These form the basis of our *Places* (node type is *Place*) where related additional properties are downloaded.

From each of the identified *Places*, we then query associated nodes based on all relations from Wikidata, this new set of nodes we refer to as *Knowledge*. This could be for example, the architect of a building. We repeat this step searching for new relationships a predefined number of times referred to as hops. We use the same approach to ingest data from Europeana. For a given city, we retrieve its coordinates using OSM, then a query is performed against Europeana API to identify content within the GPS region to add nodes in the KG. Figure 2 shows a small subset of our geolocated KG.

3.3 NER for Data Integration

We utilize Named Entity Recognition (NER) technique in order to link KG entities from different sources of information. More precisely, we use spaCy¹¹, a Python library with wide linguistic features. SpaCy annotates text with different types of named entities. We focus on the types listed in Table 1. For each entity

¹¹ <https://www.spacy.io/>

Table 1. Named entity types used to integrate data from different sources.

Type	Meaning
GPE	Geopolitical entity, i.e., countries, cities, states.
LOC	Non-GPE locations, mountain ranges, bodies of water.
FAC	Buildings, airports, highways, bridges, etc.
PER	People, including fictional ones.

ingested from Europeana, we apply NER to extract named entities from the title of Europeana entity. Then, we match those extracted named entities with the titles of Wikidata entities. If there is a match found, we create a relation (we name it *related.to*) between the matching Europeana and Wikidata entities. For example, when applying NER on a Europeana entity titled “*Carmen Amaya’s last dance*”, the named entity “*Carmen Amaya*” of type PER will be extracted, and the Europeana entity will be linked to the entity titled “*A Carmen Amaya*” from Wikidata in the KG.

4 Geolocation using Multi-view Graph Embedding

Given a geolocalized KG G , composed of a large number of triples in the form of (*head entity, relation, tail entity*). All entities have associated type $\in \{Place, Knowledge\}$, title $S = [w_1, w_2, \dots, w_n]$ with n words. Additionally, most of the *Place* entities have a latitude and longitude information. From our KG we aim to predict the distance d_t between any two *Place* entities, u and v , where the latitude and longitude property is missing in at least one of those entities (i.e., the distance between is unknown). To achieve this, we propose a method to construct two types of *Place* entity correlations: 1) based on the geographical view using the *Place* nodes, and 2) based on entity semantics using the *Knowledge* nodes. An overview about our proposed approach is shown in Figure 3, and more details are described in the following sections.

4.1 Geographical View

The geographical view aims at capturing the geographical relation between two geo-entities, *Places*, in the KG. To represent this view, we extract an induced subgraph around the two target entities. The subgraph represent how the two geo-entities are connected to each other, in terms of direct links or common nodes between the two target nodes. Since we are focusing on geographical view, we restrict the common nodes to *Place* node type. We then embed the extracted subgraph with a relational graph convolution network (R-GCN) [13] to represent the geographical relation between the target entities. In more details, we perform the following steps 1. Subgraph Extraction (sec. 4.1), and 2. Node Representation (sec. 4.1).

