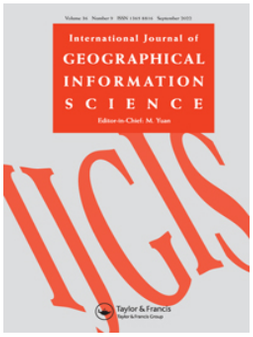

International Journal of Geographical Information Science

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/tgis20

Reasoning over higher-order qualitative spatial relations via spatially explicit neural networks

Rui Zhu, Krzysztof Janowicz, Ling Cai & Gengchen Mai

To cite this article: Rui Zhu, Krzysztof Janowicz, Ling Cai & Gengchen Mai (2022): Reasoning over higher-order qualitative spatial relations via spatially explicit neural networks, International Journal of Geographical Information Science, DOI: 10.1080/13658816.2022.2092115

To link to this article: https://doi.org/10.1080/13658816.2022.2092115

INTERNATIONAL JOURNAL OF GEOGRAPHICAL INFORMATION SCIENCE
https://doi.org/10.1080/13658816.2022.2092115

Taylor & Francis
Taylor & Francis Group

RESEARCH ARTICLE

Reasoning over higher-order qualitative spatial relations via spatially explicit neural networks

Rui Zhu[a,b], Krzysztof Janowicz[a,b,c], Ling Cai[a,b] and Gengchen Mai[d]

[a]STKO Lab, Department of Geography, University of California, Santa Barbara, CA, USA; [b]Center for Spatial Studies, University of California, Santa Barbara, CA, USA; [c]Department of Geography and Regional Research, University of Vienna, Vienna, Austria; [d]Department of Computer Science, Stanford University, CA, USA

ABSTRACT

Qualitative spatial reasoning has been a core research topic in GIScience and AI for decades. It has been adopted in a wide range of applications such as wayfinding, question answering, and robotics. Most developed spatial inference engines use symbolic representation and reasoning, which focuses on small and densely connected data sets, and struggles to deal with noise and vagueness. However, with more sensors becoming available, reasoning over spatial relations on large-scale and noisy geospatial data sets requires more robust alternatives. This paper, therefore, proposes a subsymbolic approach using neural networks to facilitate qualitative spatial reasoning. More specifically, we focus on higher-order spatial relations as those have been largely ignored due to the binary nature of the underlying representations, e.g. knowledge graphs. We specifically explore the use of neural networks to reason over ternary projective relations such as *between*. We consider multiple types of spatial constraint, including higher-order relatedness and the conceptual neighborhood of ternary projective relations to make the proposed model spatially explicit. We introduce evaluating results demonstrating that the proposed spatially explicit method substantially outperforms the existing baseline by about 20%.

ARTICLE HISTORY
Received 18 December 2020
Accepted 16 June 2022

KEYWORDS
Geospatial knowledge graphs; higher-order spatial interactions; qualitative spatial representation and reasoning; spatially explicit methods; GeoAI

1. Introduction

Humans heavily rely on qualitative spatial relations to perceive and reason over space (Lynch 1960, Freksa 1991, Egenhofer and Mark 1995b, Mark et al. 1999). Over the years, numerous researchers have studied the formalization of and reasoning over qualitative geospatial information, including qualitative distance and direction (Frank 1992, Hernandez et al. 1995, Clementini et al. 1997, Worboys 2001, Mossakowski and Moratz 2012, Freksa et al. 2018), topological relations (Egenhofer 1989, Egenhofer and Franzosa 1991, Clementini and Di Felice 1995, Klippel et al. 2013), as well as projective relations (Clementini and Billen 2006, Billen and Clementini 2004, 2005, Clementini

CONTACT Rui Zhu ruizhu@geog.ucsb.edu STKO Lab, Department of Geography, University of California, Santa Barbara, CA, USA

et al. 2010, Beall et al. 2012). It is not until recently, however, that data models, such as directed and labeled graphs (i.e., knowledge graphs), have been explored at scale to represent qualitative spatial relations that complement the strong geometric focus of GIS (Vasardani et al. 2013, Chen et al. 2018, Regalia et al. 2019, Yan et al. 2019). Moreover, most spatial inference engines work on top of a set of axioms and an underlying symbolic calculus (e.g. the region connection calculus (Renz 2002)). These systems work for relatively small datasets that are free of noise, vagueness, and error (Egenhofer 1991, Billen and Clementini 2005). Meanwhile, thanks to computational improvements and big data, a bottom-up, subsymbolic (or connectionist) approach has recently attracted extensive attention in artificial intelligence (Russell and Norvig 2002). In contrast to its symbolic counterpart, a subsymbolic approach does not represent axioms or rules explicitly; instead, they are implicitly learned through the computation of numeric vectors based on observed data (Minsky 1991). With regard to knowledge graphs particularly, plenty of advanced techniques (primarily based on neural networks such as knowledge graph embedding) have been developed to accomplish tasks such as graph completion, i.e., inferring missing links between two nodes (Wang et al. 2017, Ji et al. 2020). With qualitative spatial relations being represented within knowledge graphs, we can, therefore, leverage techniques such as *knowledge graph embedding* to perform spatial reasoning.

More concretely, in *knowledge graph embedding*, both geographic entities (originally represented as symbolic nodes) and their spatial relations (originally represented as symbolic edges) are encoded as numeric vectors (known as *embedding*). Then the validity of a relation statement, which typically consists of three components (i.e., two nodes and one edge), can be determined by computing the *relatedness* of their associated numeric vectors. There are various ways to compute such a relatedness. For example, a relation statement can be regarded as valid if the Euclidean distance between two nodes' embeddings approximates the embedding of the edge.

Nevertheless, due to its inherent binary nature (i.e., using an edge to link two entities), a graph-based data model leads to a strong preference over pairwise relations (e.g., *is part of*, *touches*, and *overlaps*), while failing to address higher-order relations that simultaneously involve more than two spatial features (Zhu et al. 2017, 2019), such as *between* and *surrounding*. Consequently, syntactic structures such as concept reification (Adams and Janowicz 2011) must be adopted to enable the representation of higher-order spatial relations when using graph-based data models. But such a structural modification poses a challenge to typical subsymbolic models in learning representations of nodes and edges as they are mostly designed for pairwise relations.

Hence, this paper aims at developing subsymbolic neural network models that are capable of reasoning over higher-order qualitative spatial relations. We specifically focus on ternary projective relations (i.e., relations that are invariant to projective transformation), including *between*, *before*, *after*, *on the left of*, and *on the right of* (Billen and Clementini 2004). By learning subsymbolic representations (i.e., embeddings) of both places and ternary projective relations based on an incomplete geospatial knowledge graph, the model is designed to predict (1) ternary spatial relations (e.g., *what is the spatial relation between Santa Barbara, Goleta, and Los Angeles?*), and

(2) places (e.g., *which landmarks are visible on the left side while traveling from Santa Barbara to Thousand Oaks?*).

Following this mindset, we first choose to represent qualitative spatial relations using the Resource Description Framework (RDF), which is capable of capturing the semantics of geographic information in a human and machine understandable way (Hitzler et al. 2009, Kuhn et al. 2014). Furthermore, instead of representing a relation as an edge, a class node is constructed as a reification of the higher-order relation that connects with more than two place-typed nodes (also called *values*), each of which has a *role* in the relation. For instance, a ternary projective relation statement (e.g. *Santa Barbara is between San Francisco and Los Angeles*) can be represented as three role-value pairs (e.g. *target* – Santa Barbara, *start* – San Francisco, and *destination* – Los Angeles). With such a geospatial knowledge graph, we then introduce a neural network architecture built from a sequence of operations (e.g. convolution and regression) to learn the embedding for both *roles* and *values*, which are further concatenated to form the embedding of a *role-value pair*. In contrast to classic symbolic approaches, such a subsymbolic approach allows for quantitative computations over role-value pairs (e.g. comparison and transformation). Consequently, the reasoning process over ternary projective relations can boil down to computing the *relatedness* of the three role-value pair embeddings. Specifically, if three role-value pairs are closely related, they are more *compatible* to compose a *valid* ternary relation statement. In order to learn such embeddings from data, we train the model by feeding both valid and invalid sample statements and design a learning objective function to ensure a high compatibility score (computed from the embedding) for a valid relation while keeping such a score low for an invalid one. Finally, for a specific spatial reasoning task (either on spatial relation or place prediction), one can then *rank* the compatibility score of different candidate relation statements using the trained embedding, and the one(s) with the highest score(s) will be selected as the answer to the question.

It is worth emphasizing that such a subsymbolic approach is fundamentally different from its symbolic counterpart in reasoning over qualitative spatial relations. Classic symbolic approach, such as the commonly used composition rules, provides a deterministic set of results while subsymbolic approaches conduct inference based on a ranked list of numerical compatibility scores (can also be interpreted as likelihoods). This difference makes the introduced subsymbolic approach more robust to noise and uncertainty. Plus, given the embedding, the compatibility score of any relation statements can be readily computed, which makes such a reasoning mechanism more computationally effective than symbolic approach, especially on large-scale and loosely-connected graphs.

The proposed method is built upon the state of the art of general n-ary link prediction (Guan et al. 2020), which is a knowledge graph embedding method using deep neural networks. However, in contrast to aspatial n-ary relations, ternary projective relation has its inherent cognitive characteristics, including *strict higher-order relatedness*, *mutual exclusivity*, and *conceptual neighborhood* (Billen and Clementini 2004). Therefore, we design a set of spatially explicit models that take into account these characteristics as spatial constraints. Experimental results demonstrate that a spatially

explicit method substantially outperforms its baseline by about 20%. In summary, we hope to contribute to a better support of higher-order qualitative spatial relations, thus, to a richer means of representing and reasoning of human knowledge in AI systems.

The remainder of this paper is structured as follows: Section 2 reviews related work on qualitative spatial representation, qualitative spatial reasoning via neural networks, and geospatial knowledge graphs. Section 3 proposes ontology design patterns to represent ternary projective relations, providing a platform for Section 4, which further defines two reasoning tasks and introduces a set of spatially explicit neural network-based methods to address them. In order to validate the proposed methods, Section 5 introduces three newly built geospatial knowledge graphs of various sizes to conduct experiments. Section 6 concludes our work and outlines future research directions.

2. Related work

In this section, we discuss related work and background readings relevant to the presented work.

2.1. Qualitative spatial representation and reasoning

With regard to qualitative spatial representation, Vasardani et al. (2013)'s place graphs model is the closest to the geospatial knowledge graph approach studied in this paper. To assist the modeling of human descriptions about the environment, a place graph represents the identified places from texts as nodes and their relations as edges. Following this work, Chen et al. (2018) proposed an extension to the basic *place graphs* model with higher-order relations, such as *betweenness*, being explicitly taken into account. In contrast to their work, which mainly uses labelled property graphs, our work applies the Resource Description Framework (RDF) to simultaneously represent both the data schema (ontology) and data itself as sets of triples. Moreover, the roles in a higher-order relation are distinguished through an attribute (position) associated with the property edge in the extended *place graphs* model while ours directly defines the semantics of edges without using any associated attributes. The geospatial knowledge graph built from our data model can be efficiently integrated with other open graph data repositories such as Wikidata.

To reason over ternary projective relations, Clementini et al. (2010) systematically proposed rules of permutation and composition. Permutation rule further contains converse and rotation. An example of permutation-based reasoning is if P_1 is *on the right side of* P_2 to P_3, then P_1 is *on the left side of* P_3 to P_2. Similar to composition table of topological relations such as RCC8 (Renz 2002), ternary projective relations' composition table aims to find new relations from two given ones. For example, if P_1 is *on the left side of* P_2 to P_3 and P_2 is *on the right side of* P_3 to P_4, then P_1 can be *on the right side, left side, or before* P_3 and P_4. On contrary to these symbolic reasoning, subsymbolic approach via neural networks has advantages including (1) robustness to noise and vagueness, (2) scalability to large data set, as well as (3) capability of implicitly mining more complicated large-scale rules (Russell and Norvig 2002, Calegari et al. 2020).

2.2. Spatial reasoning via neural networks

Rooted in the field of robotics, spatial reasoning using neural networks is one of the main research tracks in AI (Landsiedel et al. 2017). Works in this fields can be principally categorized into text-based, image-based, or hybrid spatial reasoning. Text-based reasoning focuses on answering natural language-based questions that require inferences of spatial relations. Neural networks are typically trained using pre-labeled question and answer pairs (Peng et al. 2015, Mirzaee et al. 2021). Image-based reasoning aims to infer spatial relations between objects based on their visual properties (Hudson and Manning 2019, Krishnaswamy et al. 2019, Peyre et al. 2019). Convolutional neural networks are typically applied in this group of methods. Lastly, combining texts and images to answer questions about spatial relations is an active research topic (Janner et al. 2018, Chen et al. 2019, Huang et al. 2019). The key idea underlying those models is to learn representation of objects from both visual and textual perspectives, and such a model often adopts a functional component that learns the interaction between objects (Santoro et al. 2017). While these methods manage to reason spatial relations from unstructured data (texts and images), our work attempts to do so based on structured data (i.e., geospatial knowledge graphs). Also, to enable a fair comparison with traditional symbolic spatial inference engines, our work does not, even though it is feasible, use any aspatial contextual information to assist spatial reasoning like what the aforementioned text- and image-based approaches do. Finally, none of these works is designed and tested specifically for reasoning over higher-order spatial relations.

