ARTICLE TEMPLATE

A Machine Learning based Approach for Generating Point Sketch Maps from Qualitative Directional Information

Zhiguo Long[a,b], Qingqian Li[a,b], Hua Meng[c], and Michael Sioutis[d]

[a] School of Computing and Artificial Intelligence, Southwest Jiaotong University, China
[b] Manufacturing Industry Chains Collaboration and Information Support Technology Key Laboratory of Sichuan Province, China
[c] School of Mathematics, Southwest Jiaotong University, China
[d] LIRMM UMR 5506, University of Montpellier & CNRS, France

ARTICLE HISTORY
Compiled June 13, 2024

ABSTRACT
People often use qualitative relations to describe locations or directional information, especially in written communication, such as "the restaurant is located at the *southeast* corner of the square". However, when a large number of spatial entities are involved, qualitative relations alone are not intuitive enough for people to understand a spatial configuration. In fact, many applications, e.g., pertaining to sharing travel experiences, use *sketch maps*, i.e., maps focusing on the main features of an area whilst abstracting exact scale measurements, to help demonstrate abstract qualitative relations with more intuitive geometric points. Current approaches for generating point sketch maps from qualitative spatial relations require a high level of expertise, face inherent difficulties with efficiently processing large-scale data in bulk, and are vulnerable to inaccurate or conflicting information contained in qualitative data. To address these limitations, by incorporating machine learning techniques, we propose to translate the problem into an optimization problem of data reconstruction, enabling a novel end-to-end approach for generating point sketch maps from qualitative directional relations in bulk. Experiments on real-world datasets show that the proposed approach has very high accuracy and is robust even with a large portion of inaccurate or incomplete information.

KEYWORDS
qualitative directional relations; Laplacian Eigenmaps; point sketch map; visualization

1. Introduction

Information technologies such as social media, cloud computing, and artificial intelligence have greatly improved the capability of people to share, acquire, and analyse map data. Web map services (Wu *et al.* 2011) provide us with a *quantitative* way for describing and understanding a geographical entity. Meanwhile, *qualitative* spatial relations are abundant in web documents (He *et al.* 2015) and social media (Stock *et al.* 2022). For example, in social media, users often use qualitative spatial relations,

---

This is an Accepted Manuscript of an article published by Taylor & Francis in International Journal of Geographical Information Science on 24 May 2024, available at: https://doi.org/10.1080/13658816.2024.2358405

CONTACT Hua Meng. Email: menghua@swjtu.edu.cn

like "on the east" and "to the northwest", to describe directional information. An increasing number of scholars also show interest in such relations and have incorporated them into fields such as geographical information science (Egenhofer and Mark 1995, Worboys and Duckham 2004), qualitative spatial reasoning (Cohn and Renz 2008), computer vision (Zhan *et al.* 2020), computational geometry (Frontiera *et al.* 2008), and biomedical research (Wang *et al.* 2022).

Motivation. While qualitative spatial relations have been extensively exploited in geographical information processing as an abstract semantic layer, most of the existing works mainly directly use the existing qualitative relations or the ones calculated from numerical data. Recent examples of such usage of qualitative spatial relations include: locating POIs from descriptions involving not only place names but also their qualitative spatial relations (Cheng *et al.* 2022), navigating in GPS-denied environments with the help of qualitative place maps (Hua *et al.* 2019), generating place concepts by constructing spatial hierarchical relations from data (Wu *et al.* 2019), and matching place names in different natural language descriptions by considering their qualitative relations (Kim *et al.* 2017).

We note that existing geographical information systems (GISs) and location-based services (LBSs) have more sophisticated techniques to deal with numerical data than with qualitative relations, because the former are generally more informative and easier to process. As a result, there is still a gap of integrating qualitative relations with current numerical methods in the geographical context, that is, only very few works have considered converting qualitative relations into numerical data, so that existing numerical methods can better process such qualitative information. For example, Kim *et al.* (2016) proposed to generate a picture depiction of qualitative relations between places that can facilitate emergency responses, Belouaer *et al.* (2016) discussed generating a map from verbal (qualitative) navigation instructions that can better demonstrate the navigation knowledge, and Chen *et al.* (2018) considered georeferencing a place on the map through its qualitative relations with landmarks.

In fact, we argue that the capability of converting qualitative relations into numerical data can inspire new tools and techniques to analyse and visualize geographical data, and will help systems and services that are based on numerical methods to process qualitative information. For instance, although qualitative spatial relations alone cognitively conform to human users, visual representations such as maps have stronger intuitiveness and can benefit GISs so that they will become more user-friendly. Consider the following example when we communicate with friends about a city:

> There is a large gymnasium g to the north (N) of the city square s. To the west (W) of the city square, there is a famous park p, which is located in the southwest (Sw) of the gymnasium. Additionally, there is a famous tourist attraction t to the northwest (Nw) of the park and to the west (W) of the gymnasium, while the southeast (Se) of the city square is a famous business centre b.

Figure 1(a) lists the qualitative relations between these places, e.g. "$g\ \mathsf{N}\ s$" means that g is to the north of s. This large number of qualitative descriptions and relations are not intuitive enough for people to have an overall picture of the distribution of those places. A map (as shown in Figure 1(b)) that visually demonstrates the qualitative relations would be arguably more intuitive.

Moreover, by converting qualitative relations to numerical data like points, one can utilize advanced numerical methods for data analysis, e.g., performing topological data analysis on the numerical data, as proposed in (Corcoran and Jones 2023).

Target	Relation	Reference
g	N	*s*
p	W	*s*
p	Sw	*g*
t	Nw	*p*
t	W	*g*
b	Se	*s*

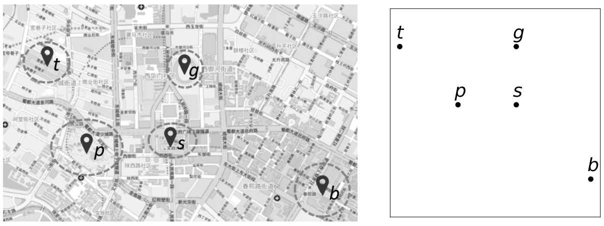

(a) (b) (c)

Figure 1. A motivational example of a point sketch map for qualitative directional relations: (a) Qualitative directional relations between some places (ground truth, input); (b) Visual distribution of these places on a real map; (c) A generated point sketch map for these places (prediction, our output).

This direction can also enhance spatial cognition by providing a more intuitive graphical representation of data than a set of qualitative relations, e.g., GISs can visualize the distribution of historical buildings by taking into account qualitative relations described in historical books. It can also help to enrich the knowledge base of GISs, by generating numerical data that align with the large amount of qualitative data existing in social media but not yet fully utilized by current systems. Other significant implications include geographical registration, location restoration, and correction of historical and modern maps. We will discuss more implications later in Section 6.

State of the art. Research on the quantitative conversion of qualitative spatial relations has achieved some preliminary results, including visualization for location descriptions (Wieczorek *et al.* 2004, Vasardani *et al.* 2013), for path descriptions (Moncla *et al.* 2016), and for scene descriptions (Kim *et al.* 2016). However, some of these studies require quantitative spatial data to be known or partially known, while some require qualitative descriptions of both distances and directions. Most of these studies rely on complex computational algorithms designed with expert knowledge to obtain sketch maps. As a result, the effectiveness and accuracy of these algorithms face great challenges when some information is fuzzy, inaccurate, or even missing. Moreover, these methods mostly deal with only a small number of geographical entities, and there is insufficient research on the quantitative conversion of numerous geographical entities in bulk.

Main idea. To address the aforementioned issues, this article proposes an end-to-end approach to obtain a *point sketch map*[1] from qualitative directional relations. A point sketch map is a set of points on the plane, where each point represents a geographical entity (e.g., place) and the points are used to demonstrate the qualitative directional relations among these entities. In sum, it utilizes a gradient descent algorithm guided by a novel loss function to continuously optimize point sketch maps; this approach is shown to be able to obtain a sketch map for a large number of geographical entities. In detail, to obtain good point sketch maps, the proposed approach introduces the Laplacian reconstruction loss (Niyogi and Belkin 2003) commonly used in the machine learning field to characterize neighbourhood information, and designs a novel directional loss to characterize qualitative relation information. This approach does

---

[1] The concept of "sketch map" used here is slightly different from the common understanding in other fields, such as spatial cognition, where it usually refers to hand-drawn intuitive pictures to depict ideas. Here, we adopt this term for a map of points to capture the intuitiveness of the map and to emphasize that the map aims to illustrate the qualitative relations between entities, rather than exact coordinates as with usual maps.

not require strong expert knowledge, nor the design of complex computational formulas. Furthermore, it has a certain degree of fault tolerance and can deal with a large number of geographical entities simultaneously. As an illustration, for the qualitative directional relations shown in Figure 1(a), the proposed approach can obtain a point sketch map as shown in Figure 1(c), which is very close to the true distribution in Figure 1(b). In fact, with respect to fault tolerance, although the angles are not exactly the same for these two figures, the qualitative relations are the same, e.g., the angle between the vector $\vec{sb}$ (from point s to point b) and the x-axis in Figure 1(b) is slightly smaller than 45°, and in Figure 1(c) it is almost 45° resulting from the numerical calculation of the approach, yet they maintain the same directional relation Se.

It is worth noting that in this article we only consider cardinal directions, but there exist other common types of directional relations, such as relative directions like *front* and *right*. Cardinal directions are a well-known type of qualitative spatial relations between geographical places, and many real-world applications including Wikipedia are interested in this type of relations. Also, for geographical places, cardinal relations can usually be easily obtained without further information other than normal geometries, whereas obtaining relative directions requires the orientations of geographical places which are usually not known. Nevertheless, the proposed idea might also be useful for other types of relations, e.g., for relative directions, when the orientations of reference points are known, one can consider that there is a cardinal direction system for each reference point; thus, this setting does not limit the potential of the proposed approach to deal with other types of relations.

Contributions. The main contributions of this article are:

(1) An end-to-end machine learning approach is employed to solve the problem of generating point sketch maps from qualitative directional relations.
(2) In addition to the neighbourhood loss function, a novel qualitative directional relation loss function is proposed to better preserve the qualitative directional information.
(3) Experimental results indicate that the proposed method can achieve an average accuracy rate close to 100% (and over 80% when there is up to about 20% of noise), demonstrating that the proposed method provides an effective solution for generating point sketch maps.

The proposed approach can also serve as a general framework of generating geometries of spatial entities in accordance with their qualitative spatial relations, which can be extended to support more complex geometries and other types of relations.

The rest of the article is organized as follows. Section 2 discuses related work. Section 3 introduces preliminary knowledge. Section 4 describes the approach in detail. Section 5 experimentally evaluates the proposed approach, in both a visual way and a metrics way. Section 6 concludes the article and discusses implications, limitations, and future work.

2. Related Work

The study of modelling quantitative information as qualitative directional relations has been a focus of research for decades. For example, Frank (1992) introduced a cone-shaped model and a projection model for directional relations. In the cone-shaped model, each direction is assigned a 45° fan-shaped acceptance region, and in the projection model, the directions on axes correspond to North, South, East, and West, respectively, while the area between two directions correspond to the directional relation between them, e.g., Northeast. Clementini (2013) proposed a qualitative representation framework for directional information and frames of reference, using a 5-intersection model. Du *et al.* (2023) defined a logic system that describes the east, west, and uncertain directions between points, providing a semantic interpretation based on the margin of error and level of indeterminacy. These qualitative models of different granularities and perspectives provide qualitative directional relations with a semantic meaning that mimics human cognition. There are many other directional relation models, and interested readers are referred to (Dylla *et al.* 2017).

