

UC Santa Barbara

Spatial Data Science Symposium 2022 Short Paper Proceedings

Title

Geoparsing comments from Reddit to extract mental place connectivity within the United Kingdom

Permalink

<https://escholarship.org/uc/item/4wg37921>

Authors

Berragan, Cillian
Singleton, Alex
Calafiore, Alessia
[et al.](#)

Publication Date

2022-09-09

DOI

10.25436/E28C7R

Peer reviewed

Geoparsing comments from Reddit to extract mental place connectivity within the United Kingdom

Cillian Berragan¹[0000-0003-2198-2245], Alex Singleton¹[0000-0002-2338-2334], Alessia Calafiore²[0000-0002-5953-2891], and Jeremy Morley³[0000-0002-3658-8796]

¹ University of Liverpool, Liverpool, United Kingdom
{c.berragan,ucfnale}@liverpool.ac.uk

² University of Edinburgh, Edinburgh, United Kingdom
acalafio@ed.ac.uk

³ Ordnance Survey, Southampton, United Kingdom
jeremy.morley@os.uk

Abstract. Place connectivity is explored between geographic locations extracted from comments on Reddit. Unlike formally structured geographic data, this corpus of unstructured text provides connections derived from co-occurring locations, capturing subconscious links between them, alongside inherent biases. Our work demonstrates the ability to link locations mentioned by unique users, building ‘mental’ place connections for over 50,000 unique locations in the United Kingdom. Sentiment regarding locations is compared against their levels of connectivity, demonstrating that user opinions regarding locations are likely drivers in mental place connectivity.

Keywords: Geoparsing · Reddit · Mental place connectivity · Sentiment

DOI: <https://doi.org/10.25436/E28C7R>

1 Introduction

Connectivity between places may be explored through the physical movement of individuals, using population movement data like transport records [1-4], or GPS information through mobile phone data [5, 6]. Due to the advent of ‘Volunteered Geographic Information’ (VGI) [7], these connections may also be explored through geotagged social media posts [8-10], with results that mirror true population movements [10, 11].

Literature discussing the role of human cognition in constructing mental images of cities [12], and how they can be represented through mental maps [13], shows that the way humans conceive spatial structures and place relationships are substantially entrenched in individuals’ experiences and geographic knowledge, which only partially derive from movements. In particular, while movements are constrained by time and euclidean distance in geographic space [14,

15], representational spaces expressed in mental maps do not necessarily correlate with these spatio-temporal boundaries.

New forms of data, especially social media and text data, offer novel opportunities to explore place connections, emerging from peoples naïve and experiential geographic knowledge. Recent studies have investigated how digital social networks, i.e. Facebook friendships [16], may be explored in this regard. While, *geo-semantic relatedness* outlines the ability to quantify relationships between geographic terms in text [17], with work applying this to co-occurring city names found in news articles, social media, and general web pages [18–21].

Our paper considers the use of a novel corpus of text data from comments on the online social media website **Reddit** as a source of VGI. A task-specific geoparsing pipeline is first used to identify place names⁴ related to the United Kingdom and resolve them to geographic coordinates [22], at a geographic resolution higher than is typically explored through geoparsing. We derive connectivity between each location in our corpus based on the number of times two distinct locations co-occur, normalised by the total number of users that mention each location. The context for co-occurrences is derived from the total collection of comments submitted by each unique user, meaning every location mentioned by a single user is treated as co-occurring, and exhibiting some implicit connectivity. A full interactive map showing place connections is available through **Unfolded.ai**.

With the connectivity between places established, we consider the ability to derive explainable characteristics from the text, to determine *why* different levels of place connectivity occur. While traditional connections between places may be influenced by factors like transport availability [2], the alternative mental place connections derived through our paper may be more heavily influenced by a subconscious bias. Connectivity is therefore examined against sentiment, expected to highlight these biases, alongside a standard measure of relative material deprivation (limited to England), through the **Indices of Multiple Deprivation** (IMD), an alternative measure that would be expected to affect traditional place connectivity.

2 Methodology

2.1 Data sources

Reddit is a public discussion, news aggregation social network, among the top 20 most visited websites in the United Kingdom. As of 2020, Reddit had around 430 million active monthly users, comparable to the number of Twitter users [23, 24]. Reddit is divided into separate independent *subreddits* each with specific topics of discussion, where *users* may submit *posts* which each have dedicated nested

⁴ In our work we consider ‘place names’ to be ambiguous, without coordinates, while ‘locations’ are place names that have attributed coordinates. We use ‘*place* connectivity’ to capture the vagueness involved in these connections.

conversation threads that users can add *comments* to. In total there are 213 subreddits that relate to ‘places’ within the United Kingdom⁵. For each subreddit, every single historic comment was retrieved using the [Pushshift](#) Reddit archive [25]. In total 8,295,591 comments were extracted, submitted by 492,123 unique users, between 2011-01-01 and 2022-04-17.