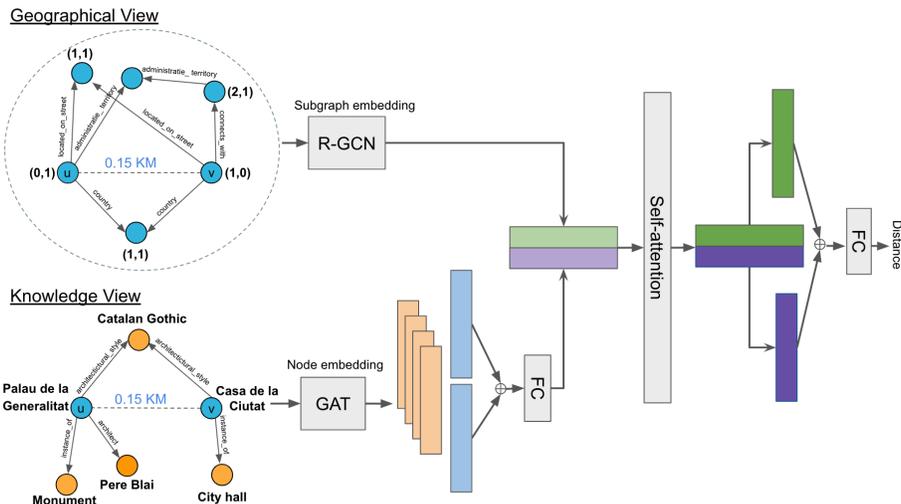


Fig. 3. The proposed approach for geographical location prediction.

Subgraph Extraction We assume that the paths (i.e., geo-entities and geo-relations) on the KG connecting the two target nodes contain the information that could imply the distance between the two target entities. Hence, as a first step, we extract an enclosing subgraph around the target nodes. The enclosing subgraph between nodes u and v is the graph induced by all the nodes that occur on a path between u and v . The subgraph set is given by the intersection of neighbors of the two target nodes followed by a pruning procedure.

More precisely, let $N_k(u)$ and $N_k(v)$ be set of nodes in the k -hop (undirected) neighborhood of the two target nodes in the KG. We compute the enclosing subgraph by taking the intersection, $N_k(u) \cap N_k(v)$, of these k -hop neighborhood sets and then prune nodes that are isolated or at a distance greater than k from either of the target nodes. This results in all the nodes that occur on a path of length at most $k + 1$ between nodes u and v , where we refer to it as the induced subgraph $G(u, v)$.

Node Representation We define an embedding for each entity (node) in the subgraph $G(u, v)$. Following [20], each node i in the subgraph is labeled with the tuple $(d(i, u), d(i, v))$, where $d(i, u)$ is the shortest distance between nodes i and u (likewise for $d(i, v)$). The two target nodes, u and v , are labeled $(0, 1)$, and $(1, 0)$ to be identifiable by the model. This scheme captures the position of each node in the subgraph with respect to the target nodes, as shown in Figure 3. The node features are defined as $[\text{one-hot}(d(i, u)) \oplus \text{one-hot}(d(i, v))]$, representing the concatenation of the one-hot embedding of the labels.

Geographical View Embedding We use multiple layers of the multi-relational R-GCN [13] to learn the embeddings of the extracted subgraph $G(u, v)$. R-GCN

adopts general message-passing scheme [19], where a node representation is iteratively updated by combining it with the aggregation of its neighbors’ representation. The subgraph representation of $G(u, v)$ is obtained by average-pooling of all the node representations: In the k -th layer of our graph neural network (GNN), a_i^k represents the aggregated message from the neighbors of node i . The aggregation function is defined as:

$$a_i^k = \sum_{r=1}^R \sum_{s \in N_r(i)} \alpha_{r(s,i)}^k W_r^k h_s^{k-1}, \quad (1)$$

where R is the total number of unique relations, $N_r(i)$ represents the neighbors of node i under relation r , W_r^k is the transformation matrix of the k -th layer over relation r , and $\alpha_{r(s,i)}^k$ is the edge attention weight at the k -th layer corresponding to the edge between nodes s and i via relation r .

The latent representation of node i in the k -th layer is:

$$h_i^k = \text{ReLU}(W_0^k h_i^{k-1} + a_i^k), \quad (2)$$

where W_0 aims at retaining the information of the node itself using self-connection, and ReLU is an activation function.

The subgraph representation of $G(u, v)$ is obtained by average-pooling of all the node representations:

$$h_{G(u,v)}^L = \frac{1}{|V|} \sum_{i \in V} h_i^L, \quad (3)$$

where V denotes the set of nodes in $G(u, v)$, and L represents the number of layers of message-passing.