Currently, researchers are increasingly interested in the fusion of qualitative directional relations and quantitative information to improve the performance of specific tasks, such as matching points of interest (POIs) (Cheng *et al.* 2022), indoor localization and navigation (Hua *et al.* 2019, Winter *et al.* 2019), and many others. For example, Cheng *et al.* (2022) integrated qualitative directional relations and quantitative distances into spatial reasoning, achieving localization of POIs through address matching and spatial reasoning. They also accelerated directional relation retrieval and improved the efficiency of spatial reasoning for large datasets by implementing compact qualitative direction representations on global equal latitude and longitude grids (Goodchild and Kimerling 2002). Winter *et al.* (2019) discussed infrastructure-independent indoor navigation methods based on smartphones, including qualitative interaction between users and devices, which are slightly less accurate but have the potential for universal application. Hua *et al.* (2019) proposed a localization and navigation method based on qualitative maps that represent important landmarks in the environment and their spatial relations. These qualitative maps are then used for localization and navigation in GPS-weak or GPS-denied environments.

As one of the aspects for the interaction and fusion of qualitative and quantitative directional information, quantitative conversion of qualitative directional relations mainly involves visualizing directional relations described in text, by drawing sketch maps of the distribution of multiple spatial entities, or of the approximate boundary of a single spatial entity. Kim *et al.* (2016) considered to iteratively update sketch maps that represent spatial relations between entities, by designing position update formulas based on qualitative direction and distance information. Since each iteration considers only a single spatial relation, this method requires continuous adjustment of the positions of drawn objects, to avoid conflicts between newly added spatial entities and existing ones. Belouaer *et al.* (2016) proposed a method for extracting and fusing qualitative spatial relations from route descriptions and conducted preliminary tests on visualizations of these relations using a genetic algorithm. However, this method suffers from low efficiency and poor performance. Long *et al.* (2022) approximated unknown boundaries of regions, by quantifying directional relations and distance information through complex formulas. This method also requires known positions of some reference regions. Graullera *et al.* (2006) introduced a map construction method for Aibo robots that can generate a basic topological map from image data collected by the robots. This method is sensitive to noise and requires long computation time and large memory space for processing large-scale maps. Schockaert *et al.* (2011) triangulated the plane and used given qualitative topological relations to generate approximate boundaries of multiple regions based on a genetic algorithm. This method requires elaborate algorithms and has high computational complexity, making it difficult to handle a large amount of spatial relations. As can be deduced from the above, existing studies can only conduct quantitative conversion for a very small number of qualitative spatial relations between geographical entities, and usually require complex calculations and a high level of expert knowledge. They also rely heavily on topological relations and distance information, and have poor fault tolerance for noisy information.

There are various data reconstruction methods in the field of data mining and machine learning, which can efficiently rearrange the spatial distribution of a large number of data points, while maintaining some relations between the original data points (such as similarity relations). For instance, Multidimensional Scaling (Cox and Cox 2008, MDS) is a classical data reconstruction method, which reconstructs data points in a lower dimensional space in such a way that the distance between every pair of points in the lower dimensional space is close to that of the respective points in the original space. Laplacian Eigenmaps (Niyogi and Belkin 2003, LE) is another typical data reconstruction method, which plays an important role in dimensionality reduction and visualization to achieve better clustering effects (Yin and Ma 2019). Compared to MDS, which focuses on the distance for every pair of points, LE is concerned more with local neighbouring information. This method and its variations, such as Locally Linear Embedding (Roweis and Saul 2000, LLE) and Locality Preserving Projection (He and Niyogi 2003, LPP), reconstruct high-dimensional data into low-dimensional data, while attempting to maintain some local similarity between neighbours. By adopting the above idea of data reconstruction, and to address the limitations of existing methods, this article proposes an end-to-end method for quantitative conversion of qualitative relations. The proposed approach does not require a high level of expertise, and only relies on qualitative directional information to generate an intuitive point sketch map for a large number of geographical entities. Specifically, a special loss function is used to guide the quantitative reconstruction of qualitative directional relations, so that the reconstructed points conform to the given relations as much as possible.

3. Preliminaries

3.1. *Cone-Shaped Cardinal Directional Relations*

Given a representative point p (here we use centroid) of a geographical entity, there can be multiple qualitative directional relations of another point q relative to p (i.e., with p as the reference point). Here, we are specifically concerned with the eight *cardinal directional relations* provided by the cone-shaped directional relation model (Frank 1992, Peuquet and Ci-Xiang 1987). The cone-shaped relation model defines eight possible directional relations between a reference point p and a target point q on the plane, i.e., North (N), Northeast (Ne), East (E), Southeast (Se), South (S), Southwest (Sw), West (W), and Northwest (Nw). For these eight directional relations, the model uniformly partitions the plane into eight cone-shaped regions (as shown in Figure 2(a)), known as the "regions of acceptance" for directional relations. If q falls inside the region of acceptance for a directional relation d, then the directional relation of q w.r.t. p is defined to be d, i.e., $(q\ R_{qp}\ p) = (q\ d\ p)$. For example, q falls within the region of acceptance for Ne in Figure 2(a), and $R_{qp} = \mathsf{Ne}$. It is worth noting that the cone-shaped relation model is commonly used but is still just an approximation of human cognition, e.g., it separates the plane into eight crisp regions, which can be error-prone around the borders, i.e., a small change in position can lead to a different relation. In general, this is the case with any qualitative model, as there is some inevitable loss of information when abstracting metric data and forming more coarse-grained descriptions. Also, in this article, the standard 45° angle is chosen for each of the regions of acceptance. This is because, when there is no further information, the eight directions have equal importance, and the regions of acceptance should be evenly distributed. In more general cases, one can use different angles for regions of acceptance, and the essential idea proposed in this article still applies.

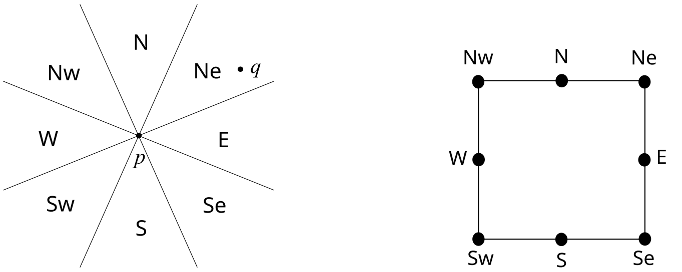

(a) cone-shaped cardinal direction relations (b) conceptual neighbourhood graph

Figure 2. Cone-shaped cardinal direction relations and their conceptual neighbourhood graph.

Some relations are more "conceptually close" than others; intuitively, if we have a reference point p and a target point q ($p \neq q$), and we assume that their directional relation is R, viz., we have $p\ R\ q$, then by continuously changing the positions of p and q, the relation R can directly transition to another relation S. Such S and R are more "conceptually close" than other relations that do not have this property. This leads to the well-known *conceptual neighbourhood graph* (Freksa 1996).

Definition 3.1 (Conceptual Neighbourhood Graph)**.** The conceptual neighbourhood graph for the cone-shaped cardinal directional relations is the undirected graph where: the vertices are the relations, and two relations (vertices) have an edge if their respective regions of acceptance share a line border.

The conceptual neighbourhood graph for the eight directional relations is shown in Figure 2(b). Two relations with an edge are said to be *conceptual neighbours.* For example, N and Ne are conceptual neighbours while N and E are not. Moreover, we can also define the *conceptual neighbourhood distance* (Condotta *et al.* 2008) between two relations based on the conceptual neighbourhood graph.

Definition 3.2 (Conceptual Neighbourhood Distance)**.** For two relations R and S in the conceptual neighbourhood graph G, the *conceptual neighbourhood distance* $d(R, S)$ between R and S is the number of edges in a shortest path between R and S in G.

For example, $d(\mathsf{N}, \mathsf{Ne}) = 1$, $d(\mathsf{N}, \mathsf{E}) = 2$, and $d(\mathsf{N}, \mathsf{S}) = 4$, and the range of the distance between two relations is $[0, 4]$.

Given a set of points on the plane, according to the above relation model, we can compute the qualitative directional relations among these points based on their coordinates. For two points p and q, the directional relation of q relative to p can be calculated as follows:

(1) As in (Deng and Li 2008), we can calculate the angle α_{pq} between the vector

from p to q and the x-axis direction using

$$\alpha_{pq} = \begin{cases} \arctan \frac{\Delta y_{pq}}{\Delta x_{pq}} & (\Delta y_{pq} \geq 0, \Delta x_{pq} > 0) \\ 90^\circ & (\Delta y_{pq} > 0, \Delta x_{pq} = 0) \\ 180^\circ + \arctan \frac{\Delta y_{pq}}{\Delta x_{pq}} & (\Delta x_{pq} < 0) \\ 270^\circ & (\Delta y_{pq} < 0, \Delta x_{pq} = 0) \\ 360^\circ + \arctan \frac{\Delta y_{pq}}{\Delta x_{pq}} & (\Delta y_{pq} < 0, \Delta x_{pq} > 0) \end{cases} \tag{1}$$

where $\Delta y_{_{pq}} = y_q - y_p$ and $\Delta x_{_{pq}} = x_q - x_p$;

(2) The directional relation is then determined by comparing the angle value with the angle range corresponding to each directional relation.

For the example in Figure 2(a), $\alpha_{pq} = 31.65^\circ$ and the angle range of Ne is $[22.5^\circ, 67.5^\circ)$. Therefore, the directional relation R_{qp} of q relative to p is Ne.

3.2. *Laplacian Eigenmaps*

Laplacian Eigenmaps (LE) proposed in (Niyogi and Belkin 2003) is a data reconstruction method targeting to preserve local similarity information of original data. Its essential idea is: if two data instances p_i and p_j (which can be considered as two points in Euclidean space) have high similarity in the original space, then their reconstructed results p_i' and p_j' should also be close in distance, thereby preserving the local similarity information. This idea is implemented in the LE model by designing a specific loss function.

Particularly, let the original data set be represented as $P = \{p_1, p_2, \ldots, p_n\}$, where a data point $p_i \in \mathbb{R}^m$ in P is an m-dimensional row vector. The similarity between p_i and p_j is given by W_{ij}. Suppose the reconstructed data points of p_i and p_j are p_i' and p_j', where $p_i', p_j' \in \mathbb{R}^d$ (d is the target dimensionality). Then, the loss function of LE is:

$$\sum_{i=1}^{n} \sum_{j=1}^{n} W_{ij} \|p_i' - p_j'\|^2. \tag{2}$$

In this way, for a pair of data points p_i and p_j with high similarity (W_{ij} is large), in order to minimize the loss, $\|p_i' - p_j'\|^2$ should be as small as possible, i.e., the reconstructed points p_i' and p_j' should be as close as possible.

