The Koja Web Mapping Application for Context-sensitive Natural Language Spatial Querying

Niloofar Aflaki[1,*], Kristin Stock[1], Christopher B. Jones[2], Hans Guesgen[1], Yukio Fukuzawa[1] and Jeremy Morley[3]

[1]Massey Geoinformatics Collaboratory, Massey University, Auckland, New Zealand
[2]Cardiff University, United Kingdom
[3]Ordnance Survey, United Kingdom

Abstract
The locations of objects are often described in natural language relative to some other object using vague and context-sensitive spatial relation terms (e.g. theatre near Trafalgar Square). Koja is a web map application that predicts the distance between a location and reference object based on the spatial relation term specified by the user and language and contextual features. That distance is used to retrieve objects of the specified type within a range of the distance. They are displayed through a map interface to make the process more intuitive and user-friendly.

Keywords
Web map, location description, linguistic features, spatial relation term, spatial preposition

1. Introduction

In natural language, people often describe location using locative expressions [1] that include three key elements: the place or object the location of which is being described (known as the locatum or trajector), a place which is the reference object (known as the relatum or landmark) and a relationship between these two (spatial relation term, often a preposition), e.g., *theatre* [locatum] *near* [spatial relation term] *Trafalgar Square* [relatum]. Koja[1] is an interactive web map which retrieves a set of candidate objects that correspond to the described locatum type, using machine learning to predict the distance between the locatum and relatum based on a range of linguistic and contextual features from the three elements in a given expression.

Web maps have been created to answer a wide range of research questions. Environmental applications such as ClimateCharts, which provides an interactive web map showing temperature and precipitation of most places on earth [2], are common, and enable researchers to explore various physical interactions. Web mapping applications that incorporate individual

GeoExT 2023: First International Workshop on Geographic Information Extraction from Texts at ECIR 2023, April 2, 2023, Dublin, Ireland
*Corresponding author.
niloofar.aflaki@massey.ac.nz (N. Aflaki); kristin.stock@massey.ac.nz (K. Stock); JonesCB@cardiff.ac.uk (C. B. Jones); h.guesgen@massey.ac.nz (H. Guesgen); y.fukuzawa@massey.ac.nz (Y. Fukuzawa); jeremy.morley@os.uk (J. Morley)
0000-0002-4643-7692 (N. Aflaki); 0000-0002-5828-6430 (K. Stock); 0000-0002-5069-8203 (Y. Fukuzawa)

[1]https://koja.io.ac.nz/

perspectives have also been developed. For example, eImage which focuses on the selection of an area based on user interests to understand how people feel about different areas [3]. However, web maps that allow users to find the location of places in response to a query rarely take account of vague and context sensitive terms, and in production systems such as Google Maps[2] and OpenStreetMap[3], places are most commonly found simply by their name, or a type of feature (e.g., hairdressers). While it is possible to include spatial relation terms (e.g., in, near, outside, north of) in Google Maps searches, they have little impact on the results returned. There have been many studies of the use of spatial relational language, e.g., [1, 4, 5, 6, 7, 8, 9], some of which focus specifically on contextual factors, e.g., [10]. The concept of spatial templates was introduced in [6] where the template models the applicability or acceptance of locations relative to a reference object, for a given spatial relation. Only limited progress has been made in automating the interpretation of spatial language in order to model or predict the locations referred to in geo-spatial locative expressions e.g. [11, 12, 13], though some progress has also been made in 'table-top' space, such as [14] in which spatial templates were modelled with deep learning. Some notable examples of considering uncertainty in the interpretation of multiple spatial relationships in locative expressions include [15, 16, 17] which focus on the geometric characteristics of objects, with limited attention to the geographical types. In the GeoLocate[4] service complex locality descriptions are interpreted by determining displacements from a reference object (relatum) using a mainly rule-based approach, but it does not take context into account and is limited in the range of spatial relations that can be interpreted. In an early use of multiple spatial templates to predict geo-spatial location [13] predicted location but did not consider feature types and only distinguished between urban and rural contexts, while [12] interpreted relative location descriptions taking account of the contextual factors of place size and prominence. Logan et al. [6] predicted distance and direction based on spatial relation term and relatum, using contextual factors but did not consider an explicit locatum. An example of an experimental search system interface that allows the user to specify one of a small set of spatial relations is that of [18] in which a search buffer adapts to the size of the reference object, taking account of predetermined parameters relating to the locatum and relatum types, and enabling search expansion if no object is found initially. The method adopted here differs from previous studies in retrieving locations of candidate locata, for a query that specifies locatum type, spatial relation term selected from a fairly large number of options, and reference object, based on a regression model that is trained on the combination of spatial relation and context features.