4.2 Knowledge View

The knowledge view aims at representing the semantic relation between two target geo-entities, *Places*, in the KG, by means of their surrounding *Knowledge* nodes. We employ a graph attention network (GAT) [16] for representing each target entity semantics based on entity textual title. Then, we concatenate the representation of the two target entities and pass the concatenated representation into a linear layer. In more details, as with Geographical View, we perform: 1. Node Representation (sec. 4.2), and 2. Node Embedding (sec. 4.2).

Node Representation For each node in G with title $S = [w_1, w_2, \dots, w_n]$ (n is different for each node since each title might have different length), we initialize the embedding of each word w with a pre-trained word vector from GloVe [9]. Let p denote the dimension of each word vector. We obtain the sentence representation \vec{S} by aggregating the embedding of each word, where $\vec{S} \in \mathbb{R}^p$. In aggregation, we use only simple averaging due to its validity [21].

Node Embedding We employ a GAT to learn representations of nodes. GAT applies attention mechanism on graph-structured data. It updates the representation of a vertex by propagating information to its neighbors, where the weights of its neighbor nodes is learned by attention mechanism automatically. Formally, given the input vertex features $\{\vec{S}_1, \vec{S}_2, \dots, \vec{S}_n\}$, a GAT layer updates the vertex representations by following steps:

$$e_{ij} = \exp\left(\text{ReLU}\left(a^T \left[W_k \vec{S}_i \oplus W_k \vec{S}_j\right]\right)\right) \quad (4)$$

$$\alpha_{ij} = \text{softmax}_j(e_{ij}) = \frac{\exp(e_{ij})}{\sum_{k \in N_i} \exp(e_{ik})} \quad (5)$$

$$h'_i = \sigma\left(\sum_{j \in N_i} \alpha_{ij} W_k h_j\right) \quad (6)$$

where W_k and a are learnable parameters, \oplus is the concatenation operation.

Knowledge View Embedding After calculating the embedding of each node in the graph, we concatenate the embeddings of the target entities, u and v . Then, we pass these concatenated representations through a linear layer, given by

$$h_{u,v} = [h'_u \oplus h'_v] W_s, \quad (7)$$

where \oplus refers to the concatenation operation, and W_s is a learnable weight parameter.

4.3 Multi-View Embedding

In order to enable the cooperation among the different views during the learning process and effectively fuse multi-view representations, our model enables information sharing across all views via a scaled dot-product self-attention [15]. Given the representations from the different views as $\{\mathcal{E}_i\}_{i=1}^2$, we stack them into a matrix $X \in \mathbb{R}^{2 \times d}$ of dimension d . We associate query, key and value matrices Q , K and V as follows:

$$Q = XW_Q, \quad K = XW_K, \quad V = XW_V \quad (8)$$

where W_Q , W_K and W_V are weight matrices. We then propagate information among all views as follows:

$$Y = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right) V \quad (9)$$

$\hat{\mathcal{E}}_i$ will then be the i -th row in matrix Y , which considered as the relevant global information for i -th view. To incorporate this information in the learning process, we concatenate the representations of the different views (geographical and

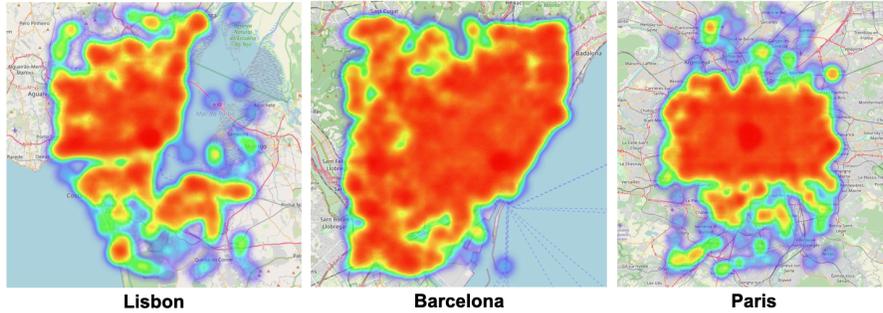


Fig. 4. Visualization of tangible cultural heritage on the map.

knowledge views) and pass it to a fully connected (FC) layer, to predict the distance \hat{d}_t as follows:

$$\hat{d}_t = [\hat{\mathcal{E}}_1 \oplus \hat{\mathcal{E}}_2]W_m, \quad (10)$$

where \oplus refers to the concatenation operation, and W_m is a learnable weight parameter.