Note that if the above model is optimized without constraints, then all p_i' and p_j' will be compressed into a single point, so that the value of the loss function is the minimum 0. To avoid that, LE adds the additional constraint $P'^T DP' = I$, where $P' = (p_1'; \ldots; p_n')$, D is an $n \times n$ diagonal matrix and $D_{ii} = \sum_{j=1}^{n} W_{ij}$. Intuitively, this constraint requires that P' should be roughly orthogonal: each column $(p_{1j}', \ldots, p_{nj}')^T$ of P' is orthogonal with other columns by a weighted factor of D, i.e., $\sum_{i=1}^{n} D_{ii} p_{ij}' p_{ik}' = 0$ $(j \neq k)$, and $\sum_{i=1}^{n} D_{ii} (p_{ij}')^2 = 1$.

The procedure of data reconstruction using LE is as follows.

(1) Obtaining the nearest neighbour structure:
    (a) ε-neighbour: if $\|p_i - p_j\| < \varepsilon$, then p_j is a neighbour of p_i.
    (b) k-nearest neighbour: if p_i is one of the k-nearest neighbours of p_j, or p_j is one of the k-nearest neighbours of p_i, then p_j is a neighbour of p_i.

(2) Computing the similarity between two points that are neighbours of each other: The commonly used similarity measure can be obtained by calculating the Gaussian kernel function $W_{ij} = \exp(-\frac{\|p_i - p_j\|^2}{\sigma^2})$, where σ controls the width of the Gaussian function and determines the range of distance values that should be significant to the similarity values (see (Romeny 2003) and also the supplementary material for more details).

(3) Solving the optimization problem in Equation 2 with the constraint $P'^T DP' = I$:
    (a) This problem can be reformulated as a Lagrangian function, whose solution is equivalent to a solution of the generalized eigenvalue problem $Ly = \lambda Dy$ (Niyogi and Belkin 2003); here, L is the Laplacian matrix $L = D - W$;
    (b) The solutions are (usually) formed by the corresponding eigenvectors of the first m' non-zero smallest generalized eigenvalues for $Ly = \lambda Dy$ (where m' is the target reconstruction dimensionality);
    (c) The m' eigenvectors are concatenated column-wise to form a matrix P', and the i-th row of P' is the reconstructed point of the original point p_i.

4. Approach

In this article, we focus on generating point sketch maps for *directional relations*, which are a well studied and utilized form of qualitative spatial relations (Ligozat 2013). Let $V = \{v_1, \dots, v_n\}$ be a set of n geographical entities. For each $v_i \in V$, there are some other entities $v_{i_1}, \dots, v_{i_{k_i}} \in V$, whose qualitative directional relations with respect to v_i are known. Let $N^+(v_i) = \{v_{i_1}, \dots, v_{i_{k_i}}\}$ be the set of *directed* neighbours of v_i. Then $\mathcal{R}(v_i) = \{(v_j\ R_{ji}\ v_i) \mid R_{ji} \in \{\mathsf{N}, \mathsf{Ne}, \mathsf{E}, \mathsf{Se}, \mathsf{S}, \mathsf{Sw}, \mathsf{W}, \mathsf{Nw}\}, v_j \in N^+(v_i)\}$ is the set of directional relations between v_i and its directed neighbours. For instance, $(v_j\ \mathsf{N}\ v_i)$ denotes that v_j is to the north of v_i. Note that these directional relations are just text descriptions, whose semantic meaning might not exactly be the cone-shaped qualitative directional relations defined in Section 3.

Problem Formulation The problem of *generating a point sketch map from qualitative directional relations* studied in this article is formalized as follows:

- **Input:** a set of geographical entities $V = \{v_1, \dots, v_n\}$ whose geometric information might not be known, and $\mathcal{R} = \bigcup_{v_i \in V} \mathcal{R}(v_i)$, i.e., information about directional relation between geographical entities (as defined earlier);
- **Output:** A point sketch map of the geographical entities in V on the two-dimensional plane, where each v_i is mapped onto a unique point $p_i = (x_i, y_i)$ on the plane.
- **Objective:** For any v_i and $v_j \in N^+(v_i)$, the points p_i and p_j in the sketch map should preserve the directional relations between v_i and v_j as much as possible.

In other words, given two geographical entities v_i and v_j, if the directional relation of v_j to v_i is R_{ji} and in the sketch map their corresponding points are p_i and p_j, then it is desired that the qualitative directional relation R^*_{ji} between p_i and p_j is "conceptually close" to R_{ji} (see Definitions 3.1 and 3.2).

To achieve good reconstruction results, there are two types of information that are very important and need to be preserved as much as possible: one is *the neighbourhood information*, and the other is *the qualitative directional relation information*. The neighbourhood information tells us which entities should be neighbours in the sketch map, so that the relative distance between originally adjacent entities is maintained; and the qualitative directional relation information tells us what kind of directional relations the entities should have in the sketch map.

To this end, the general idea is first to randomly generate points (x_i, y_i); then combine the idea of data reconstruction of LE to construct a neighbour loss, and design a new qualitative directional relation loss to preserve important information; and finally minimize the loss function by gradient descent optimization, to obtain the final point sketch map of the entities in V. The following subsections discuss the above steps in detail.

4.1. ***Neighbour Loss***

In the field of machine learning, LE is a data reconstruction method commonly used to preserve the neighbourhood information between data points. The main idea of LE is to make similar points (or points with close distances) closer to each other after reconstruction, so that the neighbourhood information can be better maintained.

Inspired by this idea, to preserve the neighbourhood information among the entities in the sketch map, this article first establishes an initial neighbour structure using the given qualitative directional relation set $\mathcal{R}$ and constructs the corresponding similarity matrix W. Given two entities v_i and v_j, $W_{ij} = 1$ if $(v_j R_{ji} v_i) \in \mathcal{R}(v_i)$, and $W_{ij} = 0$ otherwise. In other words, W_{ij} for v_i and v_j is 1 if the directional relation between them is given in the relation set. The reason for doing so is to highlight the difference between neighbouring and non-neighbouring entity pairs during the optimization process, and to facilitate the preservation of proximity of entities. Note that in this article, as only directional relations are assumed to be available for the proposed approach, we consider the points that are interrelated as neighbours of one another; this assumption may not be true in general, but we have found it to work well in our setting, because to a certain extent it well captures real-world neighbourhood information, as is the case with our dataset. For example, one would position/describe a small village with respect to some nearby town, or another village, rather than a capital that could be in the other part of the country. Of course, in situations where additional information is available, such as distance information, W can be further refined to more precisely reflect the neighbour structure of entities. Since this article focuses more on the effectiveness of the idea, this extension will not be deeply discussed here.

Subsequently, inspired by the loss function of LE, *neighbour loss* is constructed as:

$$\sum_{i=1}^{n} \sum_{j=1}^{n} W_{ij} \|p_i - p_j\|^2. \tag{3}$$

By minimizing the neighbour loss, we can obtain a point sketch map that preserves the neighbourhood information. Other approaches could also work when sufficient information is available. For example, Multidimensional Scaling (Cox and Cox 2008, MDS) tries to maintain the distance in a lower dimensional space, for every pair of points in the original space, by minimizing $\sum_{i=1}^{n} \sum_{j=1}^{n} (\|p_i - p_j\| - d_{ij})^2$ (d_{ij} is the given distance between entities v_i and v_j). This is different from LE in that LE does not maintain the distances but tries to keep similar points still close to each other after reconstruction.

However, neighbour loss alone is not enough for a good result, as preserving the neighbour information does not necessarily ensure the preservation of qualitative directional relations. For example, Figure 3(a) depicts the actual distribution of coordinate points for some geographical entities, while Figure 3(b) shows the point sketch map obtained using only the neighbour loss. From the figure, it can be observed that although the generated points are relatively close to each other, the directional relations based on the reference point of "St Aubyn" do not agree with the original relations. For example, "Thornville" was originally to the west of "St Aubyn", i.e., ("Thornville" W "St Aubyn"), but it becomes E or Ne in the sketch map. Similar inconsistencies in the qualitative directional relations can also be observed for "Emu Creek" and "Mount Binga".

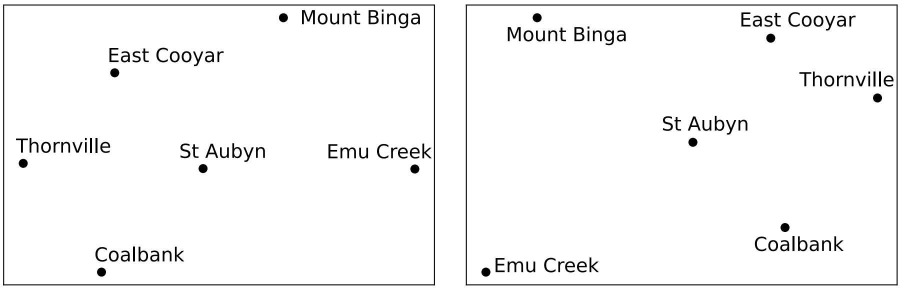

(a) Visual distribution of place names on a real map

(b) Generated point sketch map

Figure 3. An example of the problem of generating point sketch maps using only the neighbourhood loss for optimization.

4.2. *Qualitative Directional Relation Loss*

For any $v_i \in V$ and a neighbour $v_j \in N^+(v_i)$, suppose the generated points in a sketch map are $p_i = (x_i, y_i)$ and $p_j = (x_j, y_j)$, respectively. To quantify the consistency of the qualitative directional relations, in this article we first quantify an input qualitative directional relation, by setting a *canonical vector* $\overrightarrow{d_R}$ for each directional relation $R \in \{\mathsf{N}, \mathsf{Ne}, \mathsf{E}, \mathsf{Se}, \mathsf{S}, \mathsf{Sw}, \mathsf{W}, \mathsf{Nw}\}$, as shown in Figure 4. For example, the canonical vector $\overrightarrow{d_{\mathsf{E}}}$ of E is $(1, 0)$, and $\overrightarrow{d_{\mathsf{Sw}}} = (-1, -1)$.

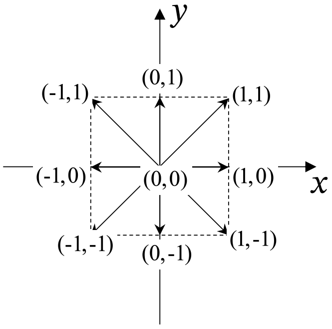

Figure 4. The canonical vectors of qualitative directional relations.