2. Koja Web Map

As an interactive web map, Koja uses location and language features in a machine learning model to predict the distance between locatum and relatum for particular spatial relational terms, in order to select a set of possible features of specified type that lie within a range of the inferred distance from the relatum. It is intended as a simple prototype for future spatial search

[2]https://www.google.com/maps/
[3]https://www.openstreetmap.org/
[4]https://www.geo-locate.org/

systems that attempt to interpret the meaning of relative spatial relational terminology in a query. We used a regression model for this prediction, with SMO a regression version of support vector machine (SVM). The training data contains about 700 location descriptions extracted from Geograph [5] and Foursquare [6] websites which include locatum, geospatial preposition and relatum. For a structured query that specifies locatum type, spatial relation and a named relatum, the derived input features for the regression model include the following:

- GloVe embeddings [7] of the feature types of locatum and relatum.
- The characteristics of the locatum and relatum including geometry type (point, line, polygon), scale, area, elongation and whether it is liquid or solid. These characteristics can affect the interpretation of a spatial relation term (e.g., "church beside post office" vs "road beside river").
- The density of objects in the area and surrounds.
- The semantics of the spatial relation term based on matching the query spatial relation with spatial relations in training data. We exploit semantic similarity of spatial relation terms to support improved learning from training data with similar spatial relations (e.g., "beside", "next to").

The spatial relation term in combination with the other features are used by the classifier to infer a distance which will vary for an individual spatial relation depending on the types of feature (e.g., "park bench near post office" vs "village near London") [19]. The distance is converted to a buffer zone to retrieve objects of the specified locatum feature type that lie inside that zone. Koja supports search based on the specified feature type of the locatum that is selected from OpenStreetMap types, a choice of 24 different spatial relation terms (e.g. "*near*", "*on*", "*at*", "*beside*", "*opposite*") [20] and, for the reference objects, named places in the City of London, UK. It is implemented with Leaflet with OpenStreetMap as a base map, and uses models trained on a set of <locatum-spatial relation term-relatum> triples extracted from Geograph and Foursquare.

The interface allows the user to select the relatum by typing the first three characters of the name and then selecting from the available places, and then to select the spatial relation term and locatum type from a drop down list (again, the initial letters can be typed in to narrow the selection). Once these values are selected and the query submitted, the system determines the input features and runs them against the model created from training data to predict the distance that best reflects the specific context of those three terms.

Figure 1 shows the results of the query *theatre near Trafalgar Square.* The doughnut shape shows the area that the locatum might be located in, calculated using the distance predicted by the machine learning model for the expression (in this case, 366.86m), with a margin of error on either side (based on the mean absolute error divided by average of all distances across all the expressions in the dataset expressed as a percentage). The orange circles identify the theatres that fall within the doughnut, and thus that may be considered to fulfil the query *theatre near Trafalgar Square.*

[5] https://www.geograph.org.uk/
[6] https://foursquare.com/
[7] https://nlp.stanford.edu/projects/glove/