4.4 Learning Objective

Given the ellipsoidal shape of the earth’s surface, we apply the Haversine distance [12] to calculate the distance of two points represented by their latitude in range of $\{-90, 90\}$ and longitude in range of $\{-180, 180\}$. The Haversine distance is the great circle distance between two geographical coordinate pairs. We train our model to reduce the Mean Squared Error (MSE) loss based on the actual and predicted Haversine distance, d_t and \hat{d}_t .

5 Experiments

5.1 Datasets

In order to evaluate our model, we create three datasets using our ingestion tool described in section 3. The datasets are about tangible (*Place*) and *Knowledge* of three main cities in Europe: Lisbon, Barcelona and Paris.

We set the number of hops to 3, since 3-hop contains sufficient information about each tangible cultural heritage entity. Statistics about the datasets, including the maximum distance between geo-entites representing cultural heritage related places are shown in Table 2. The maximum distance between the places in Barcelona dataset is less than in Lisbon and Paris. Lisbon dataset is the most challenging since it contains fewer number of nodes and relations. Moreover, Lisbon has the highest maximum distance between places. In Figure 4, we visualise the location of the tangible cultural heritage of the different cities.

Table 2. Statistics of the city datasets.

	Lisbon	Barcelona	Paris
# <i>Place</i> Nodes	2,780	5,989	26,214
# <i>Knowledge</i> Nodes	5,860	9,140	35,885
#Links	21,253	37,384	196,898
#Relation types	377	708	578
# <i>Place-Place</i> links	5,077	6,907	48,954
Maximum distance	35.32 KMs	16.76 KMs	34.69 KMs

We randomly split the links between *Places* into 80% and 20% sets with corresponding distances for training and testing, respectively. Hyper-parameters are optimised for the training set.

5.2 Baselines and Implementation Details

We compare our model against GAT [16] and R-GCN [13] as baselines when textual title of the *Place* and *Knowledge* nodes are used for message-passing to generate node embeddings. We also evaluate our model when using only the geographical view (Ours - geographical), utilizing enclosing subgraph embeddings with R-GCN. Moreover, we experiment our model without the attention layer (Ours - without att).

The latent embedding sizes used in all our models are set to 32. We set the number of layers in GAT and R-GCN to 3. To train our model, we use Adam optimizer [6] with a learning rate of 1e-4. We run our experiments on a machine with two Intel Xeon Gold 6230 CPUs running at 2.10 GHz with 128 GB of memory, and Nvidia Quadro RTX 5000 GPU with 16 GB of memory. Finally, we use Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) to measure the prediction errors.

5.3 Results

In Table 3, we present the results of the different prediction methods in terms of MAE and RMSE. As shown, our proposed method outperform other baseline methods. Best results are achieved on Barcelona dataset, since the maximum distance between the places is less than the one in Lisbon and Paris.

The results show that GAT is outperforming R-GCN when applying the message-passing technique on the whole KG for extracting node embeddings. This indicates that GAT is able to focus on certain neighbors to represent a *Place* node due to its attention mechanism. However, when applying R-GCN on enclosing subgraphs (Ours - geographical), the performance is outperforming the R-GCN baseline since the model is focusing on fewer unique relation types representing the spatial relation. Finally, the results show that applying an attention layer helps to fuse the multiple views in our proposed model.

Table 3. The results showing MAE and RMSE of Haversine distance in *km*. Best results are in **Bold**.