Then the qualitative directional relation of p_j to p_i in the sketch map is quantified by using the vector $\overrightarrow{p_i p_j}$, and the degree of consistency can be computed by comparing the difference with the vector $\overrightarrow{d_R}$, where R corresponds to the original directional relation between p_i and p_j, which can be characterized by the cosine value of the angle between them:

$$\cos(\overrightarrow{p_i p_j}, \overrightarrow{d_{R_{ji}}}) = \frac{p_j - p_i}{\|p_j - p_i\|} \cdot \frac{\overrightarrow{d_{R_{ji}}}}{\|\overrightarrow{d_{R_{ji}}}\|}. \tag{4}$$

A larger cosine value corresponds to a smaller angle, indicating that the difference between $\overrightarrow{p_i p_j}$ and the canonical vector is smaller, and the reconstruction is more consistent with the given qualitative direction relation. Thus, we propose the following as the overall *qualitative directional relation loss*:

$$-\sum_{i=1}^{n} \sum_{j=i_1}^{i_{k_i}} \frac{p_j - p_i}{\|p_j - p_i\|} \cdot \frac{\overrightarrow{d_{R_{ji}}}}{\|\overrightarrow{d_{R_{ji}}}\|}. \tag{5}$$

4.3. *Model and Algorithm*

Taking both losses into consideration, we propose to model the problem of generating a point sketch map from qualitative directional relations as the following optimization problem:

$$\underset{p_1,\ldots,p_n}{\operatorname{argmin}} \underbrace{\sum_{i=1}^{n} \sum_{j=1}^{n} W_{ij} \|p_i - p_j\|^2}_{\text{Neighbour Loss}} \underbrace{- \lambda \sum_{i=1}^{n} \sum_{j=i_1}^{i_{k_i}} \langle \frac{p_j - p_i}{\|p_j - p_i\|}, \frac{\overrightarrow{d_{R_{ji}}}}{\|\overrightarrow{d_{R_{ji}}}\|} \rangle}_{\text{Qualitative Directional Relation Loss}} . \tag{6}$$

Note that for the “Neighbour Loss”, smaller values are better, while for the summation term in “Qualitative Directional Relation Loss”, as it is a summation of cosine values, larger values are better. So in the optimization problem we subtract the latter from the former and take the argmin. Here, we also introduced a weight λ to balance the “Neighbor Loss” and “Qualitative Directional Relation Loss”.

Compared to the classical LE, the objective function in Equation 6 is more complex and cannot be directly solved by spectral decomposition. Instead, we consider using a gradient descent algorithm (Amari 1993): we first randomly initialize the coordinates of the points, and then the coordinates of the points are iteratively updated based on the gradient direction of the loss function. Different gradient descent methods exist, such as stochastic gradient descent method, batch gradient descent method, and Adam gradient descent method (see (Ruder 2016) for a more detailed introduction). It is worth noting that none of these gradient descent methods can guarantee finding the global optimal solution. Nevertheless, gradient descent methods have been widely used to approximately solve optimization problems, such as optimizing neural networks. Generally, the updating of coordinates can be done by the standard gradient descent method:

$$P := P - \alpha \frac{\bigtriangledown \mathcal{L}(P_{n\times 2}, \mathcal{R})}{\bigtriangledown P}, \tag{7}$$

where the learning rate α determines the step size of each gradient descent step. In the experiments, we implement the algorithm with a commonly used extension of the standard gradient descent method, called Adaptive Moment Estimation (Adam) by (Kingma and Ba 2014), which can adaptively control the learning rate and have better convergence guarantee.

By iteratively updating the current solution, gradient descent finds a solution that minimizes the objective function locally. However, note that the second term of the loss function is only related to the directions of the vectors, which has scale invariance, that is, scaling up or down the coordinates of the data points by a constant factor will not affect its value. On the other hand, the first term of the loss function prefers to have pairs of points with non-zero similarity value W_{ij} be constructed as close as possible. Therefore, by combining the two terms in the loss function, the data points tend to be over-compressed when being very close to one another, which is not desired.

Traditional LE prevents over-compression by adding a weighted orthogonal constraint to the optimization problem. The proposed model can also apply a similar constraint, which essentially normalizes the coordinates of the points, but will distort the distribution of the points. Two other options can be used to stop the updating of coordinates before over-compression: (1) the second term of the loss no longer decreases with respect to a threshold, and (2) the number of iterations exceeds a given number. In this article, the second method is used to stop the updating by controlling the number of iterations, and in practice one can choose other stopping criteria (like a combination of the two aforementioned methods) as appropriate.

The proposed steps to generate a point sketch map is shown in Algorithm 1. The input of the algorithm is a set of qualitative directional relations among $n = |V|$ entities $\mathcal{R} = \{(v_j R_{ji} v_i) | R_{ji} \in \{\mathsf{N}, \mathsf{Ne}, \mathsf{E}, \mathsf{Se}, \mathsf{S}, \mathsf{Sw}, \mathsf{W}, \mathsf{Nw}\}, v_i \in V, v_j \in N^+(v_i)\}$, and the output is an $n \times 2$ matrix, where the i-th row corresponds to the generated point of the i-th entity.

Algorithm 1: Point Sketch Map Generation Algorithm

```
Input: A set $\mathcal{R}$ of qualitative directional relations between $n$ geographical
       entities, a set $\mathcal{D}$ of the canonical vectors (see Figure 4), number of
       iterations $K$, and initial iteration number epoch.
Output: A matrix $P_{n\times 2}$ of generated points
Randomly generate an $n \times 2$ matrix $P_{n\times 2}$;
foreach epoch $\in (1, 2, \dots, K)$ do
    $\mathcal{L} \leftarrow (\sum_{i=1}^{n} \sum_{j=1}^{n} W_{ij} \|p_i - p_j\|^2 - \lambda \sum_{i=1}^{n} \sum_{j=i_1}^{i_{k_i}} \langle \frac{p_j - p_i}{\|p_j - p_i\|}, \frac{\overrightarrow{d_{R_{ji}}}}{|\overrightarrow{d_{R_{ji}}}|} \rangle)$;
        // Compute the loss using the $P$ obtained in the previous
        step.
    $P \leftarrow P - \alpha \frac{\nabla \mathcal{L}(P_{n\times 2}, \mathcal{R})}{\nabla P}$; // Compute the gradient of $\mathcal{L}$ with respect
        to $P$, and update $P$ using gradient descent.
return $P_{n\times 2}$;
```

5. Experimental Analysis

5.1. *Datasets and Settings*

To evaluate the proposed approach, we would need some ground-truth points corresponding to geographical entities and qualitative directional relations between these entities. The dataet of ground-truth points (GEO) in this article consists of the centroids of some suburbs in Australia, derived from the Australian Bureau of Statistics.[2] Figure 5(a) illustrates the distribution of some points in GEO.

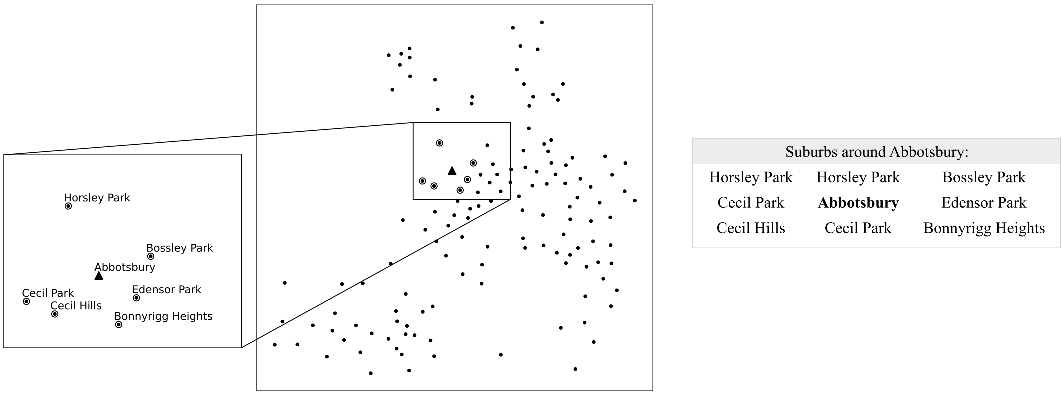

Suburbs around Abbotsbury:		
Horsley Park	Horsley Park	Bossley Park
Cecil Park	**Abbotsbury**	Edensor Park
Cecil Hills	Cecil Park	Bonnyrigg Heights

(a) (b)

Figure 5. Illustration of the datasets: (a) Some ground-truth points in GEO; (b) An infobox representation of neighbours and qualitative directional relations.

There are two sources of qualitative directional relations used in our experiments: the ones calculated from the points in GEO by a relation model (GEO-REL), and the ones annotated by users (WIKI-REL). The first source can be regarded as having no noise or errors (at least with respect to the centroids of Australian suburbs), and having stable semantic interpretations of the relations, which is mainly used to verify the effectiveness of the proposed approach and to establish an upper bound. In this article, we use the cone-shaped model in Section 3.1 to calculate the qualitative directional relations between each point and its neighbours to obtain the data set GEO-REL.

The second source, however, contains some noise and errors, and has varying semantic interpretations of the relations, as they are annotated by different users. It is mainly used to verify the robustness and the ability of the approach in dealing with inaccurate information. In this article, the data set of annotated qualitative directional relations is from the DBpedia[3] collection of Wikipedia infobox data. An infobox contains a 3×3 table that describes the qualitative directional relations between neighbouring place names, where the central cell represents the reference place name, and the eight surrounding cells correspond to the eight directional relations respectively (e.g., Nw, Se). As shown in Figure 5(b), the entity "Abbotsbury" has "Horsley Park" as its northwestern (Nw) neighbour and northern (N) neighbour, and "Bonnyrigg Heights" as its southeastern (Se) neighbour. Figure 5(a) also illustrates the ground-truth points corresponding to these entities in GEO.

Note that the aforementioned GEO-REL datasets use the centroids of the places from Wikipedia, but the places actually have much more complex polygon shapes than a single point. This simplification may result in mismatched cardinal directions between the same pair of places in these two datasets, as the cardinal direction between polygons could be different from the one between two points in these polygons (Regalia *et al.* 2016). Nevertheless, here we consider this difference as a source of noise to verify that this approach can deal with inaccurate information.

An infobox may contain multiple directional relations between the same pair of entities, such as "Horsley Park" and "Abbotsbury" in Figure 5(b). Therefore, for the

---

[2] `https://www.abs.gov.au/statistics/standards/`

[3] `https://www.dbpedia.org/`

sake of convenience of comparison, in WIKI-REL, for each pair of entities, we only keep the directional relation which is conceptually closest to the calculated directional relation in GEO-REL, i.e., the relation with the smallest distance (Definition 3.2) to the relation in GEO-REL. For example, Table 1 shows the directional relations between "Abbotsbury" and its neighbours in GEO-REL and WIKI-REL after this pre-processing. We can see that the relation N between Horsley Park and Abbotsbury is removed in WIKI-REL while Nw is kept.

Table 1. Qualitative directional relations between "Abbotsbury" and its neighbours in WIKI-REL and GEO-REL.

Name	Neighbour	WIKI-REL	GEO-REL
Abbotsbury	Horsley Park	Nw	Nw
	Bossley Park	Ne	E
	Edensor Park	E	Se
	Bonnyrigg Heights	Se	Se
	Cecil Hills	Sw	Sw
	Cecil Park	W	W

Finally, we select four subsets of entities, such that each entity in a subset has at least one neighbour, forming four subsets of GEO-REL and four subsets of WIKI-REL, as shown in Table 2. In the table, we also include a column called "consistency rate" which represents the proportion of direction relations in a dataset that are consistent with those in GEO-REL. This reflects the inaccuracy of the directional relations in the corresponding dataset, where a lower consistency rate indicates greater inaccuracy.

Table 2. Datasets extracted from WIKI-REL and GEO-REL.