2.3. Spatially explicit knowledge graph embedding methods

In past years, several spatially explicit embedding methods have been proposed to implicitly learn representations of geospatial knowledge graphs. For example, Mai et al. (2019) proposed a spatially explicit knowledge graph embedding model to relax and rewrite unanswerable geographic questions in which a distance decay function was used to resample geospatial triples during model training. Similarly, Qiu et al. (2019) modified the loss function in classical translation-based knowledge graph embedding approaches to predict links between spatial features with their distance being considered. Furthermore, to summarize geospatial knowledge graphs, Yan et al. (2019) introduced a reinforcement learning model coupled with new spatial relations to summarize graphs. Despite focusing on distinct applications, these approaches all incorporate geospatial information by leveraging distance decay. In contrast, Mai et al. (2020a) proposed a location-aware knowledge graph embedding model that directly encodes spatial footprints of geographic entities into the embedding space. Unlike the above mentioned work which mainly focuses on *simple* geospatial knowledge graphs which include only binary spatial relations, this work seeks new methods to represent and reason about higher-order spatial relations.

3. Knowledge representation for ternary projective relations

The Resource Description Framework (RDF) [1] is a W3C standard to represent resources – ranging from webpages to actual mountains – by making statements

about them in the form of $\langle subject, predicate, object \rangle$ triples. The graph forms, so to speak, by a set of such triples sharing common elements, e.g. objects. As far as geographic and earth science data is concerned, RDF and ontologies have been successfully used for various problems of data integration, interoperability, conflation, and geographic information retrieval (Bennett et al. 2008, Kuhn et al. 2014, Zhu et al. 2016, Claramunt 2020, Scheider et al. 2020). Due to the binary nature of RDF's triple-based form, *reification* is introduced in this section to represent higher-order spatial relations within RDF. Before that, we first discuss a family of higher-order spatial relations: *ternary projective relation*. In line with the literature, we will use the term knowledge graphs here to signify RDF graphs describing the world around us together with ontologies that formalize the used vocabulary.

3.1. Ternary projective relations

By considering its cardinality, Clementini (2019) categorized spatial relations into unary, binary, ternary, and n-ary. A unary relation is defined as the geometric property of a single spatial feature and a binary relation is about the interaction between two features. Topological relations (Egenhofer 1989), describing the connectedness of features, are the most common binary relations. In contrast, there are only few investigations into ternary (Billen and Clementini 2004, Bloch et al. 2006, Clementini and Billen 2006) as well as n-ary ($n > 3$) spatial relations (Dube and Egenhofer 2014).

Our work focuses on ternary projective relations and follows the formalization proposed by Billen and Clementini (2004). Namely, based on the collinearity of three spatial features, their ternary projective relations are grouped into one of the five types: *between, before, after, right,* and *left* (see the left in Figure 1 as an example)[2]. There are two key properties associated with ternary projective relations. First, the relation among the three spatial features is invariant in terms of *projective transformation*. Secondly, the relation *simultaneously* involves three spatial features and cannot be decomposed into a set of pairwise relations. The latter point becomes particularly important as ternary relations are often conceptually confused with the type of spatial relations that involves three features but can be decoupled into a set of pairwise ones. For example, three features can all overlap each other. This does not constitute a higher-order relation as they can be simply considered as three pairs of binary

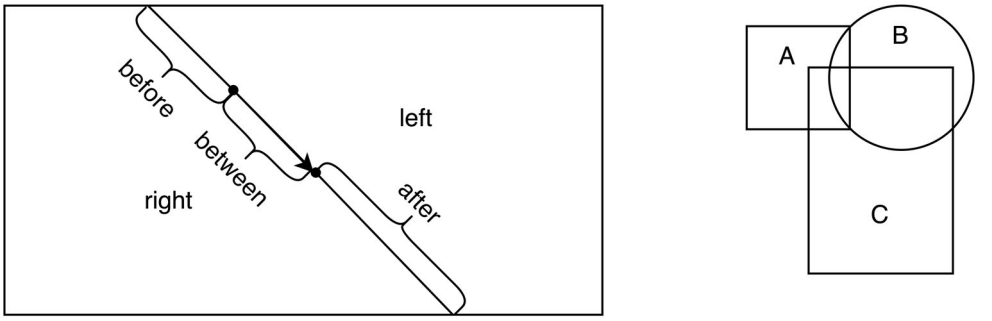

Figure 1. Left: the five ternary projective relation of points. Right: overlap among three spatial features.

overlaps (see the right of Figure 1). Put differently, topological relations like overlap are fundamentally binary while projective relations like betweeness are inherently ternary.

3.2. Relation reification for ternary relations

Unary (e.g., A is a building) and binary spatial relations (e.g., building A is next to building B) can be directly modelled in a knowledge graph as nodes and triples, respectively. On the contrary, representing higher-order (e.g. ternary projective) spatial relations in knowledge graphs is non-trivial and has rarely been studied in the literature. This work leverages *reification* techniques (Noy and Rector 2006) from the Semantic Web to address this challenge. Specifically, instead of representing relations as property edges between two nodes, we construct a new class which is a reification of the spatial relation, thereby the class can further have multiple property edges (usually more than two). An instance of the class linking multiple objects, consequently, represents one higher-order spatial relation statement.

Figure 2 shows an ontology fragment for modeling spatial relations (up to the third-order) in knowledge graphs. Particularly for ternary projective relations, we reify them into five classes, i.e., *Left*, *Right*, *Before*, *After*, and *Between*, and they share two modular design patterns (Hitzler and Shimizu 2018). First of all, for spatial relations of *left, right, before*, and *after*, they can all be represented by the pattern shown on the bottom left of Figure 2 (with the *Left* class as an example), in which three edges: *has_target, has_origin*, and *has_destination* are designed to express the role of three spatial features of interest (e.g. referred as places in this paper). Specifically, *has_target* links to the place that is to be located while *has_origin* and *has_destination* together define the origin and direction of the frame of reference using the other two places (Zhu

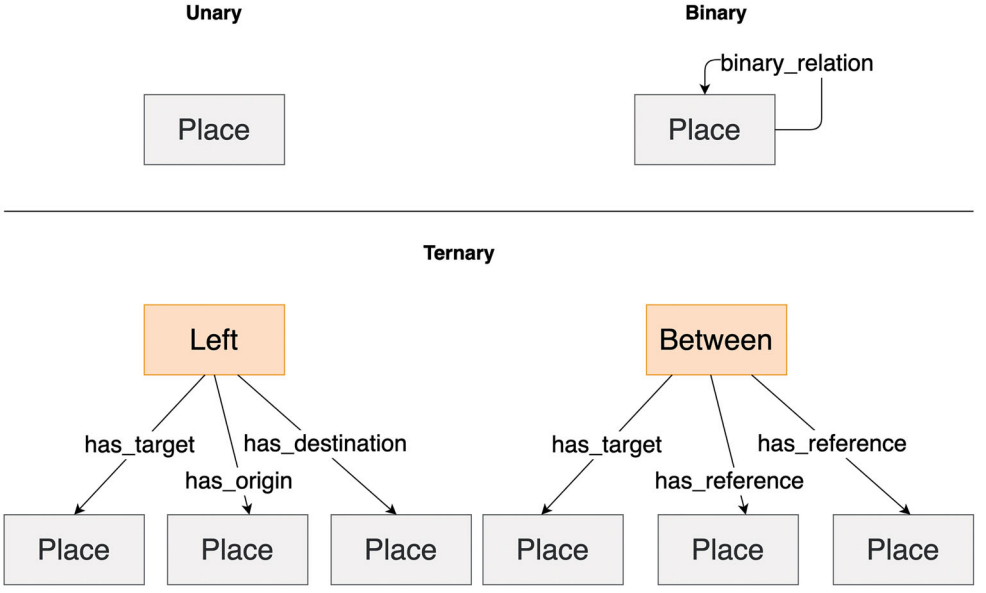

Figure 2. Ontology for unary (top left), binary (top right), and ternary relations (bottom).

et al. 2019). Secondly, due to its symmetric property (i.e., "A is between B and C" means the same to "A is between C and B"), we present a different design pattern to model the relation of *between* (see the bottom right of Figure 2). Since no distinction exists for the direction of the reference system as in the former pattern, we expect an equal role for the origin and destination, thereby we label both as *has_reference*. Through relation reification, we are able to preserve the semantics of higher-order qualitative spatial relations on one hand and on the other keep using triples as defined in RDF to build the geospatial knowledge graph.

3.3. Ternary projective relation example

Figure 3 illustrates an example of using this representational pattern for ternary projective relations. Five cities in California (i.e., Santa Barbara, Ventura, Fillmore, Thousand Oaks, and Los Angeles) are labelled on the map as points. They yield ${}^5P_3 = \frac{5!}{(5-3)!} = 60$ (P is the permutation) ternary relation statements at most (but less due to symmetry), among which three are extracted and represented as a knowledge graph fragment in Figure 3. The orange nodes, *before_1*, *between_1*, and *left_1* represent instances of the ternary relation classes: *Before, Between*, and *Left*, respectively. These relation statements (the instances) define three different ternary projective relations among the five cities, which are instances of the *Place* class and shown as grey nodes

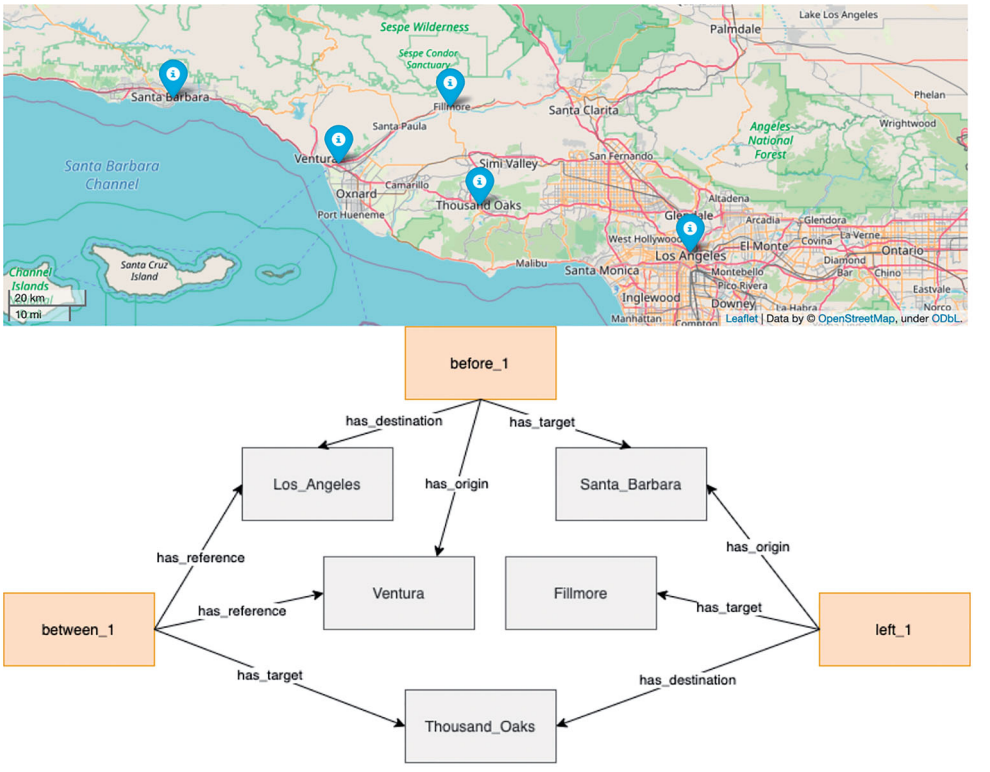

Figure 3. Top: spatial organization of five cities: Santa Barbara, Ventura, Thousand Oaks, Fillmore, and Los Angeles. Bottom: exemplary graph of ternary relation statements among the five cities.

in the graph. It is worth highlighting that even though this example only shows ternary projective relations, other types of higher-order spatial relations, e.g. *surroundedness*, can be represented in a similar way.

4. Reasoning over ternary relations through neural networks

Instead of relying on traditional spatial inference engine, we investigate reasoning over ternary relations from a bottom-up, subsymbolic (or connectionist) perspective. Specifically, we leverage deep learning to propose a neural network architecture that explicitly incorporates cognitive principles as spatial constraints to perform spatial reasoning over ternary projective relations.

4.1. Definitions

Our work focuses on reasoning about ternary projective relations to answer two types of questions. The first type is about inferring spatial relation among three observed places (or locate-able entities more broadly). We call the related task *relation prediction*. The second type aims at predicting a place given the other two and their ternary relation. We call this task *place prediction*. In this section, we formally define the involved relations and these two tasks. Jointly, they form the basis for the neural network architecture that will be introduced in Section 4.2.