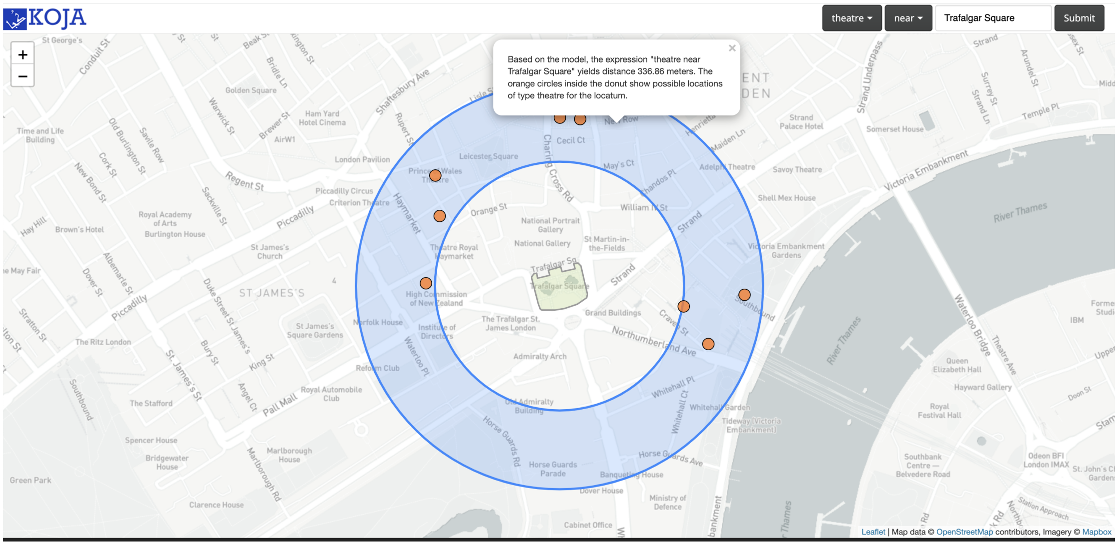

Figure 1: Predicted locations for *theatre near Trafalgar Square*

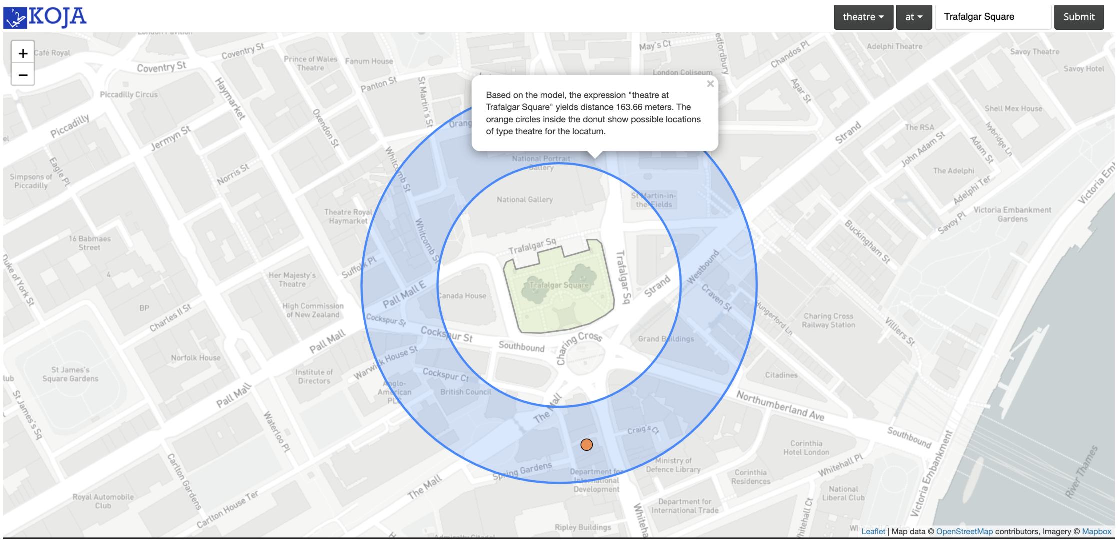

Figure 2: Predicted distance and possible locations for the expression *theatre at Trafalgar Square*

This distance is different for each selection. Fig 2 shows the results from the query *theatre at Trafalgar Square*. As can be seen, the distance that is predicted for the expression using *at* (163.66m) is much smaller than for that using *near*, since the preposition *at* usually refers to a closer distance than *near*.