Dataset	#Names	#Relations	Consistency Rate
GEO-REL-1	148	653	100
GEO-REL-2	146	651	100
GEO-REL-3	125	558	100
GEO-REL-4	137	593	100
WIKI-REL-1	148	653	81.32
WIKI-REL-2	146	651	83.72
WIKI-REL-3	125	558	74.55
WIKI-REL-4	137	593	72.18

The experiments are conducted on a computer with an AMD-EPYC-7763 2.45GHz CPU, NVIDIA 3090Ti GPU, and an Ubuntu 20.04 system with Python 3.9 and PyTorch 2.0.0. The parameters are set as follows. The weight λ in the loss function balances the effect of the two loss terms in Eq. 6, and, as the main objective is to optimize the relation loss, we give more focus on it and hence set λ to 10. The number of iterations controls when the optimization should be stopped. Based on our observations, the loss values for the used datasets usually start to converge after 1 000 iterations, so to ensure sufficient optimization we set universally $K = 5\,000$. The Adam algorithm in PyTorch is used, instead of the standard gradient descent method, and the learning rate α is set to 0.01 as is common practice. The initial randomly generated point coordinates are in the range of $[0, 1)$.

5.2. *Intuitive Evaluation of Point Sketch Map Generation*

To intuitively evaluate the effectiveness of the proposed point sketch map generation approach, we will visually compare the generated point sketch maps of GEO-REL and WIKI-REL with the distribution of the ground-truth points in GEO. Please note that a metrics-based evaluation follows immediately after, and with the intuitive evaluation here we aim to help the reader to better grasp the results of our approach.

5.2.1. Point Sketch Maps of GEO-REL

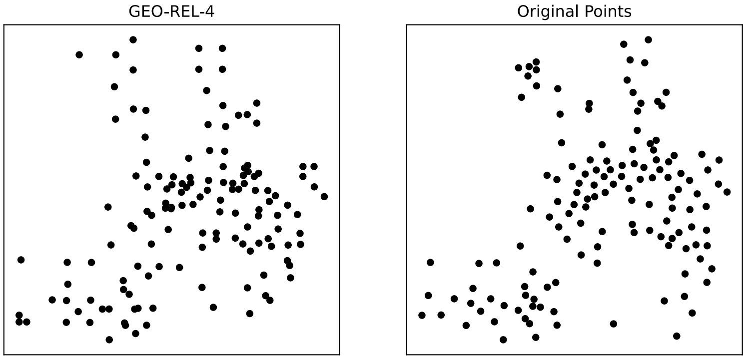

Figure 6. A case study of point sketch map generation for GEO-REL-4. Left is the sketch map, and right is the ground-truth points.

To conserve space, we only show results for GEO-REL-4, as qualitatively similar results were obtained for GEO-REL-1/2/3 (see the supplementary material for the remaining results). From the overall distribution of points in Figure 6, the proposed approach well preserves local neighbouring information and qualitative directional relations. The neighbour loss and the qualitative directional relation loss help to avoid cases where non-neighbouring points are too close and directional relations are distorted among neighbouring points. It is worth noting that although only the relations between local neighbours are given, due to the transitivity of relation constraints, the positions between non-neighbours are also mutually constrained. For instance, when B is east of A and C is east of B, the positions of A and C are constrained by each other even if there is no relation between them, e.g., C cannot be to the west of A, otherwise the relation between C and B will be violated. In fact, our additional experiments on fewer available relations show that this approach can perform well even when the relational information is severely limited, e.g., it can still have an accuracy over 0.8 after removing 40% of the relations in the GEO-REL datasets (see the supplementary material for details).

In terms of the point distribution around specific points, Figure 7 shows the sketch maps of two reference place names and their neighbours in GEO-REL-4. The reference point 26 has 6 neighbours, and the qualitative directional relations between them given in GEO-REL-4 are: 8 Nw 26 (8 is in the northwest direction of 26), 9 W 26, 25 Ne 26, 27 W 26, 43 E 26, and 44 Se 26. From the figure, it can be seen that these qualitative directional relations are well-preserved by the approach. For example, point 8 does fall in the Nw direction of point 26, and the direction between point 43 and point 26 is also very close to the given E. On the other hand, the reference point 85 has 4 neighbours, and the qualitative directional relations between them given in GEO-REL-4 are: 47 W 85, 84 Ne 85, 86 W 85, and 105 Se 85. Most of the qualitative directional relations are also well-preserved in the sketch map (marked as circled points), e.g. point 47 does fall in the west direction of point 85, and the direction between point 84 and point 85 is also very close to the northeast direction. However, the direction between point 86 (marked as a crossed point) and point 85 has a small deviation from the given direction W, and is closer to Nw instead. Nevertheless, this deviation is within an acceptable range in terms of human cognition.

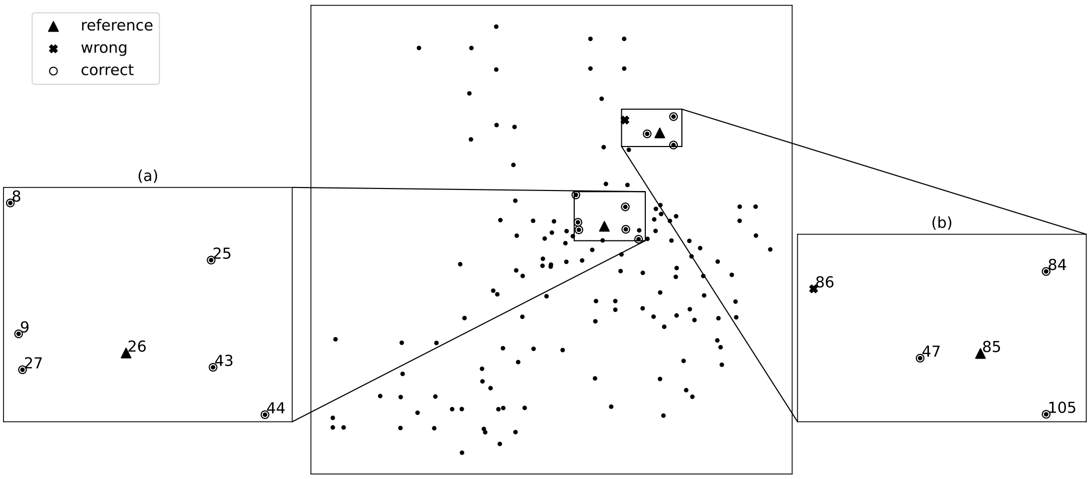

Figure 7. A detailed inspection on the point sketch map of GEO-REL-4. (a) An example where all directional relations in the sketch map are correct; (b) An example where one directional relation in the sketch map is wrong.

5.2.2. Point Sketch Maps of WIKI-REL

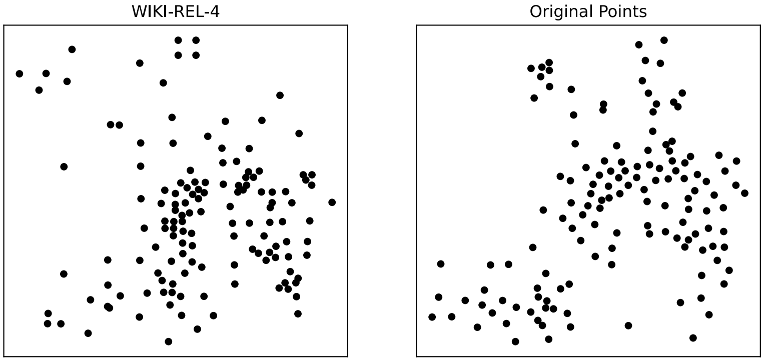
WIKI-REL-4

Original Points

Figure 8. A case study of point sketch map generation for WIKI-REL-4. Left is the sketch map, and right is the ground-truth points.

Figure 8 shows the generated point sketch map of WIKI-REL-4 and the corresponding distribution of ground-truth points in GEO (results for WIKI-REL-1/2/3 are qualitatively similar and can be found in the supplementary material). Note that even though the relation consistency rate of WIKI-REL-4 is 72.18% due to inaccuracy and inconsistency introduced by users, the overall result still shows that the proposed approach can maintain the neighbour information and qualitative directional relations to a good extent.

In Figure 9, we selected three cases to investigate in more depth: (a) All directional relations are correct in the sketch map. The qualitative directional relations between point 108 and its surrounding neighbours in WIKI-REL-4 are: 78 N 108, 88 Nw 108, 89 Se 108, and 90 Sw 108. (b) Few directional relations are wrong in the sketch map. The qualitative directional relations between point 85 and its surrounding neighbours in WIKI-REL-4 are: 47 W 85, 84 Ne 85, 86 Nw 85, and 105 E 85. However, in the generated point sketch map, point 105 is more to the southeast of point 85, and the actual difference from the original relation E is small. (c) Most directional relations are wrong in the sketch map. The qualitative directional relations between point 11 and its surrounding neighbours in WIKI-REL-4 are: 3 Nw 11, 4 W 11, 12 Sw 11, 19 N 11, 20 Ne 11, 21 E 11, and 22 Se 11. In the sketch map, some relations are slightly offset, such as point 20, which is more towards the E direction of point 11, but the offset is still acceptable. This might be because some relations in the data are not very accurate (Table 2), which causes some points in the sketch map to be placed at wrong positions, and in turn makes the method incapable of placing other points at the right positions. Moreover, as we will show in the evaluation later on, our sketch maps often fall under cases (a) and (b). This further illustrates that our method is effective even when the given qualitative information is less accurate.

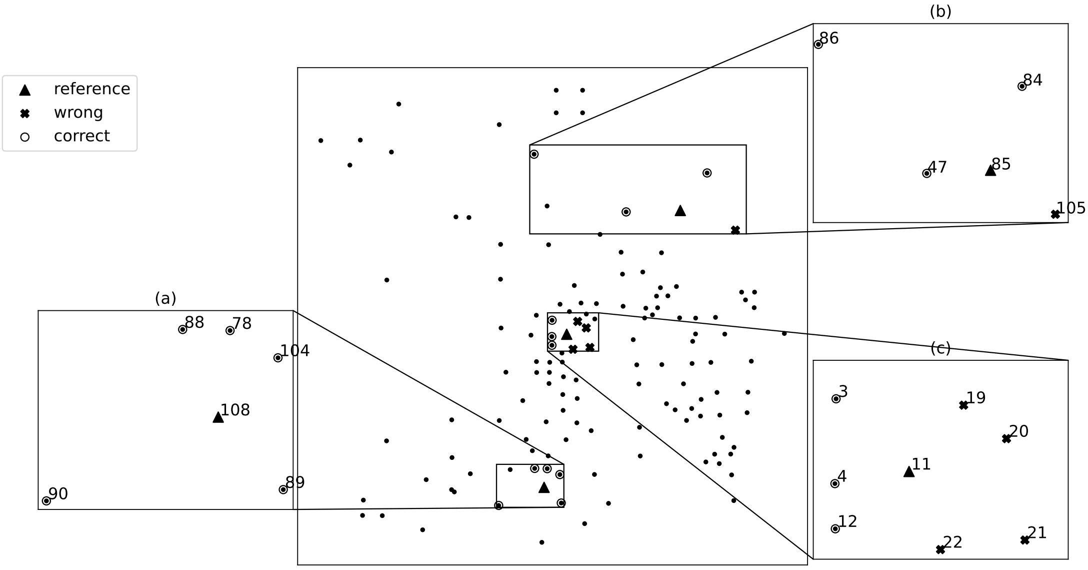

Figure 9. A detailed inspection on the point sketch map of WIKI-REL-4. (a) An example where all directional relations in the sketch map are correct; (b) An example where one directional relation in the sketch map is wrong; (c) An example where most of the relations are wrong in the sketch map.