Definition 1 (Ternary relations and statements about them). Each reified ternary relation t_i is associated with exactly three role-value pairs - $\langle$*target:place, origin:place, destination:place*$\rangle$ (for the relation *between*, both *origin* and *destination* will be replaced by *reference*). Accordingly, statements can be denoted as a tuple of the form: $t_i = \langle r_{i1} : v_{i1}, r_{i2} : v_{i2}, r_{i3} : v_{i3} \rangle$ with $\langle r_{i1}, r_{i2}, r_{i3} \rangle \in R^*$ and $v_{i1}, v_{i2}, v_{i3} \in V$. The value set V is defined as $\{place_1, \ldots, place_M\}$, where M is the number of places in the knowledge graph. The role set R^* for ternary projective relations is defined as $R^* = \{R_{left}, R_{right}, R_{before}, R_{after}, R_{between}\}$, where $R_{left} = \langle$*left_target, left_origin, left_destination*$\rangle$ (note that R_{right}, R_{before}, and R_{after} are denoted by replacing the prefix *left* with *right, before*, and *after*, respectively) and $R_{between} = \langle$*between_target, between_reference, between_reference*$\rangle$. Note that R^* is a set of grouped roles based on their dependency in the flat role set: $R = \{$*left_target, right_origin, before_destination,* $\ldots$ *, between_reference*$\}$. For ternary projective relations, there are 14 such roles in total. Subsequently, a ternary knowledge graph T is defined as a collection of ternary relation statements: $T = \{t_1, \ldots, t_i, \ldots, t_N\}$, where $1<i<N$ and N is the overall number of relation statements. For the sake of simplicity, the subscript i in t, r and v will be skipped whenever no confusion arises.

Example 1. The reified relation statement *between_1* in Figure 3 can be stated as: $\langle$*between_target: Thousand_Oaks, between_reference: Ventura, between_reference: Los_Angeles*$\rangle$. In plain English it corresponds to *Thousand Oaks is between Ventura and Los Angeles*.

Definition 2 (Relation Prediction). Given an incomplete relation statement $t = \langle ?r_1 : v_1, ?r_2 : v_2, ?r_3 : v_3 \rangle$, where $?r_1, ?r_2, ?r_3$ are missing roles, the task of relation prediction is defined as the process of inferring $\langle ?r_1, ?r_2, ?r_3 \rangle$ given the observed values v_1, v_2, v_3.

Example 2. Using the spatial organization of cities in Figure 3 as an example, one relation prediction task can be to predict the missing roles in $\langle ?r_1 : Santa\_Barbara, ?r_2 : Ventura, ?r_3 : Los\_Angeles \rangle$. Put into natural language, *what is the spatial relation among Santa Barbara, Ventura, and Los Angeles?*. One correct prediction would be: $\langle ?r_1, ?r_2, ?r_3 \rangle = R_{before}$ (*Santa Barbara is before the route from Ventura to Los Angeles*).

Definition 3 (Place Prediction). Given an incomplete relation statement $t = \langle r_1 : ?v_1, r_2 : v_2, r_3 : v_3 \rangle$, where $?v_1$ is a missing value and its role r_1 is a *target* (e.g., if the relation is about *right*, $r_1 = right_{target}$), the task of place prediction is to predict the missing value $?v_1$ given the other two values (v_2 and v_3) and their roles $\langle r_1, r_2, r_3 \rangle$.

Example 3. Using Figure 3 as an example, one place prediction task would be to infer the missing place $?v_1$ in $\langle$*left_target:*$?v_1$, *left_origin:Santa_Barbara, left_destination:Thousand_Oaks*$\rangle$. This can be phrased in natural language as *which city is on the left side from Santa Barbara to Thousand Oaks?* There can be many correct predictions for this question, and $?v_1 =$ *Fillmore* is one of them (*Fillmore is on the left side from Santa Barbara to Thousand Oaks*).

4.2. Baseline model

Our baseline model is inspired by the state of the art in n-ary link prediction, which is developed to predict general (mostly aspatial) relations in knowledge graphs (Guan et al. 2020). The neural network architecture, specifically for ternary relations, is illustrated in Figure 4. The main goal is to learn a *compatibility score* for each ternary relation statement t so that if t is true, the model would result in a high compatibility score; otherwise, the score would be low. The compatibility score is evaluated based on the *relatedness* of role-value pairs that are involved in t, which is further computed through the *embedding* of the involved roles and values. The embedding for roles and values are looked up from their embedding matrix: $\boldsymbol{E_R} \in \mathbb{R}^{|R| \times k}$ and $\boldsymbol{E_V} \in \mathbb{R}^{|V| \times k}$, respectively, where k is the dimension of the embedding, which is a hyper-parameter of the model, and $|R|$ and $|V|$ indicate the number of roles and values, respectively. The whole process is trained through a neural network architecture that includes three main components: (1) ternary relation statement embedding; (2) relatedness computation; and (3) scoring. Moreover, in order to provide the model with both positive and negative training samples, the technique of (4) negative sampling is adopted (see Wang et al. (2017) for a comprehensive review), which is not explicitly depicted in the architecture but is a key component for the training process. We discuss each of these components in detail in this section, together with two essential properties of this model at the end.

4.2.1. Ternary relation statement embedding

For one ternary relation statement t, we compose its embedding $\boldsymbol{E_t}$ by concatenating the embedding of its three role-value pairs, each of which is further a concatenation of its role embedding $\boldsymbol{R_t} \in \mathbb{R}^{1 \times k}$ and value embedding $\boldsymbol{V_t} \in \mathbb{R}^{1 \times k}$, which are the corresponding rows in $\boldsymbol{E_R}$ and $\boldsymbol{E_V}$, respectively. To enhance expressivity, a convolution layer

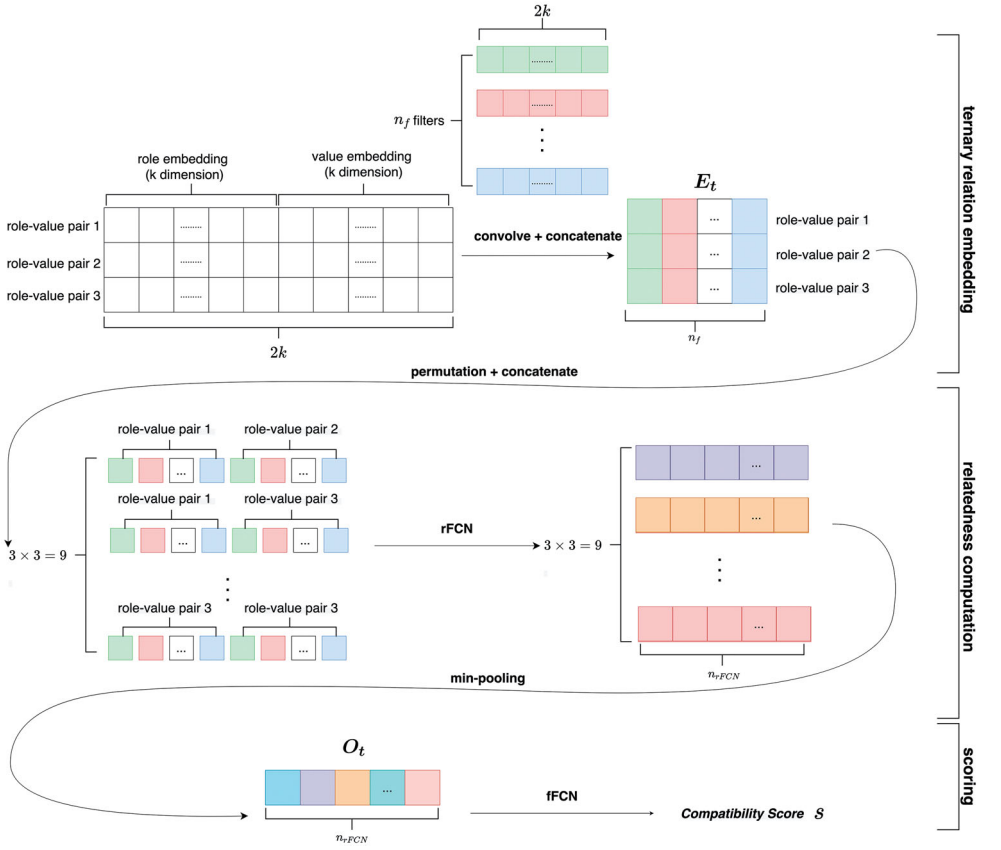

Figure 4. Neural network architecture of the baseline model.

of n_f filters (denoted as $\Omega \in \mathbb{R}^{n_f \times 2k}$) is applied on the concatenated role-value pair embeddings and Rectified Linear Units (ReLU) are utilized as the activation function (Nair and Hinton 2010), which lead to the feature matrix $\boldsymbol{E_t} \in \mathbb{R}^{3 \times n_f}$. Formally, the embedding of a ternary relation statement t is computed as:

$$\boldsymbol{E_t} = concat(ReLU(concat(\boldsymbol{R_t}, \boldsymbol{V_t}) * \Omega)) \tag{1}$$

where $*$ denotes convolution operation, and *concat* indicates the process of concatenation.

4.2.2. Relatedness computation

Guan et al. (2020) transfers the problem of validating a relation statement into measuring the *overall relatedness* of the involved role-value pairs. Intuitively, if all the three role-value pairs are closely related, the relation statement tends to be true. However, when the number of role-value pairs in a relation is greater than two (like the ternary projective relation), the computation becomes nontrivial. Therefore, Guan et al. (2020) simplifies the computation into first evaluating the *pairwise relatedness* of role-value pairs, based on which the *overall relatedness* is subsequently measured by *min-pooling*.

In terms of computing *pairwise relatedness*, the three role-value pair embeddings $\boldsymbol{E}_{\boldsymbol{t}}^{(j)} \in \mathbb{R}^{1\times n_f}$ ($j = 1, 2, 3$), which indicates the j^{th} row of matrix $\boldsymbol{E}_{\boldsymbol{t}}$, are first concatenated pairwisely. As Figure 4 illustrates, the three role-value pairs are concatenated into nine (3×3) possible permutations. Subsequently, a fully connected layer (FCN), coupled with a ReLU activation function, is imposed on each of the nine concatenated permutations, which results in nine feature vectors, each with a dimension of $1 \times n_{rFCN}$, where n_{rFCN} indicates the number of layers in this relatedness FCN (*rFCN*). The feature vector represents the relatedness of a pair in some respects, which are indicated by the value at each dimension. Finally, the *overall relatedness* of the nine permutations is consolidated through a *min-pooling*, where the minimal value in each dimension is extracted to build the overall relatedness feature vector $\boldsymbol{O}_{\boldsymbol{t}} \in \mathbb{R}^{1\times n_{rFCN}}$. In summary, this component can be formalized as:

$$\begin{aligned} \boldsymbol{O}_{\boldsymbol{t}} &= min_{j,k=1}^{3}(rFCN(concat(\boldsymbol{E}_{\boldsymbol{t}}^{(j)}, \boldsymbol{E}_{\boldsymbol{t}}^{(k)}))) \\ &= min_{j,k=1}^{3} ReLU(concat(\boldsymbol{E}_{\boldsymbol{t}}^{(j)}, \boldsymbol{E}_{\boldsymbol{t}}^{(k)})\boldsymbol{W}_{\boldsymbol{rFCN}} + \boldsymbol{b}_{\boldsymbol{rFCN}}) \end{aligned} \tag{2}$$

where $\boldsymbol{W}_{\boldsymbol{rFCN}} \in \mathbb{R}^{2n_f \times n_{rFCN}}$ and $\boldsymbol{b}_{\boldsymbol{rFCN}} \in \mathbb{R}^{n_{rFCN}}$ are the weight matrix and bias vector of *rFCN*, respectively.

Using a fully connected layer (FCN) to compute the relatedness of two features of interest is a widely adopted practice in deep learning, especially in computer vision (Santoro et al. 2017). In addition, the rationale behind applying min-pooling rather than other types of pooling techniques lies in the assumption that if the minimum relatedness value along a dimension is large enough, then the overall relatedness of this dimension would be large.

4.2.3. Scoring

Once the overall relatedness vector $\boldsymbol{O}_{\boldsymbol{t}}$ is computed, another fully connected layer (*fFCN*) is utilized to evaluate the final compatibility score s of a ternary relation statement t. Formally, the process is represented as:

$$\begin{aligned} s(t) &= fFCN(\boldsymbol{O}_{\boldsymbol{t}}) \\ &= \boldsymbol{O}_{\boldsymbol{t}}\boldsymbol{W}_{\boldsymbol{fFCN}} + \boldsymbol{b}_{\boldsymbol{fFCN}} \end{aligned} \tag{3}$$

where $\boldsymbol{W}_{\boldsymbol{fFCN}}$ and $\boldsymbol{b}_{\boldsymbol{fFCN}}$ are the weight matrix and bias vector of *fFCN*, respectively.