Figure 3 shows results of the query *theatre near St-Martin-in-the-Fields*, illustrating the effect of different context. St-Martin-in-the-Fields is a church, quite close to Trafalgar Square, but the predicted distance is 149.20m, much smaller than for Trafalgar Square. This is because

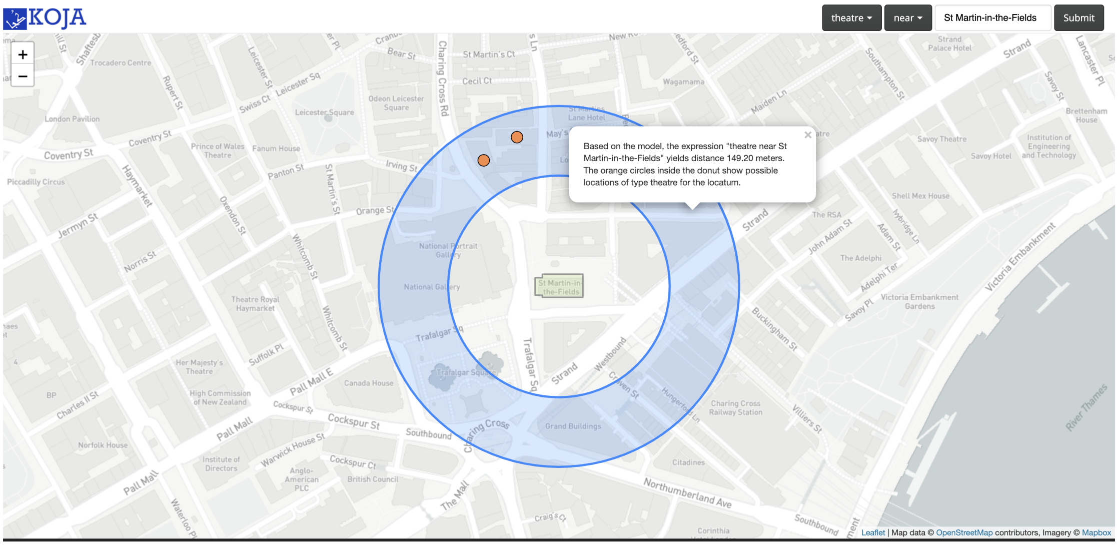

Figure 3: Predicted distance and possible locations for the expression *theatre near St Martin-in-the-Fields*.

the characteristics of the relatum, which is a church (St-Martin-in-the-Fields) differ from the characteristics of the relatum of Trafalgar Square (in Figure 1), which is a square.

3. Conclusions and Future Work

The Koja web mapping application offers an interactive interface for structured spatial queries that incorporates context into the interpretation of vague spatial relation terms, providing distance predictions that use machine learning models with multiple semantic and linguistic input features. It was trained using examples of georeferenced locative expressions from the volunteered data sources Geograph and Foursquare. This provides some significant advances beyond current web mapping applications that either do not cater for spatial relation terms, or use much simpler models for them.

In future work, the authors plan to work with unstructured queries and to further advance the models of spatial language by exploiting deep learning approaches, as well as using richer sets of training data that reflect the geographic context that humans use when formulating natural language location descriptions.

References

[1] A. Herskovits, Language and spatial cognition, Cambridge university press Cambridge, 1986.

[2] L. Zepner, P. Karrasch, F. Wiemann, L. Bernard, ClimateCharts.net – an interactive climate analysis web platform, International Journal of Digital Earth 14 (2020) 338–356. doi:10.1080/17538947.2020.1829112.