5.3. *Metrics-based Evaluation of Point Sketch Map Generation*

Note that once the points for geographical entities are obtained by the algorithm, we can use the obtained points to calculate the directional relations with the cone-shaped model as detailed in Section 3.1. Then the performance can be evaluated by comparing the relations from the sketch map with the ground-truth ones. The following two metrics are selected for evaluation: "accuracy rate" and "total error distance". The metric "accuracy rate" (ACC) refers to the proportion of correctly preserved relations to the total number of relations in the dataset. The "error distance" of a relation from the sketch map to the ground-truth one refers to the conceptual neighbourhood distance (see Definition 3.2) between these two relations. Note that the range of the error distance for a single relation is $[0, 4]$. The "total error distance" of a whole point

Table 3. Metrics for generated point sketch maps of GEO-REL and WIKI-REL datasets.

Datasets	**GEO-REL Datasets**				**WIKI-REL Datasets**			
	1	2	3	4	1	2	3	4
#Names	148	146	125	137	148	146	125	137
#Relations	653	651	558	593	653	651	558	593
#Correct/#Incorrect	646/7	644/7	550/8	574/19	556/97	564/87	489/69	507/86
Total Error Distance	7	7	8	19	99	87	69	87
ACC (%)	98.93	98.92	98.57	96.80	85.15	86.64	87.63	85.50

sketch map is the sum of the distances between all relations from the sketch map and their corresponding ground-truth relations.

Table 3 quantitatively presents the quality of the generated point sketch maps for GEO-REL-1/2/3/4 and WIKI-REL-1/2/3/4. For GEO-REL datasets, compared to the calculated directional relations based on points from GEO, the results show that the proposed approach has an accuracy of nearly 100%. Moreover, the error distance of each wrong relation does not exceed 1 (note that the range of error distance of a single relation is $[0, 4]$). For example, a ground-truth relation W might have become Nw in the sketch map, which however is usually acceptable in many real-world applications. In addition to what is presented in Table 3, the average error in angles, when comparing the generated point sketch maps with the original points, is about 9.5° for GEO-REL datasets and 11.6° for WIKI-REL datasets.

For WIKI-REL datasets, compared to the provided relations in infoboxes, the accuracy is in the range between 85.15% and 87.63%. Regarding the error distance, we note that most of the incorrect relations have an error distance of 1, and only very few have an error distance of 2, indicating that the offset of the incorrectly demonstrated relations is small. It is worth noting that the annotated directional relations in the WIKI-REL datasets are originally not accurate enough, having a consistency rate of about 80% (as shown in Table 2). This indicates that the proposed approach can accurately preserve the qualitative directional relations within distance 1, which is acceptable in many real-world scenarios, and thus reflects the practical value of this method.

As a more specific illustration, Figure 10 presents the detailed numbers of correct/incorrect relations for each entity and its neighbours in GEO-REL-4 (see the supplementary material for the remaining results). The height of each bar represents the total number of relations for the corresponding entity, with the grey portion indicating the number of incorrect relations and the white portion indicating the number of correct relations. The results show that for each entity with incorrect directional relations, the number of incorrect relations is actually small (no more than 2 for this dataset).

Figure 11 presents the detailed numbers of correct/incorrect relations for each entity and its neighbours in WIKI-REL-4 (see the supplementary material for the remaining results). The results show that most of the qualitative directional relations between entities can be correctly demonstrated. There are individual entities with more incorrect relations (≤ 4), but the error distances are small (most of the distances are 1 and only one is 2). Overall, the proposed approach can preserve the qualitative directional relations well even with inaccurate directional information.

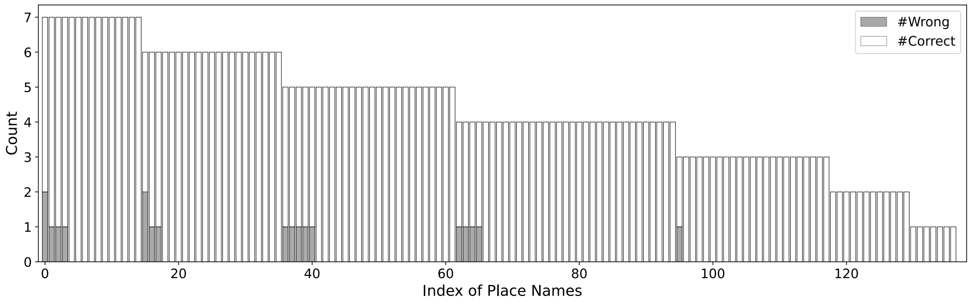

Figure 10. Correctly/incorrectly preserved directional relations compared to the original ones for each entity in GEO-REL-4.

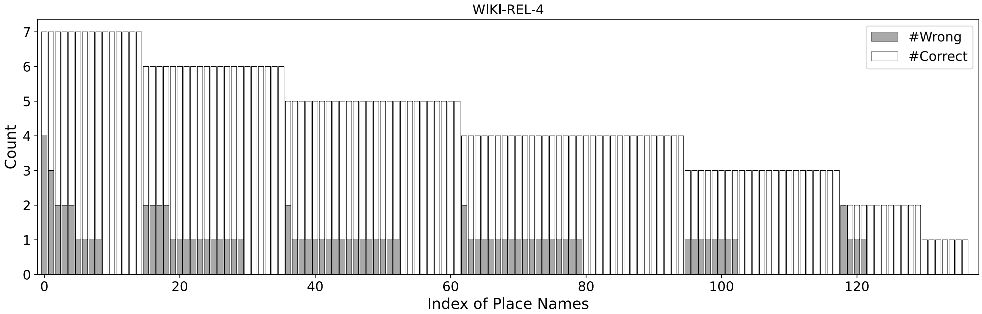

Figure 11. Correctly/incorrectly preserved directional relations compared to the original ones for each entity in WIKI-REL-4.

5.4. *Comparison with Baselines*

Following the discussion in Section 2, the methods below are chosen as baselines for generating sketch maps, in order to obtain an as-complete-as-possible view across different techniques for solving our novel problem; clearly, and as we detail below, we had to adapt these techniques to our problem formulation.

- Kim (Kim *et al.* 2016): This method iteratively generates objects and dynamically adjusts their positions and sizes to satisfy new relations. Note that this method is originally designed for generating sketch maps with rectangles instead of points, and we re-implement its idea for generating sketch maps of points.
- Belouaer (Belouaer *et al.* 2016): This method makes use of a genetic algorithm to generate sketch maps of points. Following (Belouaer *et al.* 2016), we use a matrix of size $n \times 2$ to represent the points, and the population consists of 100 such matrices as individuals. We use the best individual as the final points after 100 generations.
- SparQ (Wolter and Wallgrün 2012): This is a qualitative spatio-temporal reasoner that can obtain a quantitative instantiation of variables (e.g., points) that satisfies a given set of relations. It is worth noting that SparQ uses a different directional relation model as in this article, viz., the projection-based model (Frank 1992). For example, N holds when the reference point and the target point are on the y-axis, which is a stricter requirement than when considering a region of acceptance as in the considered model here. Therefore, directly using the relations in the datasets can create inconsistencies and make SparQ fail to find a solution (clearly, the two models do not generally align). So we used the relation model in SparQ to first re-calculate the relations between the ground-truth points in GEO, and then use the new datasets (GEO-REL-SparQ) to compare with the

Table 4. ACC metrics for different methods on GEO-REL and WIKI-REL datasets. Note that SparQ used a different relation model for the directional relations.

Method	**GEO-REL Datasets**				**WIKI-REL Datasets**			
	1	2	3	4	1	2	3	4
Ours	98.93	98.92	98.57	96.80	85.15	86.64	87.63	85.50
Kim	47.17	42.24	45.34	45.70	42.57	43.47	43.01	44.18
Belouaer	15.47	15.48	16.00	16.26	15.48	15.42	15.72	15.67
SparQ	100.00	100.00	100.00	100.00	-	-	-	-

proposed approach on the original GEO-REL datasets.

The results are shown in Table 4. It is important to note that when there is no conflicting information—as is the case for the GEO-REL-SparQ datasets obtained from GEO, SparQ will always correctly identify points that satisfy all the relations, as it relies on a correct procedure of generating a quantitative solution (valuation) from a qualitative configuration, so this result is nothing new. However, when there is conflicting information—as is the case for the WIKI-REL datasets, it will fail completely. Both Kim and Belouaer methods have relatively poor performance, with only about 45% and 16% of the relations being able to be correctly realized, respectively. This demonstrates the superiority of the proposed approach on generating point sketch maps for qualitative directional relations, especially when there is conflicting and noisy information.

Regarding the runtimes, for Belouaer, it is about 600 seconds for each dataset; for Kim and SparQ, it is around 1 second. For ours, it is about 5 seconds for 5 000 iterations.

5.5. *Scalability*

To see how the proposed approach scales, for different values of n (from 1 000 to 10 000 with a step size of 1 000), we randomly select n points from GEO, and for each of the n points, we calculate the directional relations between a point and its k neighbours (from 3 to 21 with a step size of 3) in the n points, using the cone-shaped model in Section 3.1. The relations are then used as input of the proposed approach to generate sketch maps with 5 000 iterations. The results are averaged over 3 runs for each n and each k and are shown in Figure 12.

It can be seen that the proposed approach, which has a runtime complexity approximately in $O(n^2)$, can scale to thousands of places and relations. For the case with 10 000 places and up to 21 relations for each place, the average time used to generate sketch maps with 5 000 iterations is 509 seconds. Also, the number of neighbours seems to have no significant effect on the runtime. A possible reason could be that the calculation of gradient dominates the runtime, for which different numbers of neighbours have a similar cost.

5.6. *Effect of Incorrect Relations*

To further demonstrate that the proposed approach can deal with a large proportion of incorrect/inconsistent relations, we randomly changed a certain proportion of relations to be corresponding opposite relations (e.g., N is changed to S when selected) and also to be relations with an error distance of 1 or 2 (e.g., N is changed to one of Ne, Nw, E, W) for the eight datasets of GEO-REL and WIKI-REL. The ratio of opposite or wrong

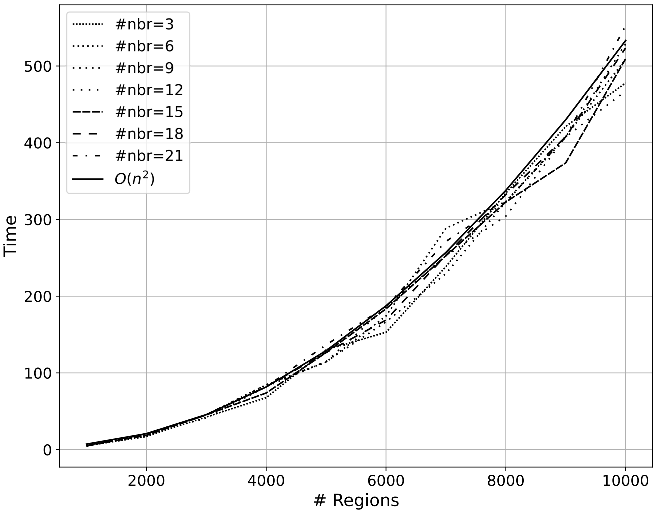

Figure 12. The runtime changes of the proposed approach for different number of places and neighbours.

relations varies from 0.1 to 0.9 with a step size of 0.1. The ACC metric values are averaged over 5 repeated constructions of sketch maps, comparing against the original relations. The results are shown in Figure 13.