4.2.4. Negative sampling

In addition to the aforementioned three components that are directly related to the neural network architecture, negative sampling is another essential step in the training process. Negative sampling is not only applied in the baseline model, but has also been intensively used in almost all types of knowledge graph embedding methods, such as TransE (Bordes et al. 2013) and DisMult (Yang et al. 2015). In order to enable the model to comprehensively learn what is correct and what is wrong so that it can accurately produce inferences on new inputs, the technique of negative sampling generates negative samples based on what is known from the positive ones (i.e., training data). For example, given a positive ternary relation statement $t^+ = \langle r_1 : v_1, r_2 : v_2, r_3 : v_3 \rangle$, negative samples can be generated by replacing its roles or values, by some randomly selected entities from the flat role set R, or value set V. Formally, the full set T^-

of candidate negative samples for one positive statement t^+ can be defined as:

$$\begin{aligned} T^-(t^+) \quad &= \{\langle r'_1 : v_1, r'_2 : v_2, r'_3 : v_3\rangle | r'_1 \in R \wedge r'_2 \in R \wedge r'_3 \in R \wedge \langle r'_1 : v_1, r'_2 : v_2, r'_3 : v_3\rangle \notin T^+\} \\ &\cup \{\langle r_1 : v'_1, r_2 : v_2, r_3 : v_3\rangle | v'_1 \in V \wedge \langle r_1 : v'_1, r_2 : v_2, r_3 : v_3\rangle \notin T^+\} \\ &\cup \{\langle r_1 : v_1, r_2 : v'_2, r_3 : v_3\rangle | v'_2 \in V \wedge \langle r_1 : v_1, r_2 : v'_2, r_3 : v_3\rangle \notin T^+\} \\ &\cup \{\langle r_1 : v_1, r_2 : v_2, r_3 : v'_3\rangle | v'_3 \in V \wedge \langle r_1 : v_1, r_2 : v_2, r_3 : v'_3\rangle \notin T^+\} \end{aligned} \tag{4}$$

where T^+ represents the original knowledge graph with $+$ emphasizing that all statements in T^+ are positives. Note that in contrast to Guan et al. (2020), the proposed baseline model replaces the three roles all together (role-replacement as shown in the first line of Formula 4) rather than individually as the way of replacing values (value-replacement as illustrated in the second to fourth line of Formula 4). This modification aligns with the characteristic of ternary projective relations that if one role of the relation changes, the other two will also change accordingly.

While training the baseline model, for each positive sample t^+, one negative sample is extracted from its candidate set $T^-(t^+)$, with a probability of it being either role-replacement or value-replacement proportional to the cardinality of role set ($|R|$) and value set ($|V|$), respectively. Having both positive and negative samples, the *loss function*, a logistic loss similar to the one used in Yang et al. (2015), is defined as:

$$\mathcal{L}(\tilde{t}) = \log\left(1 + e^{-l_{\tilde{t}}} s(\tilde{t})\right) \tag{5}$$

where $l_{\tilde{t}} = 1$ if $\tilde{t}$ is a positive sample; otherwise, $l_{\tilde{t}} = -1$. The established neural network is optimized through backpropogation, and Adam (Kingma and Ba 2015) is implemented as the stochastic optimization method.

Lastly, it is worth highlighting *two properties* of the baseline model. First, the model is permutation free. Namely, thanks to the permutation process discussed in 4.2.2, the model is agnostic to the order of input role-value pairs (see Guan et al. (2020) for a detailed proof). Secondly, even though we focus on ternary relations in this work, both binary and higher-order spatial relations, as well as their combinations, can be studied using the same framework because no matter how many role-value pairs are involved in the relation, the model outputs one single compatibility score.

4.3. Spatially explicit methods

In Section 3 we argued that spatial is special in terms of *representing* ternary projective relations in knowledge graphs. This section focuses on investigating methods to inject those spatial constraints about ternary projective relations into the introduced *reasoning* model so that it becomes spatially explicit (Yan et al. 2019, Janowicz et al. 2020, Mai et al. 2020b). We will focus on three types of spatial constraint, based on which corresponding modifications are proposed to improve the baseline model. In addition, we also introduce an attention mechanism into the neural network architecture in order to optimize the overall relatedness computation.

4.3.1. Ternary projective relations are strictly higher-order

Even though the baseline model can process higher-order relations, it is specifically designed for and merely evaluated on general knowledge bases, such as Wikidata, in

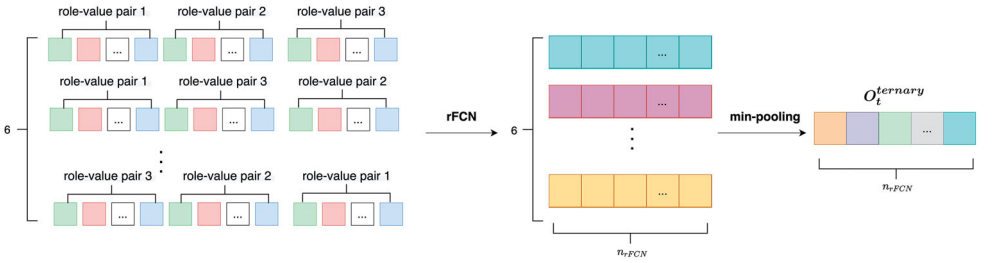

Figure 5. Higher-order relatedness computation.

which the definition of higher-order relations is relatively relaxed. Namely, all relations that involve more than two features are considered as higher-order. This definition does not distinguish cases whether the higher-order relation can be decomposed into a set of pairwise relations or not. As argued by Guan et al. (2020), the statement "Marie Curie received Nobel Prize in Physics in 1903 together with Henri Becquerel and Pierre Curie" can be represented as a higher-order relation ⟨*person:Marie Curie, award:Nobel Prize in Physics, point_in_time: 1903, together_with:Henri Becquerel, together_with:Pierre Curie*⟩. However such a relation can, in fact, be decomposed into lower-order (e.g., pairwise) relations. For instance, the binary relation ⟨*person:Marie Curie, award:Nobel Prize in Physics*⟩ still holds true. In contrast, the definition of higher-order in ternary projective relations is strict, as discussed in Section 3.1. The involved three places are *simultaneously* related in the relation; one cannot decompose them into a set of pairwise relations. For example, the binary relation ⟨*between_target: Ventura, between_reference: Los_Angeles*⟩ is invalid and its semantics is missing without the third role-value pair.

This fundamental difference has practical implications with respect to the neural network architecture design, particularly to the component of relatedness computation. If the higher-order relation can be decomposed into pairwise relations, the computation of overall relatedness can be simplified through the evaluation of pairwise relatedness (Figure 4 and Section 4.2.2). However, if a relation is strictly higher-order, measuring the relatedness through its pairwise counterparts might blindly discard the essential semantics. Therefore, this work proposes a new component for relatedness computation, in which the original pairwise permutations are replaced by a set of third-order ones (Figure 5). Furthermore, in contrast to the baseline model, we do not consider permutations where there are duplicated role-value pairs (e.g., in Figure 4, the relatedness of *role-value pair 3* and *role-value pair 3* is considered). As a result, a ternary projective relation includes in total six third-order permutations (${}^3P_3 = \frac{3!}{(3-3)!} = 6$) as shown in Figure 5.

The formula to compute higher-order relatedness is hence modified as:

$$\begin{aligned} \boldsymbol{O_t^{ternary}} &= min^3_{j,k,l=1,j\neq k\neq l}(rFCN(concat(\boldsymbol{E_t^{(j)}}, \boldsymbol{E_t^{(k)}}, \boldsymbol{E_t^{(l)}}))) \\ &= min^3_{j,k,l=1,j\neq k\neq l} Relu(concat(\boldsymbol{E_t^{(j)}}, \boldsymbol{E_t^{(k)}}, \boldsymbol{E_t^{(l)}})\boldsymbol{W_{rFCN}} + \boldsymbol{b_{rFCN}}) \end{aligned} \tag{6}$$

4.3.2. Ternary projective relations are mutually exclusive

Another characteristic of ternary projective relation is that the five relations are mutually exclusive (Billen and Clementini 2004). For example, if we know the relation among three sequenced places is *left* (i.e., the corresponding role tuple is R_{left}), the other four relations *right, before, after*, and *between* become invalid on these three places in the same sequence. This argument does not necessarily hold for general (e.g., aspatial) ternary relations.

To take advantage of this spatial constraint that is specific to ternary projective relation, we improve the negative sampling process by replacing the random strategy (discussed in Section 4.2.4). Specifically, for a positive ternary relation statement t^+ from the training data set, we generate a negative sample candidate set, which includes all possible *hard* negatives according to their mutually exclusive relation with the positive. Taking *left_1* in Figure 3 as an example, the positive sample ⟨*left_target: Fillmore, left_origin: Santa_Barbara, left_destination: Thousand_Oaks*⟩ can produce a negative sample set by iteratively replacing roles of *left* to their correspondences in R_{right}, R_{before}, R_{after}, and $R_{between}$. For instance, ⟨*right_target: Fillmore, right_origin: Santa_Barbara, right_destination: Thousand_Oaks*⟩ is one *hard* negative sample in the candidate set.

Accordingly, the randomly sampled $T^-(t^+)$ in Formula 4 is updated to a spatially constrained sample set $T^{-*}(t^+)$ as:

$$\begin{aligned} T^{-*}(t^+) \quad &= \{\langle r'_1 : v_1, r'_2 : v_2, r'_3 : v_3\rangle | \langle r'_1, r'_2, r'_3\rangle \in (R^* \backslash \langle r_1, r_2, r_3\rangle)\} \\ &\cup \{\langle r_1 : v'_1, r_2 : v_2, r_3 : v_3\rangle | v'_1 \in V \wedge \langle r_1 : v'_1, r_2 : v_2, r_3 : v_3\rangle \notin T^+\} \\ &\cup \{\langle r_1 : v_1, r_2 : v'_2, r_3 : v_3\rangle | v'_2 \in V \wedge \langle r_1 : v_1, r_2 : v'_2, r_3 : v_3\rangle \notin T^+\} \\ &\cup \{\langle r_1 : v_1, r_2 : v_2, r_3 : v'_3\rangle | v'_3 \in V \wedge \langle r_1 : v_1, r_2 : v_2, r_3 : v'_3\rangle \notin T^+\}, \end{aligned} \tag{7}$$

In contrast to Formula 4, where roles are sampled randomly from the flat role set R, the replacement of roles in Formula 7 is constrained by grouped role families in R^*. Note that the value-replacement part (second to fourth line) are the same between the two formulas.

4.3.3. Conceptual neighborhood of ternary projective relations

Freksa and Kreutzmann (2016) defined conceptual neighborhoods as "direct discrete transitions between temporal and spatial relations". Initially, they were introduced to study qualitative temporal and spatial reasoning (Freksa 1991, Freksa et al. 1991, Egenhofer and Mark 1995a). With respect to ternary projective relations, the five relations can also be correlated through the notion of their conceptual neighborhoods. As Figure 1 (left) illustrates, the five relations separate the region into five sub-regions, whose organization implies the possibility of transition between regions[3]. For example, *right* can be directly transited to *between, before*, and *after* but not *left*. So the conceptual neighborhood of relation *right* is composed of the former three. A complete list of the conceptual neighborhood for each relation is presented in Table 1. Note that this is one way to form such conceptual neighborhoods for ternary projective relations. The same is true for Freksa's temporal conceptual neighborhood.

Next, we inject this conceptual neighborhood induced constraint into our model by a new negative sampling strategy. The underlying assumption is that a relation is easier to be confused with its conceptual neighbors. Therefore, in order to allow

Table 1. Conceptual neighbourhood of ternary projective relation.

Relation	Conceptual neighbourhood
left	*before, after, between*
right	*before, after, between*
before	*left, right, between*
after	*left, right, between*
between	*left, right, before, after*

the model to improve its performance, we explicitly feed it with more difficult samples that are from the conceptual neighborhood. Taking *left_1* in Figure 3 as an example again: one negative sample constrained by conceptual neighborhood is ⟨*between_target: Fillmore, between_reference: Santa_Barbara, between_reference: Thousand_Oaks*⟩, while the sample built from *right* as demonstrated in Section 4.3.2 will not be selected anymore since *right* is not a conceptual neighbour of *left*. The formula for this negative sampling strategy is similar to Formula 7 only with R^* being replaced by a rule-based role set according to Table 1.

Finally, it is worth emphasizing that even though generated samples of the two strategies introduced in Section 4.3.2 and 4.3.3 overlap, the rationales behind them are clearly distinct.

4.3.4. Attention vs. min-pooling

When computing the overall relatedness of role-value pairs, the baseline model applies the technique of min-pooling, with an assumption that if the minimum relatedness value of a dimension across all the permutations is large enough, the overall relatedness of the corresponding dimension will be significant enough so that the final compatibility score can be appropriately computed to validate a relation. One can argue that it is an oversimplified assumption for two reasons. First, solely extracting the minimum value would result in a loss of information (i.e., relatedness values in other non-minimum permutations are discarded). Secondly, it neglects the complex interaction among all permutations, which might play roles in evaluating the overall relatedness. Therefore, we propose to replace min-pooling with an *attention* based module (Bahdanau et al. 2014), in which all relatedness values in a dimension are taken into account with weights proportional to their relative importance in measuring the overall relatedness. An updated version of Formula 2 is, therefore, illustrated as follows (the ternary version, i.e., Formula 6, can be updated in the same way):

$$\boldsymbol{O}'_{\boldsymbol{t}} = atten^{3}_{j,k=1}(rFCN(concat(\boldsymbol{E}^{(\boldsymbol{j})}_{\boldsymbol{t}}, \boldsymbol{E}^{(\boldsymbol{k})}_{\boldsymbol{t}}))) \tag{8}$$

Figure 6 explains the attention module $atten()$. To compute the feature value o^q at the q^{th} dimension ($q \in \{1, 2, \ldots, n_{rFCN}\}$) of the overall relatedness vector $\boldsymbol{O}'_{\boldsymbol{t}}$, corresponding values of all permutations e^q_r ($r \in \{1, 2, \ldots, N_p\}$, N_p is the number of permutations, which is 9 for ternary projective relations) are aggregated with weights α^q_r, which are computed based on their relative importance in the overall relatedness. The process can be formally modeled using Formula 9 and 10. We can imagine o^q as the *expected* relatedness over all possible permutations, each of which has its own probability (α^q_r) of attending to o^q. Note that this attention module is part of the neural

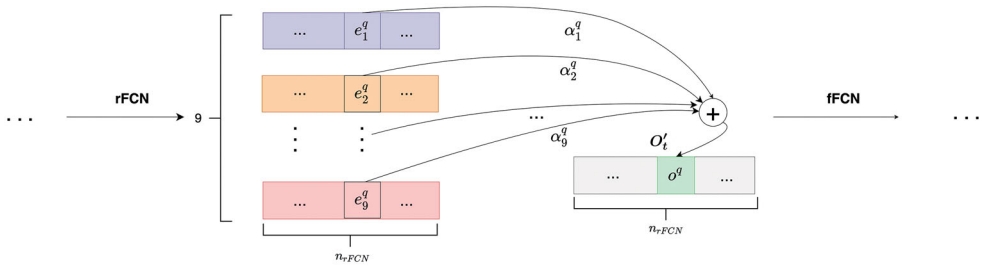

Figure 6. Attention module.

network architecture, so the gradient from the loss function can backpropagate through it allowing a joint learning of the weight (α_r^q), the relatedness computation (e_r^q), as well as the overall compatibility score (s).

$$o^q = \sum_{r=1}^{N_P} \alpha_r^q e_r^q \tag{9}$$

$$\alpha_r^q = \frac{\exp\left(e_r^q\right)}{\sum_{h=1}^{N_P} exp(e_h^q)} \tag{10}$$

5. Experiments

To evaluate the proposed spatially explicit models (Section 4.3.1–4.3.4), we perform two tasks: relation prediction and place prediction, which are formally defined in Section 4.1, and compare their performances to the baseline model (Section 4.2). In addition, to gain an understanding of the spatial reasoning capability of our neural network architecture, we further design experiments to explore the impact of graph density as well as the number of negative samples.