	Lisbon		Barcelona		Paris	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
GAT	1.95	2.64	1.43	1.87	0.70	1.23
R-GCN	2.30	3.12	1.48	1.90	1.21	1.63
Ours - geographical	2.23	3.10	0.48	0.75	0.83	1.42
Ours - without att.	2.13	2.86	0.56	0.86	0.67	1.20
Ours	1.90	2.50	0.42	0.71	0.59	1.05

6 Conclusion

We present an ingestion tool and a framework to create a geolocalized KG for contextualising cultural heritage. In addition, we propose a method that introduces a geospatial distance restriction to refine the embedding representations of geographic entities in a geolocalized KG, which fuses geospatial information and semantic information into a low-dimensional vector space. We utilize this method for a geographical distance prediction task, where we outperform baseline methods.

ACKNOWLEDGMENTS

This work was supported by MEMEX project funded by the European Union’s Horizon 2020 research and innovation program under grant agreement No 870743.

References

1. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: Dbpedia: A nucleus for a web of open data. In: The semantic web, pp. 722–735. Springer (2007)
2. Carriero, V.A., Gangemi, A., Mancinelli, M.L., Marinucci, L., Nuzzolese, A.G., Presutti, V., Veninata, C.: Arco: The italian cultural heritage knowledge graph. In: International Semantic Web Conference. pp. 36–52. Springer (2019)
3. Freire, N., Voorburg, R., Cornelissen, R., de Valk, S., Meijers, E., Isaac, A.: Aggregation of linked data in the cultural heritage domain: A case study in the europeana network. *Information* **10**(8), 252 (2019)
4. Guo, X., Qian, H., Wu, F., Liu, J.: A method for constructing geographical knowledge graph from multisource data. *Sustainability* **13**(19), 10602 (2021)
5. Haslhofer, B., Isaac, A.: data. europeana. eu: The europeana linked open data pilot. In: International Conference on Dublin Core and Metadata Applications. pp. 94–104 (2011)
6. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *CoRR abs/1412.6980* (2015)
7. Maietti, F., Di Giulio, R., Piaia, E., Medici, M., Ferrari, F.: Enhancing heritage fruition through 3d semantic modelling and digital tools: the inception project. In: IOP conference series: materials science and engineering. vol. 364, p. 012089. IOP Publishing (2018)

8. Pellegrino, M.A., Scarano, V., Spagnuolo, C.: Move cultural heritage knowledge graphs in everyone's pocket (2020)
9. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). pp. 1532–1543 (2014)
10. Qassimi, S., Abdelwahed, E.H.: Towards a semantic graph-based recommender system. a case study of cultural heritage. *J. Univers. Comput. Sci* **27**, 714–733 (2021)
11. Qiu, P., Gao, J., Yu, L., Lu, F.: Knowledge embedding with geospatial distance restriction for geographic knowledge graph completion. *ISPRS International Journal of Geo-Information* **8**(6), 254 (2019)
12. Robusto, C.C.: The cosine-haversine formula. *The American Mathematical Monthly* **64**(1), 38–40 (1957)
13. Schlichtkrull, M., Kipf, T.N., Bloem, P., Van Den Berg, R., Titov, I., Welling, M.: Modeling relational data with graph convolutional networks. In: European semantic web conference. pp. 593–607. Springer (2018)
14. Tietz, T., Waitelonis, J., Zhou, K., Felgentreff, P., Meyer, N., Weber, A., Sack, H.: Linked stage graph. In: SEMANTICS Posters&Demos (2019)
15. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
16. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
17. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Communications of the ACM* **57**(10), 78–85 (2014)
18. Vsesviatska, O., Tietz, T., Hoppe, F., Sprau, M., Meyer, N., Dessì, D., Sack, H.: Ardo: an ontology to describe the dynamics of multimedia archival records. In: Proceedings of the 36th Annual ACM Symposium on Applied Computing. pp. 1855–1863 (2021)
19. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? ArXiv [abs/1810.00826](https://arxiv.org/abs/1810.00826) (2019)
20. Zhang, M., Chen, Y.: Link prediction based on graph neural networks. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 31. Curran Associates, Inc. (2018)
21. Zhu, X., Li, T., De Melo, G.: Exploring semantic properties of sentence embeddings. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 632–637 (2018)