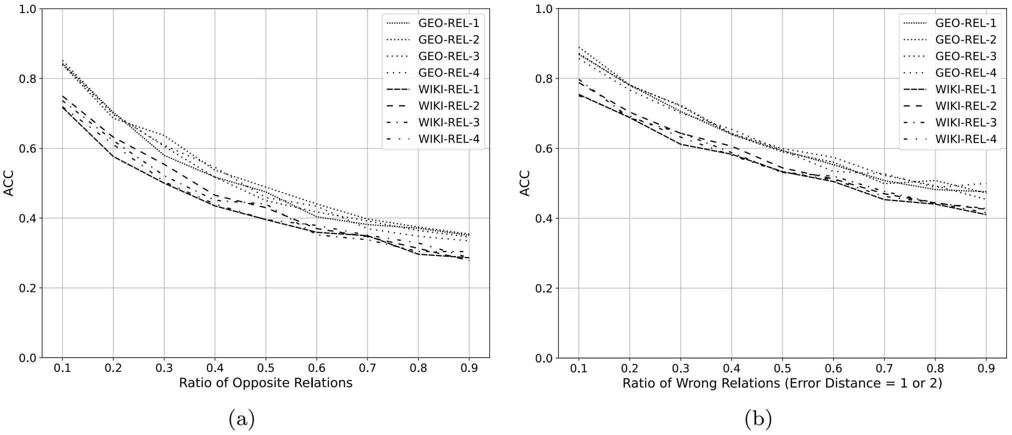

(a) (b)

Figure 13. The change of ACC values as the ratios of opposite relations and wrong relations (error distance = 1 or 2) increase.

Generally, the proposed approach can still maintain a relatively high accuracy ($>$ 0.6) when the ratio of opposite relations is smaller than or equal to 0.2. When the ratio is below 0.6, the number of unsatisfied relations is larger than the number of the provided incorrect relations. This might be because that the presence of incorrect relations can make it harder to find points that satisfy other correct relations. For wrong relations with an error distance of 1 or 2 in Fig. 13, we can see that the proposed approach can maintain a higher accuracy than in the case of opposite relations, and can achieve a 0.4 accuracy rate even with 90% of wrong relations. This demonstrates the robustness and error correction capability of the proposed approach.

Table 5. Comparison of metrics when generating sketch maps without and with qualitative proximity information for GEO-REL datasets.

Without Proximity	GEO-REL-1	GEO-REL-2	GEO-REL-3	GEO-REL-4
#Correct/#Incorrect	646/7	644/7	550/8	574/19
Total Error Distance	7	7	8	19
ACC (%)	98.93	98.92	98.57	96.80
With Proximity	GEO-REL-1	GEO-REL-2	GEO-REL-3	GEO-REL-4
#Correct/#Incorrect	650/3	649/2	554/4	589/4
Total Error Distance	3	2	4	4
ACC (%)	99.54	99.69	99.28	99.33

5.7. *Handling Proximity Information*

The proposed model takes neighbour information into account when formulating the loss function (Equation 6), by setting the similarity value of W_{ij} to 1 if two places are neighbours and to 0 otherwise. Qualitative proximity information can also be considered by varying the value of W_{ij} to reflect degrees of proximity.

In particular, to demonstrate how the proposed model can handle proximity information, we first categorize distances between points in the GEO-REL datasets, by using a threshold r, which is the average of the mean distances between 3-nearest neighbours of each point, i.e., $r = \frac{1}{n}\sum_{i=1}^{n}(\sum_{j\in N_3(i)} d_{ij})/3$. Points of distance less than or equal to r are considered as being *close*, points of distance between r and $2r$ are considered as being in the *middle*, and points of distance larger than $2r$ are considered as being *far*. For points with different qualitative proximity relations, we set different W_{ij} values to prioritize the optimization of closer points, i.e., $W_{ij} = 1$ for *close*, $W_{ij} = 0.5$ for *middle*, $W_{ij} = 0.25$ for *far*.

Table 5 compares the results obtained without and with qualitative proximity information. It can be seen that by considering qualitative proximity information in the model, more directional relations are satisfied, which means that proximity information can guide the model to find points that better align with the directional relations. This might be because the qualitative directional relations are not precise with respect to angles, and proximity information can help to correct the realized angles.

However, we find that the current model has difficulty in maintaining relative distance orders, e.g., two places that are considered close might have generated points whose distance value is not ranked in the first 1/3 of all distance values. In fact, about 2/3 of the distance values are ranked wrong in the generated point sketch map when considering qualitative proximity information, which is not better than the point sketch map that considers neighbouring points with the same global W_{ij}.

We also tried to add a distance loss in the form of a mean square loss of distance: $\sum_{i=1}^{n}\sum_{j=i_1}^{i_{k_i}}(\|p_i - p_j\| - d_{ij})^2$. Still, neither the distance values nor the relative distance relations can be satisfactorily maintained in the results. This might be because the model has difficulty in properly balancing between the distance/neighbour loss and the directional relation loss.

In summary, we only considered a very primitive strategy to characterize proximity information in the current model. There are still open and challenging problems, including how to better convert qualitative proximity information into quantitative ones, and how to balance between directional and proximity information.

Table 6. Comparison of metrics when generating sketch maps using either selected single relations or original multiple relations for WIKI-REL datasets.

Single Relation	WIKI-REL-1	WIKI-REL-2	WIKI-REL-3	WIKI-REL-4
#Correct/#Incorrect	556/97	564/87	489/69	507/86
Total Error Distance	99	87	69	87
ACC (%)	85.15	86.64	87.63	85.50
Multiple Relations	WIKI-REL-1	WIKI-REL-2	WIKI-REL-3	WIKI-REL-4
#Correct/#Incorrect	511/142	533/118	471/87	487/106
Total Error Distance	146	118	87	107
ACC (%)	78.25	81.87	84.41	82.12

5.8. *Generalizing to Multi-Relations*

As mentioned in Section 5.1, a pair of places might have multiple directional relations, whereas we only kept the relation that is conceptually closest to the calculated relation in GEO-REL in the previous experiments. In fact, the proposed model can be adapted to handle multiple relations between a pair of places. One approach is to take a (possibly weighted) average of the relational losses of these relations for a given pair, i.e., $-\sum_{i=1}^{n}\sum_{j\in N^{+}(i)}(\frac{1}{|R|}\sum_{r\in R}\frac{p_j-p_i}{\|p_j-p_i\|}\cdot\frac{\overrightarrow{d_{R_{ji}}}}{\|\overrightarrow{d_{R_{ji}}}\|})$, where $R=\{r_1,\ldots,r_m\}$ is a set of multiple relations between places i and j. This loss can balance the effect of different relations and guide the model to find points that have a relation for the corresponding pair of places that is conceptually closest to the multiple relations, that is, a consensus of multiple relations.

The results are shown in Table 6. Note that in the table the relations of the generated sketch maps in the case of multiple relations are compared against the selected single relations in the previous WIKI-REL datasets. It can be seen that the performance of the proposed model is still promising, although there is a decrease in performance when multiple relations are considered.

6. Conclusion and Future Work

In daily life, user descriptions of locations are often qualitative rather than quantitative, e.g., "the archaeological museum is to the north of the old town". However, although qualitative descriptions conform to human cognition, they are not intuitive enough compared to visualized maps. In order to convert qualitative information into a sketch map, traditional methods often design complex formulas based on some expert knowledge for such conversion, which have difficulties in handling qualitative spatial relations for a large number of entities and lack tolerance for inaccurate spatial relations. In this article, we propose an end-to-end approach that models the sketch map generation of qualitative directional relations into an optimization problem. This approach introduces the Laplacian reconstruction loss to characterize neighbourhood information, and incorporates a novel loss of qualitative directional relations to capture qualitative directional information. Gradient descent is exploited to iteratively reduce the value of the proposed loss function and thus obtain an approximate solution to the optimization problem. The proposed approach combines qualitative directional relations with machine learning techniques, and provides a new solution for quantification and visualization of qualitative relations. Extensive experiments conducted on various datasets, including automatically generated, user-annotated, noisy, and information-missing, verify the effectiveness of the proposed approach.

Implications. The proposed approach and the problem of converting qualitative spatial relations to numerical data have several important implications for geographical information science, qualitative spatial reasoning, and other relevant fields.

Firstly, by quantifying the large amount of qualitative spatial relations existing in various sources like social media and blogs, the proposed approach can enrich the knowledge base of GISs, and consequently help to better integrate qualitative information with current GISs that are mainly based on numerical data and methods. The obtained numerical representations can also be readily utilized by well-established analysis tools in GISs, and inspire new techniques from other fields to process qualitative information. For example, one can perform topological data analysis (Corcoran and Jones 2023) on the point sketch maps obtained with the proposed approach, providing a way to analysing qualitative relations and revealing patterns of geographical data in the view of topology. With geometrical representations instead of symbolic qualitative relations, it also becomes more convenient and powerful to calculate similarities between geographical places, as one can use various measurements based on numerical values. This can in turn enhance the matching of geographical entities for which originally only qualitative information was available.

Secondly, compared to pure qualitative relations, geometrical representations can help people to better understand and register geographical information, due to its intuitiveness. This can also improve the accessibility of geographical data and the experience of human-computer interaction in GISs. For example, there are plenty of qualitative descriptions of spatial relations among events and places, and by generating numerical representations of events and places, GISs can easily integrate historical information with current data, such as when visualizing historical scenes on top of modern maps. Users of location-based services, such as travellers who would like to record and share the routes and points of interest (POIs) during their travel, even when they do not have professional geographical backgrounds, can conveniently generate routes and POIs, by using the qualitative relations they recorded casually.

Thirdly, for qualitative spatial reasoning, quantifying qualitative relations offers the possibility of performing reasoning with numerical methods. For example, the provided qualitative relations between geographical entities are usually contradictory, e.g., due to the oversimplified interpretation of geometries of these entities. Deriving a plausible solution for these contradictory relations such that the maximum number of relations are satisfied is known as the MAX-QCN problem (Condotta *et al.* 2016). The proposed approach that generates points satisfying most of the relations can be considered as an efficient approximate solution or heuristic to this NP-hard MAX-QCN problem.

Lastly, the idea of the proposed approach can also help geographical education, by simplifying and visualizing the complex spatial information into intuitive maps. Besides, it is useful for other fields like the design of scenes in computer games, where designers can just specify the qualitative relations between objects in the scene and a draft can be generated automatically. For the generation of pictures or videos from texts that is popular due to the emergence of Large Language Models (LLMs), it can serve as a referee that can guide LLMs to understand the rules behind qualitative relations, such that the generated pictures or video are not against the fact.

Limitations and Future Work. Besides the aforementioned important implications and potential applications, the current approach still has several limitations that can also be topics of future work, as discussed in what follows.