5.1. Data

Since there is no existing dataset that facilitates experiments on reasoning over ternary projective relations in a large scale, we first created one by extracting significant cities (in Arizona, California, Nevada, and Oregon) from DBpedia Places[4], which is a gazetteer built from user contributed content in Wikipedia[5]. In total, we collected 764 cities. Given that every three cities can generate three ternary projective relations, we will end up having about 440 million relations, which is infeasible for most system to execute. Besides, humans certainly do not rely on memorizing a large number of relations to understand the spatial organization of our environment; we rather use a small number of relations assisted with our intuitive reasoning ability to infer the missing statements on-demand (Egenhofer and Mark 1995b, Freksa 2020). Therefore, we only sampled a subset of relations among these 764 cities and attempted to investigate the impact of the subset size on reasoning performance. We hypothesize that with denser knowledge graphs, our model would perform better.

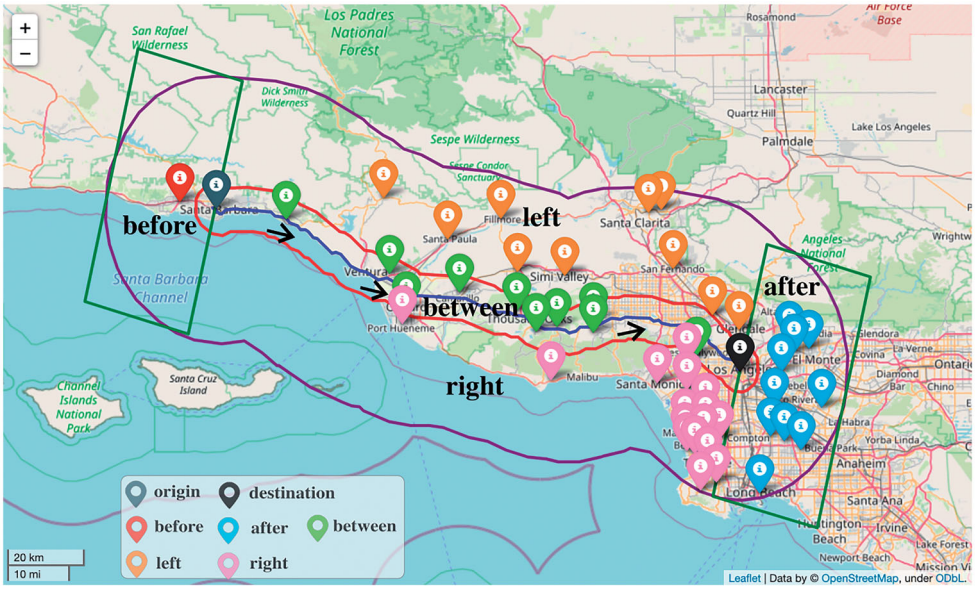

Figure 7. Example of computing ternary projective relations.

Among the 764 cities, we first sampled 2000 pairs, from which the origin and destination of the ternary projective relation were selected. In most gazetteers cities are represented as points without a spatial extent. Moreover, these cities are distributed on a geographic space rather than a Euclidean plane. Hence simply applying Billen and Clementini (2004)'s formalization of ternary projective relation (Figure 1 left) becomes problematic. Instead of drawing a straight line between two points, we then used driving route between the two points, together with its buffer zones, to separate the space in order to determine the five relations. We believe such a route-based formalization is closer to human's perception of ternary projective relations for spatial features like cities, while extensive studies have to be conducted to validate it, which is beyond the scope of this paper. More concretely, we built *five buffer zones* to determine the five relations. As Figure 7 illustrates, the five buffer zones were created by four buffers: the red and purple buffer zones were computed based on the route from the origin to destination (blue line) with different buffer distances (5 *km* and 30 *km*, respectively); while the two green zones were based on the origin and destination, respectively (buffer distance is 30 *km*). Based on the five buffer zones (named left, between, right, before, and after in Figure 7), cities were consequently grouped into five relations, represented as points of different colors in Figure 7, with respect to the origin and destination. Detailed process of computing the five buffer zones can be found in Appendix A.

Once the full set of ternary projective relation statements for each of the 2000 sampled pairs were computed, we randomly sampled L ($L \in \{10, 20, 30\}$) relation statements from the full set (if its size is less than L, the full set will be kept). Having different L is important for testing the role of graph density on the reasoning capability of proposed methods. Moreover, we split the data into 80%, 10%, and 10% for training,

Table 2. Statistics about the graph (N_{train}, N_{valid}, and N_{test} represent the number of relations for training, validation, and testing, respectively.).

Name	L	N_{train}	N_{valid}	N_{test}
West2000_10	10	15895	1987	1987
West2000_20	20	31578	3947	3947
West2000_30	30	46871	5859	5859

validation, and testing, respectively. In summary, we generated three geospatial knowledge graphs of different densities and their statistics are shown in Table 2.

5.2. Evaluation metrics

In terms of relation prediction, we remove all the three roles in a testing statement and replace them with roles in R^* to generate a set of candidate statements. Then each of the candidate statements is fed into the trained model and a compatibility score is computed. Subsequently, candidate statements are ranked according to their compatibility scores, based on which an assessment of the trained model is produced. The process to test place prediction is similar but with only one value (i.e., the target place) being replaced by entities in V when generating candidate statements. We apply several evaluation metrics to assess these two tasks.

5.2.1. Relation prediction

Given three places, relation prediction aims at inferring their relation. Similar to other work on link prediction (Wang et al. 2017), we use the mean reciprocal rank (*MRR*), as well as *Hits@n* ($n = 1$ or 3) to evaluate this task. *MRR* measures the average reciprocal rank of the correct statements in predictions, and *Hits@n* calculates the proportion of predictions where the correct statement is ranked in the top n. The performance of a model is greater with higher values of these three metrics.

5.2.2. Place prediction

In contrast to relation prediction, where there are only five families of relations (see R^*), the prediction of a target place given the origin and destination as well as their relation is more challenging, and there are often numerous correct answers. For example, the question "which city is in between San Francisco and Los Angeles?" will result in dozens of places. Therefore, we apply metrics that are commonly used in recommender systems (Bobadilla et al. 2013) to assess performance. More concretely, we use *precision@k* ($k = 1, 5,$ or 10), which indicates the proportion of correct places in top k of the ranked list. The larger *precision@k* is, the better a method performs.

5.3. Implementation details

To find the best parameter setting for the proposed method, we tuned hyperparameters using the validation data. The set of hyperparameters that fits the model the best was consequently selected. Specifically, a grid search was executed to select the learning rate $\lambda \in \{0.00001, 0.0001, 0.001, 0.01, 0.1\}$, and random searches (Bergstra and Bengio 2012) were implemented to choose the embedding dimension $k \in$

Table 3. Abbreviation (Abbr.) and corresponding section(s) of experimented methods.

Method	Abbr.	Section(s)
Baseline Model	*Pairwise*	4.2.1–4.2.3
Higher-order relatedness	*Ternary*	4.3.1
Random negative sampling	*RND*	4.2.4
Mutual exclusivity – constrained negative sampling	*ME*	4.3.2
Conceptual neighbourhood – constrained negative sampling	*CN*	4.3.3
Attention module	*Atten*	4.3.4

$\{50, 100\}$, the number of filters $n_f \in \{50, 100, 200, 400, 500\}$, as well as the dimension of the relatedness FCN $n_{rFCN} \in \{50, 100, 200, 400, 500, 800, 1000, 1200\}$. Details of the tuning process can be found in Appendix B. The finally adopted hyperparameters are: $\lambda = 0.001, k = 100, n_f = 100, n_{fFCN} = 1000$.

5.4. Results and discussions

In Section 4 we introduced multiple types of spatial constraints and their corresponding spatially explicit methods. This section discusses their performances of improving over baseline model on reasoning about ternary projective relations. Table 3 lists related sections and abbreviation of different methods.

Table 4 depicts experimental result of relation and place predictions by using different combinations of our proposed methods. The experiment is conducted on the data set of West2000_20, which has an intermediate size among the three built knowledge graphs (Table 2). From this experiment, we observe that jointly leveraging higher-order relatedness (*Ternary*), mutual exclusivity biased negative sampling (*ME*), as well as the attention module (*Atten*) leads to the best result with *Hits*@1 reaching 0.755 and *precision*@1 hitting 0.407. In comparison, the performance of the baseline model (*Pairwise* + *RND*) on these two metrics is 0.548 and 0.194, respectively. *Our best model achieves about 20% improvement for both relation prediction and place prediction.*

In order to understand distinct roles that different components play in the two tasks, we further perform a series of *ablation studies* (Meyes et al. 2019), in which comparisons are made between models with and without a key component.

5.4.1. Impact of higher-order relatedness computation

To investigate the impact of higher-order relatedness computation in predicting relations and places, we compared it with its counterpart - pairwise relatedness computation (Section 4.2.2). As illustrated on Figure 8, taking into account higher-order relatedness considerably benefits the performance of predicting places. In fact, the improvement reaches 9.7% on average for *precision*@1, with a maximum improvement hitting 16.7% for the best model comparing with its pairwise counterpart. The same pattern can be observed for *precision*@5 and *precision*@10 as well. This supports our hypothesis that the three places involved in a ternary projective relation are simultaneously related. Hence, they should not be decomposed into pairwise relations. In terms of relation predictions, there are remarkable improvements as long as the attention module is applied. However, the improvement becomes generally less significant without using attentions. In fact, the performance occasionally decreases (e.g., the *ME* related methods), even though the decrease is marginal comparing to the

Table 4. Experimental results on the data of West2000_20.

	Relation prediction			Place prediction		
Data: West2000_20	*Hits@1*	*Hits@3*	*MRR*	*Precision@1*	*Precision@5*	*Precision@10*
Pairwise + RND	0.548	0.881	0.704	0.194	0.177	0.155
Pairwise + RND + Atten	0.638	0.903	0.762	0.237	0.208	0.179
Ternary + RND	0.563	0.876	0.711	0.210	0.187	0.163
Ternary + RND + Atten	0.661	0.906	0.774	0.362	0.292	0.230
Pairwise + ME	0.591	0.895	0.730	0.215	0.185	0.161
Pairwise + ME + Atten	0.662	**0.910**	0.776	0.240	0.209	0.178
Ternary + ME	0.588	0.877	0.723	0.230	0.194	0.166
Ternary + ME + Atten	**0.755**	**0.910**	**0.833**	**0.407**	**0.318**	**0.252**
Pairwise + CN	0.620	0.893	0.746	0.211	0.182	0.160
Pairwise + CN + Atten	0.651	0.905	0.769	0.233	0.200	0.171
Ternary + CN	0.625	0.886	0.747	0.244	0.210	0.180
Ternary + CN + Atten	0.739	0.907	0.822	0.368	0.292	0.233
Best improvement	0.207	0.029	0.129	0.213	0.141	0.097

The best models under each metric are highlighted in bold.