[3] M. S. Barros, A. Degbelo, G. Filomena, Evaluative image 2.0: A web mapping approach to capture people's perceptions of a city, Transactions in GIS 26 (2021) 1116–1139. doi:10.1111/tgis.12867.

[4] K. R. Coventry, S. C. Garrod, Saying, Seeing and Acting, Psychology Press, 2004. doi:10.4324/9780203641521.

[5] A. Tyler, V. Evans, The semantics of English prepositions: Spatial scenes, embodied meaning, and cognition, Cambridge University Press, 2003.

[6] S. C. Levinson, LANGUAGE AND SPACE, Annual Review of Anthropology 25 (1996) 353–382. doi:10.1146/annurev.anthro.25.1.353.

[7] K. Stock, C. B. Jones, T. Tenbrink, Speaking of location: a review of spatial language research, Spatial Cognition and Computation 22 (2022) 185–224. doi:10.1080/13875868.2022.2095275.

[8] P. F. Fisher, T. M. Orf, An investigation of the meaning of near and close on a university campus, Computers, Environment and Urban Systems 15 (1991) 23–35. doi:10.1016/0198-9715(91)90043-d.

[9] V. Robinson, Individual and multipersonal fuzzy spatial relations acquired using human–machine interaction, Fuzzy Sets and Systems 113 (2000) 133–145. doi:10.1016/s0165-0114(99)00017-2.

[10] X. Yao, J.-C. Thill, How far is too far? - a statistical approach to context-contingent proximity modeling, Transactions in GIS 9 (2005) 157–178. doi:10.1111/j.1467-9671.2005.00211.x.

[11] M. M. Hall, C. B. Jones, Generating geographical location descriptions with spatial templates: a salient toponym driven approach, International Journal of Geographical Information Science 36 (2021) 55–85. doi:10.1080/13658816.2021.1913498.

[12] H. Chen, S. Winter, M. Vasardani, Georeferencing places from collective human descriptions using place graphs, Journal of Spatial Information Science (2018). doi:10.5311/josis.2018.17.417.

[13] M. M. Hall, P. D. Smart, C. B. Jones, Interpreting spatial language in image captions, Cognitive Processing 12 (2010) 67–94. doi:10.1007/s10339-010-0385-5.

[14] M. Malinowski, M. Fritz, A pooling approach to modelling spatial relations for image retrieval and annotation, arXiv preprint arXiv:1411.5190 (2014).

[15] Q. Guo, Y. Liu, J. Wieczorek, Georeferencing locality descriptions and computing associated uncertainty using a probabilistic approach, International Journal of Geographical Information Science 22 (2008) 1067–1090. doi:10.1080/13658810701851420.

[16] Y. Liu, Q. H. Guo, J. Wieczorek, M. F. Goodchild, Positioning localities based on spatial assertions, International Journal of Geographical Information Science 23 (2009) 1471–1501. doi:10.1080/13658810802247114.

[17] P. Doherty, Q. Guo, Y. Liu, J. Wieczorek, J. Doke, Georeferencing incidents from locality descriptions and its applications: a case study from yosemite national park search and rescue, Transactions in GIS 15 (2011) 775–793. doi:10.1111/j.1467-9671.2011.01290.x.

[18] G. Fu, C. B. Jones, A. I. Abdelmoty, Ontology-based spatial query expansion in information retrieval, in: Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2005, pp. 1466–1482. doi:10.1007/11575801_33.

[19] N. Aflaki, K. Stock, C. B. Jones, H. Guesgen, J. Morley, Y. Fukuzawa, What do you mean

you're in trafalgar square? comparing distance thresholds for geospatial prepositions, in: 15th International Conference on Spatial Information Theory (COSIT 2022), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.COSIT.2022.1.
[20] N. Aflaki, K. Stock, C. B. Jones, H. Guesgen, J. Morley, An empirical study of the semantic similarity of geospatial prepositions and their senses, Spatial Cognition and Computation 23 (2022) 132–176. doi:10.1080/13875868.2022.2111683.