Using points to approximate the places is inherently rough. As has been noted in the literature, e.g., (Peuquet and Ci-Xiang 1987), the shape, size, and even cultural and historical knowledge of places can affect the understanding of directional relations, and the relations between centroids might not align with this understanding. Still, the approach presented here introduces a promising way to quantitatively utilize qualitative spatial relations with machine learning and to visualize them in an end-to-end manner.

It is definitely interesting to investigate how other more complex geometries and other types of information can be used to generate more comprehensive sketch maps similarly. When more complex geometries, like rectangles and polygons, and other types of qualitative information, like relative directions, topological relations, hierarchy of objects, and proximity, are considered, devising a proper model to deal with complex geometries and incorporate different types of qualitative information is an open and challenging topic. For example, in this article, we consider places with directional relations among them as neighbours, while in practice places that are far away might also be involved in directional relations, e.g., landmarks in a city might have relations with places all over the city; though, we think that such relations would be comparatively less frequent (e.g., one would position the Eiffel tower with reference to the heart of Paris, rather than a suburb or a town in the south of France). This would result in perturbations in the generated sketch map. Although the constraint of directional relations could reduce the perturbations, it is interesting to see if one can detect and correct such inconsistency of neighbouring information to make the sketch map more concise. Obtaining directional relations from less structured sources (like social media posts, blogs, and books) with techniques including natural language processing and large language models is also important and challenging. Sometimes, crucial information, such as the directions of references for relative directions, is not explicitly available. Being able to derive such information by combining different sources is crucial for real-world applications. For example, to infer the directions of references, one can combine sensor data from mobile phones or photos at the same time of the generation of the descriptions, and/or make use of qualitative spatial reasoning techniques.

Although the current approach can scale to thousands of places and relations, it still can be improved to meet the needs of the era of big data. An emerging topic in Artificial Intelligence research is to combine symbolic reasoning with quantitative computation, and this can also help the current approach scale better. Specifically, the current approach did not consider the logical implications of the relations, which can actually guide the generation of points. For example, some relations might be redundant in the sense that satisfying some other relations can ensure this relation to be satisfied, and thus such relations can be removed to improve efficiency, cf. (Peng *et al.* 2023). Reasoning can also help to give better initial positions of points, which may help the optimization process to finish in far fewer iterations.

Acknowledgement

This work was supported by the National Natural Science Foundation of China under Grant number 61806170; the Fundamental Research Funds for the Central Universities under Grant numbers 2682022ZTPY082 and 2682023ZTPY027; French National Research Agency under Grant number SA21PD01; and University of Montpellier under Grant number PP21PD01-RM06.

We would also like to thank the editor and the reviewers for their invaluable help to improve the quality of this work.

Disclosure Statement

The authors report there are no competing interests to declare.

Notes on Contributors

Zhiguo Long is a Lecturer with SCAI - Southwest Jiaotong University, China His research interests include fundamental and practical techniques in geographical knowledge representation and reasoning. He contributed to the concept, implementation and writing of this article.
Qingqian Li is a master student at SCAI - Southwest Jiaotong University, China. Her dissertation topics include qualitative spatial representation and machine learning. She contributed to the visualization, data curation, and writing of this article.
Hua Meng is an Associate Professor with the School of Mathematics, Southwest Jiaotong University, China. His research interests include machine learning and topological data analysis. He contributed to the concept and writing of this article.
Michael Sioutis is a tenure-track Junior Professor with the Faculty of Sciences of the University of Montpellier. His research interests include artificial intelligence, data mining, and logic programming. He contributed to the investigation, validation, and writing of this article.

Data and Codes Availability Statement

The data and codes that support the findings of this study are available with the identifier at link (https://doi.org/10.6084/m9.figshare.25037462).

References

Amari, S.i., 1993. Backpropagation and stochastic gradient descent method. *Neurocomputing*, 5 (4-5), 185–196.

Belouaer, L., Brosset, D., and Claramunt, C., 2016. From verbal route descriptions to sketch maps in natural environments. *In*: *ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 1–10.

Chen, H., Vasardani, M., and Winter, S., 2018. Georeferencing places from collective human descriptions using place graphs. *Journal of Spatial Information Science*, 17 (1), 31–62.

Cheng, R., Liao, J., and Chen, J., 2022. Quickly locating pois in large datasets from descriptions based on improved address matching and compact qualitative representations. *Transactions in GIS*, 26 (1), 129–154.

Clementini, E., 2013. Directional relations and frames of reference. *GeoInformatica*, 17, 235–255.

Cohn, A.G. and Renz, J., 2008. Qualitative spatial representation and reasoning. *In*: *Handbook of knowledge representation*. vol. 3. 551–596.

Condotta, J.F., Kaci, S., and Schwind, N., 2008. A framework for merging qualitative constraints networks. *In*: *International Florida Artificial Intelligence Research Society Conference*. 586–591.

Condotta, J., Nouaouri, I., and Sioutis, M., 2016. A SAT approach for maximizing satisfiability in qualitative spatial and temporal constraint networks. *In*: *International Conference on Principles of Knowledge Representation and Reasoning*. 432–442.

Corcoran, P. and Jones, C.B., 2023. Topological data analysis for geographical information science using persistent homology. *International Journal of Geographical Information Science*, 37 (3), 712–745.

Cox, M.A.A. and Cox, T.F., 2008. *Multidimensional scaling*. Berlin, Heidelberg: Springer Berlin Heidelberg, 315–347.

Deng, M. and Li, Z., 2008. A statistical model for directional relations between spatial objects. *GeoInformatica*, 12, 193–217.

Du, H., *et al.*, 2023. A logic of east and west. *Journal of Artificial Intelligence Research*, 76, 527–565.

Dylla, F., *et al.*, 2017. A survey of qualitative spatial and temporal calculi: algebraic and computational properties. *ACM Computing Surveys*, 50 (1), 1–39.

Egenhofer, M.J. and Mark, D.M., 1995. Naive geography. *In*: A.U. Frank and W. Kuhn, eds. *International Conference on Spatial Information Theory*. 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., 1996. Temporal reasoning based on semi-intervals. *Artificial Intelligence*, 54 (1–2), 199–227.

Frontiera, P., Larson, R., and Radke, J., 2008. A comparison of geometric approaches to assessing spatial similarity for GIR. *International Journal of Geographical Information Science*, 22 (3), 337–360.

Goodchild, M.F. and Kimerling, A.J., 2002. Discrete global grids: A web book. Retrieved from `https://escholarship.org/uc/item/9492q6sm`.

Graullera, D.A., Moreno, S., and Escrig, M.T., 2006. Map building including qualitative reasoning for aibo robots. *Frontiers in Artificial Intelligence and Applications*, 146, 211.

He, L., *et al.*, 2015. Fuzzy spatial relation ontology driven detection of complex geospatial features in a web service environment. *Earth Science Informatics*, 8, 63–76.

He, X. and Niyogi, P., 2003. Locality preserving projections. *Advances in neural information processing systems*, 16.

Hua, H., Zhang, P., and Renz, J., 2019. Qualitative place maps for landmark-based localization and navigation in gps-denied environments. *In*: *ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 23–32.

Kim, J., Vasardani, M., and Winter, S., 2016. From descriptions to depictions: A dynamic sketch map drawing strategy. *Spatial Cognition & Computation*, 16 (1), 29–53.

Kim, J., Vasardani, M., and Winter, S., 2017. Similarity matching for integrating spatial information extracted from place descriptions. *International Journal of Geographical Information Science*, 31 (1), 56–80.

Kingma, D.P. and Ba, J., 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Ligozat, G., 2013. *Qualitative spatial and temporal reasoning*. John Wiley & Sons.

Long, Z., *et al.*, 2022. Approximating region boundaries based on qualitative and quantitative information. *IEEE Intelligent Systems*, 37, 38–47.

Moncla, L., *et al.*, 2016. Reconstruction of itineraries from annotated text with an informed spanning tree algorithm. *International Journal of Geographical Information Science*, 30 (6), 1137–1160.

Niyogi, P. and Belkin, M., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15 (6), 1373–1396.

Peng, J., Sioutis, M., and Long, Z., 2023. On neuro-symbolic challenges in directional relation prediction. *In*: M. Sioutis, Z. Long, J.H. Lee and M. Bhatt, eds. *International Workshop on Spatio-Temporal Reasoning and Learning (STRL 2023, co-located with IJCAI 2023)*. CEUR-WS.org, vol. 3475.

Peuquet, D.J. and Ci-Xiang, Z., 1987. An algorithm to determine the directional relationship between arbitrarily-shaped polygons in the plane. *Pattern Recognition*, 20 (1), 65–74.

Regalia, B., Janowicz, K., and Gao, S., 2016. Volt: A provenance-producing, transparent sparql proxy for the on-demand computation of linked data and its application to spatiotemporally dependent data. *In*: H. Sack, E. Blomqvist, M. d'Aquin, C. Ghidini, S.P. Ponzetto and C. Lange, eds. *The Semantic Web. Latest Advances and New Domains*. 523–538.

Romeny, B.M., 2003. *The gaussian kernel*. Dordrecht: Springer Netherlands, 37–51.

Roweis, S.T. and Saul, L.K., 2000. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290 (5500), 2323–2326.

Ruder, S., 2016. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.

Schockaert, S., Smart, P.D., and Twaroch, F.A., 2011. Generating approximate region boundaries from heterogeneous spatial information: An evolutionary approach. *Information Sciences*, 181 (2), 257–283.

Stock, K., *et al.*, 2022. Detecting geospatial location descriptions in natural language text. *International Journal of Geographical Information Science*, 36 (3), 547–584.

Vasardani, M., Winter, S., and Richter, K.F., 2013. Locating place names from place descriptions. *International Journal of Geographical Information Science*, 27 (12), 2509–2532.

Wang, Y., *et al.*, 2022. A spatiotemporal graph attention network based on synchronization for epileptic seizure prediction. *IEEE Journal of Biomedical and Health Informatics*.

Wieczorek, J., Guo, Q., and Hijmans, R., 2004. The point-radius method for georeferencing locality descriptions and calculating associated uncertainty. *International Journal of Geographical Information Science*, 18 (8), 745–767.

Winter, S., *et al.*, 2019. Infrastructure-independent indoor localization and navigation. *ACM Computing Surveys*, 52 (3), 1–24.

Wolter, D. and Wallgrün, J.O., 2012. *Qualitative spatio-temporal representation and reasoning: Trends and future directions*. IGI Global, Ch. Qualitative Spatial Reasoning for Applications: New Challenges and the SparQ Toolbox, 336–362.

Worboys, M.F. and Duckham, M., 2004. *GIS - a computing perspective (2. ed.)*. Taylor & Francis.

Wu, H., *et al.*, 2011. Monitoring and evaluating the quality of web map service resources for optimizing map composition over the internet to support decision making. *Computers & Geosciences*, 37 (4), 485–494.

Wu, X., *et al.*, 2019. A fuzzy formal concept analysis-based approach to uncovering spatial hierarchies among vague places extracted from user-generated data. *International Journal of Geographical Information Science*, 33 (5), 991–1016.

Yin, W. and Ma, Z., 2019. LE & LLE regularized nonnegative tucker decomposition for clustering of high dimensional datasets. *Neurocomputing*, 364, 77–94.

Zhan, Y., *et al.*, 2020. Multi-task compositional network for visual relationship detection. *International Journal of Computer Vision*, 128, 2146–2165.