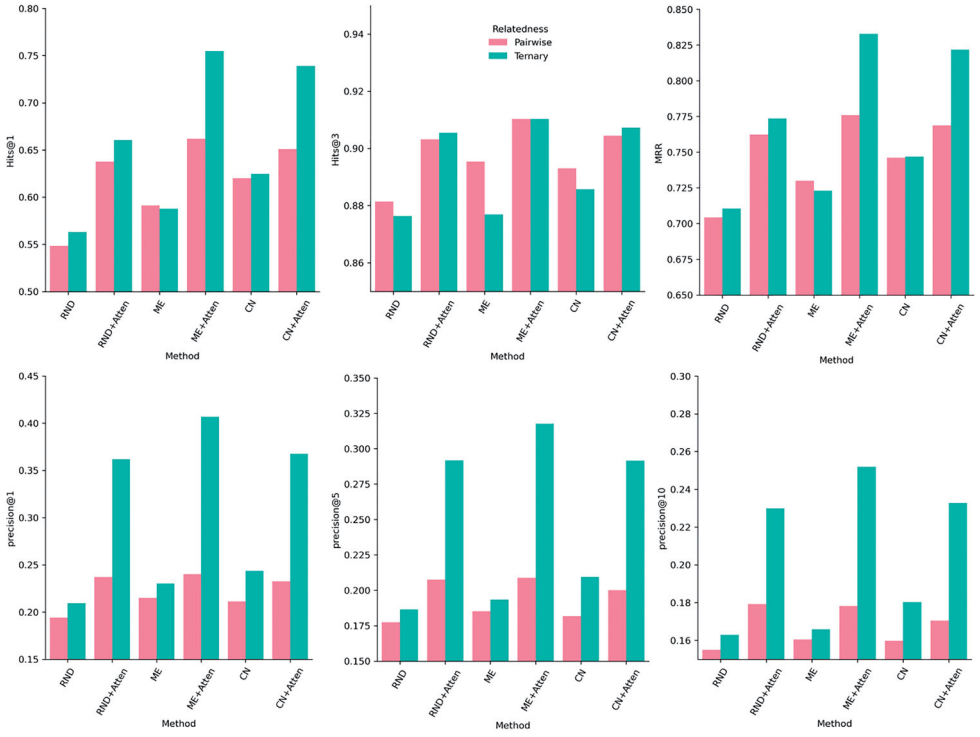

Figure 8. Comparison between pairwise and higher-order relatedness computations for relation prediction (top) and place prediction (bottom).

improvement of leveraging higher-order relatedness (i.e., 0.7% for *MRR*, 1.8% for *Hits*@3, and 0.35% for *Hits*@1). One possible explanation for this effect is that higher-order interaction among roles is implicitly encoded in the training data, where roles from the same family of ternary relations (see R^*) always co-occur; hence, repeatedly

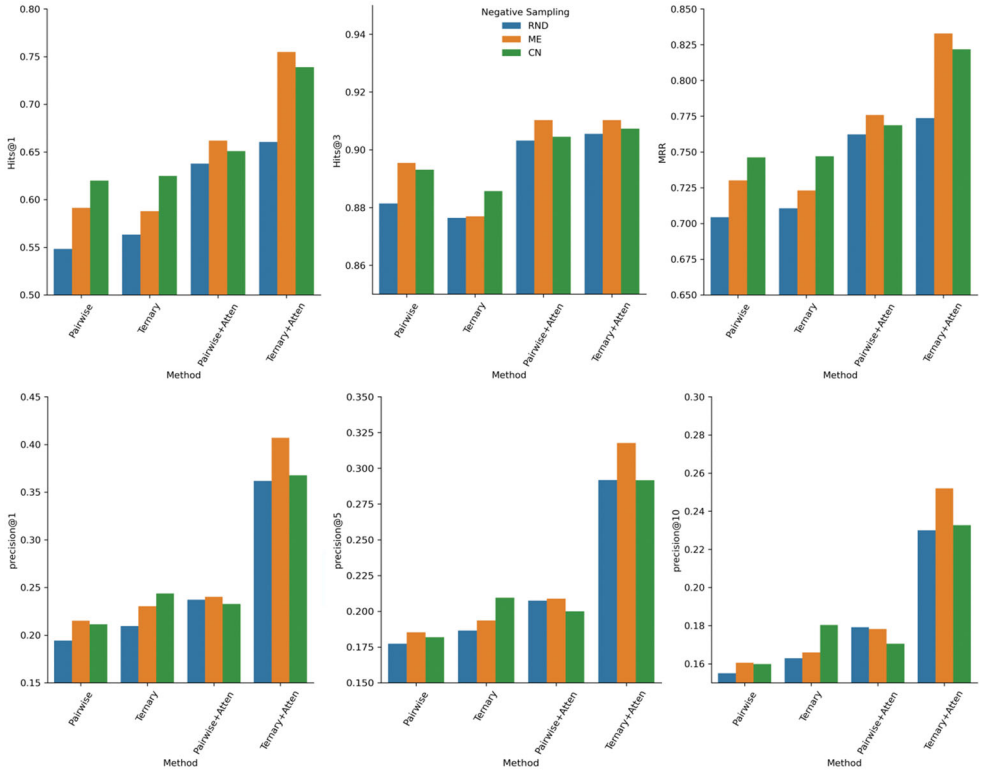

Figure 9. Comparison between random and spatially constrained (i.e., by mutual exclusivity or conceptual neighbourhood) negative samplings for relation prediction (top) and place prediction (bottom).

adding higher-order relatedness of roles into the model would not help substantially increase the performance of relation prediction.

5.4.2. Impact of spatially constrained negative sampling

Likewise, we compare the impact of three different negative sampling strategies across various methods, which are shown in Figure 9. In general, for both relation and place predictions, spatially constrained sampling strategies outperform the baseline strategy. More concretely, when mutual exclusivity constraint is applied, the *MRR* of relation prediction achieves a 2.78% improvement on average across all the methods; while the improvement increases to 3.32% when conceptual neighbourhood constraint is used. In terms of place prediction, the improvement reaches to 2.24 and 1.31% in terms of *precision*@1 for mutual exclusivity constraint and conceptual neighbourhood constraint, respectively. Even though we observe consistent improvement by using any of the two spatial constraints, their differences are subtle in general (i.e., within 1.6% across all metrics of the two tasks), which justifies the discussion in Section 4.3.3 that there are remarkable overlaps between these two spatially constrained negative sampling strategies.

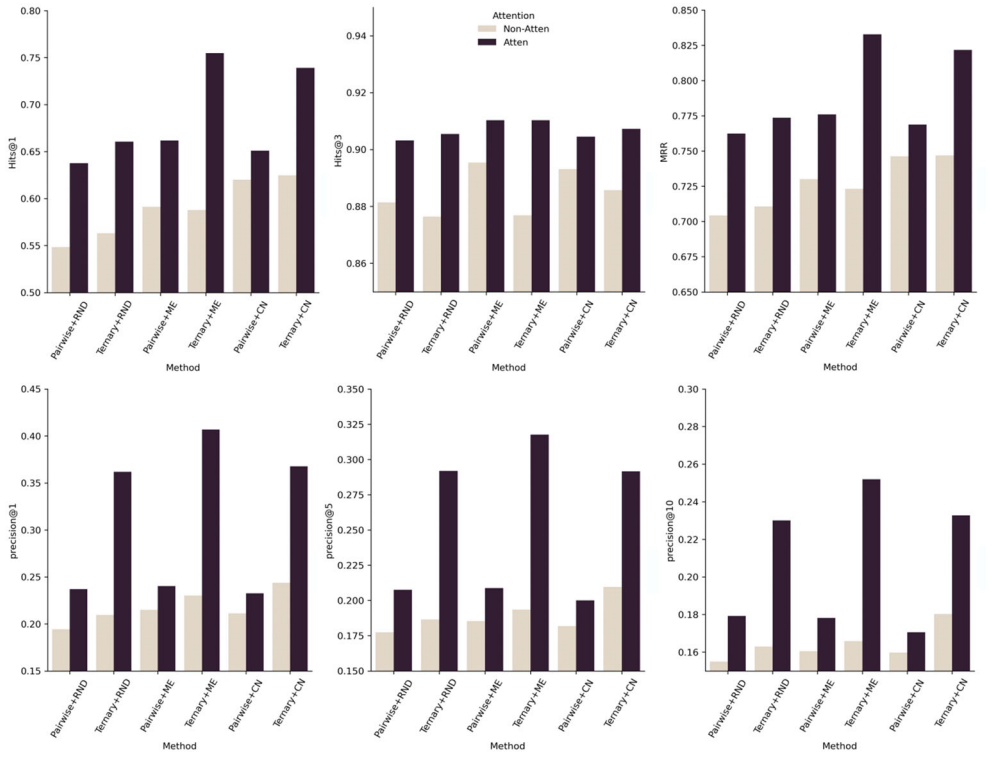

Figure 10. Comparison between attention and non-attention models for relation prediction (top) and place prediction (bottom).

5.4.3. Impact of attention

According to Figure 10, we observe that leveraging the attention module constantly outperforms its non-attention counterpart, where the min-pooling is used to evaluate the overall relatedness of role-value pairs. Different metrics under the two tasks benefit from this method in various strengths. For instance, the average improvement of *Hits*@1 of relation predictions reaches 9.49%, while it is only 2.20% and 6.23% for *Hits*@3 and *MRR*, respectively. One reason underlying this observation is that *Hits*@1 (the top one in the ranked predictions has to be correct) is a stricter metric compared to *Hits*@3 (the correct one exist in the top three). Simpler methods can therefore achieve a comparatively greater performance on *Hits*@3 (already around 0.9) than *Hits*@1 (only around 0.6). Consequently, there is more potential to improve *Hits*@1 than *Hits*@3 using new methods. Besides, it shows that when coupled with higher-order relatedness (*Ternary*), the attention module (*Atten*) significantly fulfills its potential in predicting places (see the bottom of Figure 10). Specifically, the average improvement when *Atten* is used together with *Ternary* are 11.32%, 7.79%, and 5.14% for *precision*@1, *precision*@5, and *precision*@10, respectively. On the contrary, it is 2.23%, 1.80%, and 1.31%, respectively, when *Atten* is coupled with *Pairwise*-based methods.

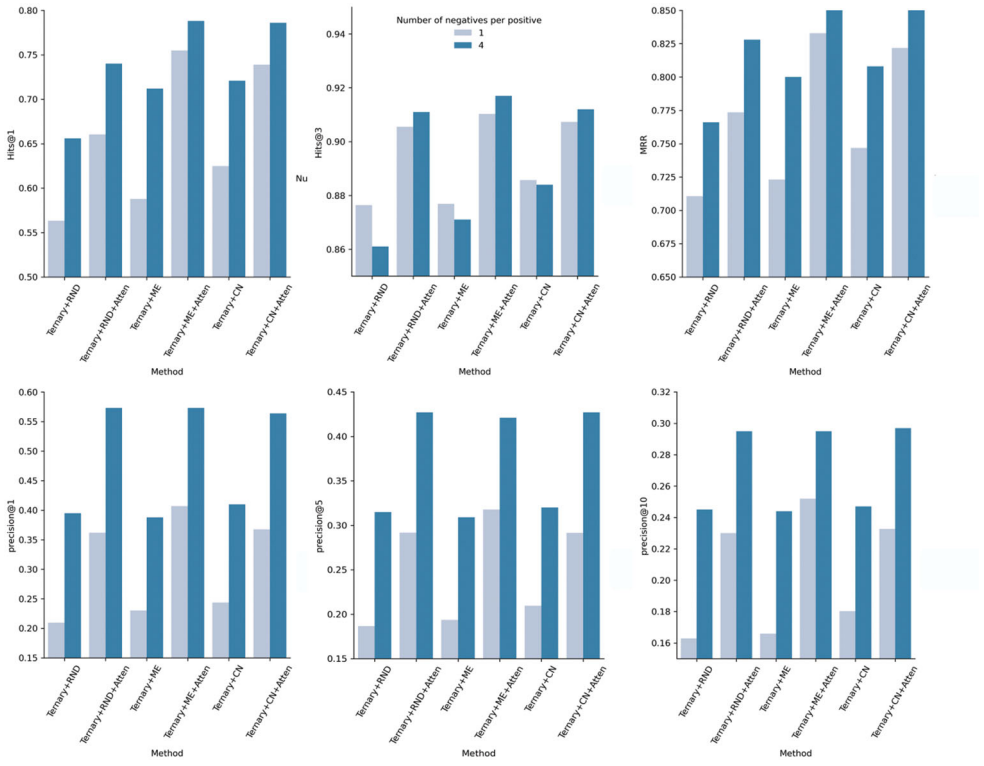

Figure 11. Comparison between one and four negative samples per positive for relation prediction (top) and place prediction (bottom). The experiment is conducted on ternary-related methods.

5.4.4. Impact of negative sample size

So far, all tested methods sample one negative based on one positive relation statement. As discussed in Sections 4.3.2 and 4.3.3, however, one positive statement is able to generate multiple *hard* negatives. Therefore, this experiment investigates the impact of negative sample size on reasoning over ternary projective relations. Specifically, we increased the number of negative samples per positive from one to four, which is the largest number of negatives that both *ME* and *CN* can produce. The complete experimental results are listed in Appendix C and comparisons across *Ternary*-related methods are depicted in Figure 11. On average, the improvements are around 5.06%, 7.88%, and 0.67%, for *MRR*, *Hits*@1, and *Hits*@3, respectively. Compared with relation prediction, we observed even greater benefits of applying a larger negative sample size for place prediction. Namely, *precision*@1, *precision*@5, and *precision*@10 improved about 18.05%, 12.14%, and 6.65%, respectively. This achievement can be attributed to the fact that having more negative samples enables the model to better capture the complex interaction among places because each place, in relation with others in either positive or negative way, will have a higher frequency (four times over the baseline) to be observed by the model.

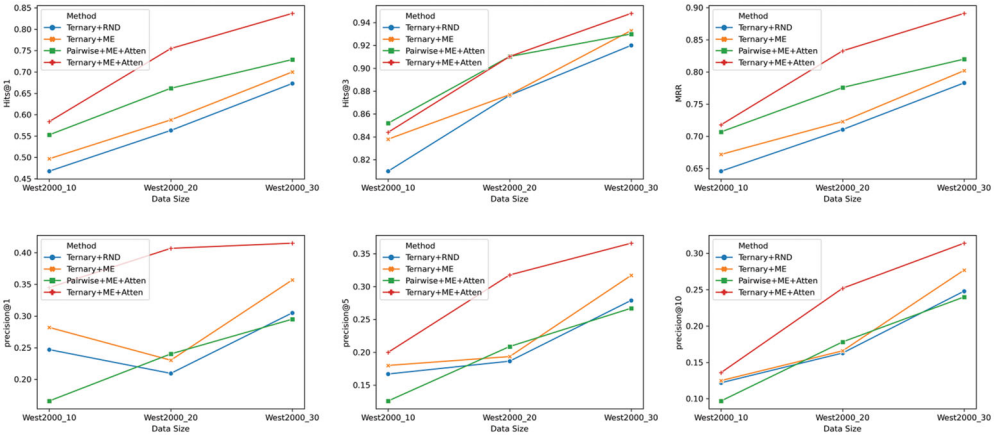

Figure 12. Comparison among three different graph densities on selected methods.

5.4.5. Impact of knowledge graph density

In this experiment, we hypothesize that feeding a denser geospatial knowledge graph into the model would improve the prediction. To test it, we experiment on the impact of graph densities in reasoning about ternary projective relations. Figure 12 illustrates the trend of different methods' performances working on the three graphs listed in Table 2 (results for dataset *West2000_10* and *West2000_30* are detailed in Appendix D). We specifically select four methods for the comparison by removing key components (i.e. higher-order relatedness, spatially constrained negative sampling, and the attention module) one by one from the best method (i.e. *Ternary* + *ME* + *Atten*). As Figure 12 shows, when increasing the number of sampled relation statements in a graph, the six metrics are notably improving, with only exceptions for *precision*@1 of methods *Ternary* + *ME* and *Ternary* + *RND*. This exceptional drop of *precision*@1, as well as the relatively marginal increase of *precision*@5 and *precision*@10 from *West2000_10* to *West2000_20*, are only observed for the two methods where no attention module is leveraged. Therefore, it appears that when the number of relation statements in a geospatial knowledge graph is small, applying attention module becomes more promising in predicting places than simply adding more samples into the training data. This observation can be explained by the fact that the attention module preserves all computed relatedness information among role-value pairs while its counterpart - min-pooling - simply dismisses much of them by only keeping the minimum ones (see Section 4.3.4 for a detailed discussion).

6. Conclusion and future work

Knowledge graphs and graph-based representations more broadly are powerful means to publish geospatial information. However, both the existing geospatial knowledge graphs and the ongoing related research are mainly focusing on modeling binary spatial relations such as *part of* and *overlap*. Studies on representing higher-order relations are rare. This work, particularly focusing on ternary projective relations, presented two

ontology design patterns to represent higher-order qualitative spatial relations. We further developed a neural network architecture to predict the relation among three places, as well as to infer the missing place that is in a relation with the other two. Without being satisfied with the performance of a totally data-driven method, we further explored whether explicitly informing the model with simple rules (i.e., spatial constraints) that human frequently use in reasoning would improve the model performance. Three types of spatial constraints, together with multiple techniques, were identified to transfer the neural network architecture to be spatially explicit: (1) ternary projective relations are *strictly higher-ordered*, (2) *mutually exclusive*, and (3) each relation has a *conceptual neighbourhood*. Three geospatial knowledge graphs with various densities were subsequently created and experimented to evaluate proposed methods, whose results validated artificial neural network's capability of spatial reasoning on ternary relations and demonstrated remarkable benefits of injecting geospatial domain knowledge into the method, even if they are rather simple.

This work brings up several interesting insights that are worth addressing in the future. First of all, what is the "sweet" point between feeding more data into the model and injecting more spatial theories such as the applied spatial constraints? Our experiments only illustrate that reasoning performance can improve by either adding more relations into the graph or by introducing more spatial constraints into the model. Complex interactions between these two, however, remain to be investigated. We conjecture that given a dense enough graph, adding more domain knowledge will not further advance the performance as the model can comprehensively learn those spatial theory-deduced constraints from the graph. Nevertheless, the quality of the training graph would play a key role as well but how and to what degree are to be explored in the future. Secondly, it would be worth researching on whether the proposed neural model can produce new knowledge related to qualitative spatial reasoning. For example, can a neural architecture produce the composition table for ternary projective relations purely based on training data? Next, how do ternary projective relations, together with the binary ones, advance downstream tasks such as geospatial question answering and place summarization is still an open question. Last but not least, a systematic comparison between classic symbolic approaches and our proposed subsymbolic approaches on reasoning over higher-order qualitative spatial relations is necessary in the future, especially from perspectives of addressing aforementioned downstream applications.

Notes

1. https://www.w3.org/RDF/.
2. Note that the straight-line based segmentation in Figure 1 is replaced by a path-based segmentation in order to consider the complexity of real world data when generating relation statements for experiments in this paper. More details will be discussed in Section 5.1.
3. Strictly speaking, the figure includes line segments but as these also have a spatial extent in geographic space we will consider them as regions here.
4. https://wiki.dbpedia.org/projects/dbpedia-places.
5. https://www.wikipedia.org/.
6. https://developers.google.com/maps/documentation/directions/overview.
7. https://docs.ray.io/en/latest/tune/index.html.

Data and codes availability statement

The data and codes that support the findings of this study are available with the identifier(s) at the link https://doi.org/10.6084/m9.figshare.13350737.v2.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work is funded by the National Science Foundation's Convergence Accelerator Program under [Grant No. 1936677 and No. 2033521].

Notes on contributors

Rui Zhu is a Postdoctoral Scholar at the Center for Spatial Studies, University of California, Santa Barbara. He obtained a PhD in Geography from the University of California, Santa Barbara. He also holds a master degree in Information Sciences from the University of Pittsburgh and a bachelor degree in Information Management and Information Systems from Shanxi University of Fiance and Economics. Rui has expertise in spatial statistics, geospatial semantics, knowledge graphs, and GeoAI. His work has been applied to urban planning, global health, environmental intelligence, as well as humanitarian relief.

Krzysztof Janowicz is a Professor for Geoinformatics at the University of California, Santa Barbara and director of the Center for Spatial Studies. He is also a Professor at the Department of Geography and Regional Research, University of Vienna. His research focuses on how humans conceptualize the space around them based on their behavior, focusing particularly on regional and cultural differences with the ultimate goal of assisting machines to better understand the information needs of an increasingly diverse user base. Janowicz's expertise is in knowledge representation and reasoning as they apply to spatial and geographic data, e.g. in the form of knowledge graphs.

Ling Cai is a PhD Candidate at the Space and Time Knowledge Organization Lab, Department of Geography, University of California, Santa Barbara. She obtained her M.S. degree in Geographical Information Science from Chinese Academy of Sciences and B.S. degree from Wuhan University. Her research interests include qualitative spatial temporal reasoning, temporal knowledge graph, neuro-symbolic AI, and urban computing.

Gengchen Mai is a Postdoctoral Scholar at the Stanford AI Lab, Department of Computer Science, Stanford University. He is also affiliated with the Stanford Sustainability and AI Lab. He obtained his PhD degree in Geography from the Department of Geography, University of California, Santa Barbara in 2021 and B.S. degree in GIS from Wuhan University in 2015. His research mainly focuses on spatially-explicit machine learning, geospatial knowledge graph, and geographic question answering.

References

Adams, B., and Janowicz, K., 2011. Constructing geo-ontologies by reification of observation data. *In*: *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, GIS '11, Chicago. New York: ACM, 309–318.

Bahdanau, D., Cho, K., and Bengio, Y., 2014. Neural machine translation by jointly learning to align and translate. *In: 3rd International Conference on Learning Representations, ICLR 2015*, 7–9 May 2015, Conference Track Proceedings, San Diego, CA, USA.

Beall, J., *et al.*, 2012. On the ternary relation and conditionality. *Journal of Philosophical Logic*, 41 (3), 595–612.

Bennett, B., Mallenby, D., and Third, A., 2008. An ontology for grounding vague geographic terms. *In*: C. Eschenbach and M. Gruninger, eds. *Proceedings of the 5th International Conference on Formal Ontology in Information Systems*. Saarbrücken: IOS-Press, 280–293.

Bergstra, J., and Bengio, Y., 2012. Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13, 281–305.

Billen, R., and Clementini, E., 2004. A model for ternary projective relations between regions. *In*: E. Bertino, S. Christodoulakis, D. Plexousakis, V. Christophides, M. Koubarakis, K. Bohm, E. Ferrari, eds. *International Conference on Extending Database Technology, EDBT 2004*. LNCS, vol. 2992. Heidelberg: Springer, 310–328.

Billen, R., and Clementini, E., 2005. Introducing a reasoning system based on ternary projective relations. *In*: P. Fisher, ed. *Developments In Spatial Data Handling 11th International Symposium on Spatial Data Handling*. Berlin, Heidelberg: Springer, 381–394.

Bloch, I., Colliot, O., and Cesar, R.M., 2006. On the ternary spatial relation between. *IEEE Transactions on Systems, Man, and Cybernetics. Part B, Cybernetics*, 36 (2), 312–327.

Bobadilla, J., *et al.*, 2013. Recommender systems survey. *Knowledge-Based Systems*, 46, 109–132.

Bordes, A., *et al.*, 2013. Translating embeddings for modeling multi-relational data. *Advances in Neural Information Processing Systems 26*, December 2013. Red Hook, NY, USA: Curran, 2787–2795.

Calegari, R., Ciatto, G., and Omicini, A., 2020. On the integration of symbolic and sub-symbolic techniques for Xai: a survey. *Intelligenza Artificiale*, 14 (1), 7–32.

Chen, H., *et al.*, 2019. Touchdown: Natural language navigation and spatial reasoning in visual street environments. *In*: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019, 12538–12547.

Chen, H., *et al.*, 2018. A graph database model for knowledge extracted from place descriptions. *ISPRS International Journal of Geo-Information*, 7 (6), 221.

Claramunt, C., 2020. Ontologies for geospatial information: progress and challenges ahead. *Journal of Spatial Information Science*, 2020 (20), 35–41.

Clementini, E., 2019. A conceptual framework for modelling spatial relations. *Information Technology and Control*, 48 (1), 5–17.

Clementini, E., and Billen, R., 2006. Modeling and computing ternary projective relations between regions. *IEEE Transactions on Knowledge and Data Engineering*, 18 (6), 799–814.

Clementini, E., and Di Felice, P., 1995. A comparison of methods for representing topological relationships. *Information Sciences – Applications*, 3 (3), 149–178.

Clementini, E., Di Felice, P., and Hernández, D., 1997. Qualitative representation of positional information. *Artificial Intelligence*, 95 (2), 317–356.

Clementini, E., *et al.*, 2010. A reasoning system of ternary projective relations. *IEEE Transactions on Knowledge and Data Engineering*, 22 (2), 161–178.

Dube, M.P., and Egenhofer, M.J., 2014. Surrounds in partitions. *In*: *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, *SIGSPATIAL '14*, 4–7 November. Dallas, TX: ACM, 233–242.

Egenhofer, M.J., 1989. A formal definition of binary topological relationships. *In*: W. Litwin and H.-J. Schek, eds. *Third International Conference on Foundations of Data Organization and Algorithms*, Paris, France, Lecture Notes in Computer Science, vol. 367. New York: Springer-Verlag, 457–472.

Egenhofer, M.J., 1991. Reasoning about binary topological relations. *In*: O. Gunther and H. J. Schek, eds. *Proceedings of the Second Symposium on Large Spatial Databases, SSD'91*, Zurich, Switzerland, Lecture Notes in Computer Science, vol. 525. Springer. 141–160.

Egenhofer, M.J., and Franzosa, R.D., 1991. Point-set topological spatial relations. *International Journal of Geographical Information Systems*, 5 (2), 161–174.

Egenhofer, M.J., and Mark, D.M., 1995a. Modelling conceptual neighbourhoods of topological line-region relations. *International Journal of Geographical Information Systems*, 9 (5), 555–565.

Egenhofer, M.J., and Mark, D.M., 1995b. Naive geography. *In*: A. Frank & W. Kuhn, eds. *Spatial Information Theory: A Theoretical Basis for GIS, COSITí 95*, Lecture Notes in Computer Science, vol. 988. Berlin: Springer, 1–15.

Frank, A.U., 1992. Qualitative spatial reasoning about distances and directions in geographic space. *Journal of Visual Languages & Computing*, 3 (4), 343–371.

Freksa, C., 1991. Qualitative spatial reasoning. *In*: D. M. Mark, and A. Frank, eds. *Cognitive and Linguistic Aspects of Geographic Space*. Dordrecht: Springer, 361–372.

Freksa, C., 2020. Beyond spatial reasoning: challenges for ecological problem solving. *Journal of Spatial Information Science*, *2020*, 43–49.

Freksa, C., and Kreutzmann, A., 2016. *Neighborhood, conceptual*. *In: International Encyclopedia of Geography: People, the Earth, Environment and Technology*. Hoboken: Wiley-Blackwell, 1–12.

Freksa, C., van de Ven, J., and Wolter, D., 2018. Formal representation of qualitative direction. *International Journal of Geographical Information Science*, 32 (12), 2514–2534.

Freksa, C., *et al.*, 1991., Conceptual neighborhood and its role in temporal and spatial reasoning. In: M. Singh and L. Trave-Massuves, eds. *Decision support systems and qualitative reasoning*. Amsterdam: North-Holland, 181–187.

Guan, S., *et al.*, 2020. NeuInfer: Knowledge inference on N-ary facts. *In*: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Online. Association for Computational Linguistics, 6141–6151.

Hernandez, D., Clementini, E., and Di Felice, P., 1995. Qualitative distances. In: A. Frank and W. Kuhn, eds. *Spatial Information Theory: A Theoretical Basis for GIS*, Lecture Notes in Computer Science, vol. 988. Berlin: Springer-Verlag, 45–58.

Hitzler, P., *et al.*, 2009. *Foundations of semantic web technologies*. Chapman and Hall/CRC.

Hitzler, P., and Shimizu, C., 2018. Modular ontologies as a bridge between human conceptualization and data. *In*: P. Chapman, D. Endres, N. Pernelle, eds. *Proceedings of the 23rd International Conference on Conceptual Structures, ICCS 2018*, Lecture Notes in Computer Science, vol. 10872. Springer, 3–6.

Huang, Z., *et al.*, 2019. Geosqa: a benchmark for scenario-based question answering in the geography domain at high school level. *In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, 5866–5871.

Hudson, D.A., and Manning, C.D., 2019. Gqa: a new dataset for real-world visual reasoning and compositional question answering. *In*: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Long Beach, CA, USA, 6700–6709.

Janner, M., Narasimhan, K., and Barzilay, R., 2018. Representation learning for grounded spatial reasoning. *Transactions of the Association for Computational Linguistics*, 6, 49–61.

Janowicz, K., *et al.*, 2020. Geoai: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science*, 34 (4), 625–636.

Ji, S., *et al.*, 2020. A survey on knowledge graphs: representation, acquisition and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33 (2), 494–514.

Kingma, D.P., and Ba, J., 2015. Adam: a method for stochastic optimization. *In: The International Conference on Learning Representations (ICLR)*. San Diego, CA, USA, 1–15.

Klippel, A., *et al.*, 2013. The Egenhofer–Cohn hypothesis: Or, topological relativity? *In*: M. Raubal, A.U. Frank, and D.M. Mark, eds. *Cognitive and Linguistic Aspects of Geographic Space*, Lecture Notes in Geoinformation and Cartography. Berlin: Springer, 195–215.

Krishnaswamy, N., Friedman, S., and Pustejovsky, J., 2019. Combining deep learning and qualitative spatial reasoning to learn complex structures from sparse examples with noise. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 2911–2918.

Kuhn, W., Kauppinen, T., and Janowicz, K., 2014. Linked data-a paradigm shift for geographic information science. *In*: M. Duckham, E. Pebesma, K. Stewart, A. U. Frank, eds. *Geographic Information Science, GIScience 2014*, Lecture Notes in Computer Science, vol. 8728. Cham: Springer. 173–186.

Landsiedel, C., *et al.*, 2017. A review of spatial reasoning and interaction for real-world robotics. *Advanced Robotics*, 31 (5), 222–242.

Li, L., *et al.*, 2018. A system for massively parallel hyperparameter tuning. arXiv preprint arXiv: 1810.05934.

Lynch, K., 1960. *The image of the city*. Vol. 11. Cambridge, MA: MIT Press.

Mai, G., *et al.*, 2020a. Se-kge: a location-aware knowledge graph embedding model for geographic question answering and spatial semantic lifting. *Transactions in GIS*, 24 (3), 623–655.

Mai, G., *et al.*, 2020b. Multi-scale representation learning for spatial feature distributions using grid cells. In: *Proceedings of the Eighth International Conference on Learning Representations, ICLR 2020*, 26–30 April 2020. Addis Ababa, Ethiopia.

Mai, G., *et al.*, 2019. Relaxing unanswerable geographic questions using a spatially explicit knowledge graph embedding model. *In*: P. Kyriakidis, D. Hadjimitsis, D. Skarlatos, A. Mansourian, eds. *Geospatial Technologies for Local and Regional Development, AGILE 2019*, Lecture Notes in Geoinformation and Cartography. Cham: Springer, 21–39.

Mark, D.M., *et al.*, 1999. Cognitive models of geographical space. *International Journal of Geographical Information Science*, 13 (8), 747–774.

Meyes, R., *et al.*, 2019. Ablation studies in artificial neural networks. CoRR. Retrieved from arXiv: abs/1901.08644.

Minsky, M.L., 1991. Logical versus analogical or symbolic versus connectionist or neat versus scruffy. *AI Magazine*, 12, 34–34.

Mirzaee, R., *et al.*, 2021. Spartqa: a textual question answering benchmark for spatial reasoning. *In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.

Mossakowski, T., and Moratz, R., 2012. Qualitative reasoning about relative direction of oriented points. *Artificial Intelligence*, 180, 34–45.

Nair, V., and Hinton, G.E., 2010. Rectified linear units improve restricted boltzmann machines. *In*: *Proceedings of the 27th International Conference on Machine Learning, ICML-10*, 21–24 June, Haifa, Israel, 807–814.

Noy, N., Rector, A., 2006. Defining n-ary relations on the semantic web. Working group note, W3C, April 2006. Available from: http://www.w3.org/TR/swbp-n-aryRelations/

Peng, B., *et al.*, 2015. Towards neural network-based reasoning. Available from: https://arxiv.org/abs/1508.05508

Peyre, J., *et al.*, 2019. Detecting unseen visual relations using analogies. *In*: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 1981–1990.

Qiu, P., *et al.*, 2019. Knowledge embedding with geospatial distance restriction for geographic knowledge graph completion. *ISPRS International Journal of Geo-Information*, 8 (6), 254.

Regalia, B., Janowicz, K., and McKenzie, G., 2019. Computing and querying strict, approximate, and metrically refined topological relations in linked geographic data. *Transactions in GIS*, 23 (3), 601–619.

Renz, J., 2002. Computational properties of RCC-8. *In*: J. Renz, eds. *Qualitative Spatial Reasoning with Topological inFormation*, Lecture Notes in Computer Science, vol. 2293. Berlin, Heidelberg: Springer.

Russell, S.J., and Norvig, P., 2002. Artificial intelligence: a modern approach. New Jersey: Prentice Hall, Upper Saddle River.

Santoro, A., *et al.*, 2017. A simple neural network module for relational reasoning. *In*: I. Guyon, *et al.*, eds. *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 4967–4976.

Scheider, S., *et al.*, 2020. Ontology of core concept data types for answering geo-analytical questions. *Journal of Spatial Information Science*, 2020, 167–201.

Vasardani, M., *et al.*, 2013. From descriptions to depictions: a conceptual framework. *In*: T. Tenbrink, J. Stell, A. Galton, Z. Wood, eds. *International Conference on Spatial Information Theory, COSIT 2013*, Lecture Notes in Computer Science, vol. 8116. Cham: Springer, 299–319.

Wang, Q., *et al.*, 2017. Knowledge graph embedding: a survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29 (12), 2724–2743.

Worboys, M.F., 2001. Nearness relations in environmental space. *International Journal of Geographical Information Science*, 15 (7), 633–651.

Yan, B., *et al.*, 2019. A spatially explicit reinforcement learning model for geographic knowledge graph summarization. *Transactions in GIS*, 23 (3), 620–640.

Yang, B., *et al.*, 2015. Embedding entities and relations for learning and inference in knowledge bases. *In: Proceedings of the International Conference on Learning Representations, ICLR 2015*.

Zhu, R., *et al.*, 2016. Spatial signatures for geographic feature types: examining gazetteer ontologies using spatial statistics. *Transactions in GIS*, 20 (3), 333–355.

Zhu, R., Janowicz, K., and Mai, G., 2019. Making direction a first-class citizen of tobler's first law of geography. *Transactions in GIS*, 23 (3), 398–416.

Zhu, R., Kyriakidis, P.C., and Janowicz, K., 2017. Beyond pairs: generalizing the geo-dipole for quantifying spatial patterns in geographic fields. *In*: A. Bregt, T. Sarjakoski, R. van Lammeren, F. Rip, eds. *Societal Geo-innovation, AGILE 2017*, Lecture Notes in Geoinformation and Cartography. Cham: Springer, 331–348.

Appendix A. Computation of the five buffer zones for ternary projective relations

There are multiple steps involved in determining ternary projective relations among three places. First, we regard the two places in each sampled pair as either the origin or destination. Secondly, we utilize the Google Directions API[6] to obtain the route between the origin and destination places. Since no city will exactly lie on the route so as to be considered as *between* the origin and destination, we apply a buffer zone on the route (red line in Figure 7) and whichever city (green points in Figure 7) falls into the zone will be taken as in *between*. Furthermore, we extend the buffer zone with larger distances (zone with purple boundary in Figure 7), which results into the *left* and *right* zones (orange and pink points in Figure 7). Next, for both *before* and *after*, we first build a rectangle buffer zone for each, which is vertical to the first (for *before* zone) or the last (for *after* zone) segment of the route. Lastly, the *before* and *after* zones are determined by the intersection of green rectangle zones with the purple one.

Appendix B. Hyperparameter tuning details

In hyperparameter tuning, we use *precision*@1 of place prediction as the main evaluation metric because its performance is generally worse than relation predictions. Namely, the set of hyperparameters that achieves the best result in terms of *precision*@1 will be kept to further train the model. In addition, to improve the efficiency of the optimization process, we apply a scheduler to stop the searching early. Specific algorithm used in this work is ASHA (Li et al. 2018). The scheduler terminates those trials that cannot keep improving model performance in terms of selected metrics, and allocates more resource to those trials that are more promising. The maximum number of trials is set to 10 and the maximum number of epochs for each trial is set to 100 (the number of epochs for training the model is 200). Such a hyperparameter tuning process is implemented using the library Ray Tune[7].

Appendix C. Experiment results about the impact of negative sampling size.

Data: West2000_20 (four negative samples per positive)	Relation prediction			Place prediction		
	Hits@1	*Hits@3*	*MRR*	*Precision@1*	*Precision@5*	*Precision@10*
Pairwise + RND	0.664	0.895	0.774	0.322	0.275	0.225
Pairwise + RND + Atten	0.692	0.903	0.795	0.374	0.317	0.250
Ternary + RND	0.656	0.861	0.766	0.396	0.315	0.245
Ternary + RND + Atten	0.740	0.911	0.828	**0.573**	**0.427**	0.295
Pairwise + ME	0.717	0.890	0.809	0.332	0.277	0.225
Pariwise + ME + Atten	0.733	0.911	0.822	0.374	0.309	0.245
Ternary + ME	0.712	0.871	0.800	0.388	0.309	0.244
Ternary + ME + Atten	**0.788**	**0.917**	**0.856**	**0.573**	0.421	0.295
Pairwise + CN	0.688	0.884	0.786	0.333	0.277	0.226
Pairwise + CN + Atten	0.744	0.907	0.828	0.377	0.309	0.247
Ternary + CN	0.721	0.884	0.808	0.410	0.320	0.247
Ternary + CN + Atten	0.786	0.912	0.854	0.564	**0.427**	**0.297**

The best models under each metric are highlighted in bold.

Appendix D. Experiment results about the impact of knowledge graph density.

Data: West2000_10 (one negative sample per positive)	Relation prediction			Place prediction		
	Hits@1	*Hits@3*	*MRR*	*Precision@1*	*Precision@5*	*Precision@10*
Pairwise + RND	0.473	0.817	0.653	0.193	0.149	0.112
Pairwise + RND + Atten	0.532	0.833	0.692	0.172	0.131	0.102
Ternary + RND	0.468	0.810	0.646	0.247	0.167	0.122
Ternary + RND + Atten	0.548	0.837	0.695	0.312	0.198	**0.137**
Pairwise + ME	0.484	0.833	0.664	0.206	0.155	0.155
Pariwise + ME + Atten	0.553	**0.852**	0.707	0.166	0.126	0.097
Ternary + ME	0.497	0.838	0.672	0.282	0.180	0.125
Ternary + ME + Atten	**0.584**	0.844	**0.718**	**0.345**	**0.200**	0.136
Data: West2000_30 (one negative sample per positive)	Relation prediction			Place prediction		
	Hits@1	*Hits@3*	*MRR*	*Precision@1*	*Precision@5*	*Precision@10*
Pairwise + RND	0.661	0.918	0.779	0.261	0.250	0.226
Pairwise + RND + Atten	0.705	0.935	0.806	0.268	0.253	0.230
Ternary + RND	0.673	0.920	0.783	0.305	0.279	0.248
Ternary + RND + Atten	0.743	0.935	0.830	0.313	0.284	0.253
Pairwise + ME	0.680	0.928	0.788	0.274	0.255	0.232
Pariwise + ME + Atten	0.729	0.930	0.820	0.295	0.267	0.240
Ternary + ME	0.700	0.933	0.802	0.357	0.317	0.277
Ternary + ME + Atten	**0.837**	**0.948**	**0.891**	**0.415**	**0.366**	**0.314**

The best models under each metric are highlighted in bold.