To train a model to identify place names from comments, the WNUT-17 corpus was used [26], keeping only ‘location’ labels. In total this corpus covers 5,690 individual documents from Reddit, Twitter, YouTube, and StackExchange.

Two gazetteers were selected to geocode place names, chosen to be UK centric, and at a high resolution. We were specifically interested in a gazetteer that did not include country names external to the UK, but included fine-grained named locations like street names. For our gazetteer we combined [OS Open Names](#), and non-settlements from the [Gazetteer of British Place Names](#).

2.2 Geoparsing

A custom named entity recognition (NER) model was built using a RoBERTa-based transformer language model, pre-trained using Twitter data⁶, as this architecture has given good results on the WNUT-17 corpus [27].

With place names identified, we developed a method for attributing each name to a single set of geographic coordinates. Place names typically appear multiple times in gazetteers, especially when grounding fine-grained locations like street names, meaning a disambiguation method is required. We disambiguate place names by finding their minimum distance to a collection of contextual locations. Contextual locations in this case refer to all gazetteer entries matching place names that appear in sentences with this target place name, in the same subreddit. This works under the assumption that each unique place name in a single subreddit is likely to refer to the same location, and that locations mentioned in surrounding text are likely close together [28].

Sentiment was attributed to each sentence in our corpus containing a place name, using an existing fine-tuned sentiment classification transformer model⁷. Each place name identified by our NER model was assigned sentiment based on its context sentence.

2.3 Measuring place connectivity

Our place connectivity methodology considers the co-occurrence between each place mentioned by every unique user in our corpus. The following equation described by [10] outlines this concept:

$$PCI_{ij} = \frac{\mathbf{S}_{ij}}{\sqrt{\mathbf{S}_i \mathbf{S}_j}} \quad i, j \in [1, n]$$

⁵ https://www.reddit.com/r/unitedkingdom/wiki/british_subreddits

⁶ <https://huggingface.co/cardiffnlp/twitter-roberta-base>

⁷ <https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest>

PCI_{ij} is the place connectivity index between locations i and j , \mathbf{S}_{ij} is the total number of users that mention *both* locations (i.e. the *intersect* in set theory). This is normalised, given locations with higher populations are expected to be mentioned by a larger number of users, using $\sqrt{\mathbf{S}_i \mathbf{S}_j}$, the total number of users mentioning \mathbf{S}_i multiplied by the total number of users mentioning \mathbf{S}_j , taking the square root. n is the total number of unique locations found in all comments.

3 Results

To assess the performance of our trained NER model we manually annotated 498 randomly selected Reddit comments with place names. On this test dataset, our model achieved an F_1 performance of 0.845, a recall of 0.91 and precision of 0.85, above the expected performance on the WNUT test dataset. Similarly, we annotated 200 comments as neutral, positive or negative sentiment, and found the pre-built sentiment model performed as expected [29], with an F_1 , recall and precision of 0.70.

In total 26% of all comments contained at least one place name, 4,475,800 place names were identified, with 2,816,072 (63%) attributed with a set of coordinates. From these locations, 57,682 were found to be unique.

3.1 Place connectivity

Figure 1 demonstrates mental place connections when aggregated to Local Authority Districts (LAD). This aggregation is performed by treating a place equivalent to the LAD it is contained in. Aggregation allows for a higher level view of place connections to be observed, for example combining several points of interest within a major city produces a single strong connection, rather than several weaker connections. Figure 1 (a) shows connections within England and Wales. Clusters emerge where isolated urban areas and their surrounding LADs exhibit strong connectivity, for example around London, Liverpool and Manchester, and Bristol. Notably Wales appears reasonably isolated from England, but inter-connectivity between LADs is strong, including in more rural areas. This isolation between Wales and the rest of the UK has also been observed through semantic analysis of Twitter [8].

Figure 1 (b) shows connections within Scotland, these are much stronger generally than the rest of the UK, and highly inter-connected. These connections are also less isolated compared with Wales, with links between Glasgow, Edinburgh and London, as well as strong links with Durham and Newcastle. Links in Scotland are not restricted to urban areas, with strong connections between the rural Highland LAD and all major urban centres. These connections may however be influenced by the varying levels of ambiguity between English place names compared with names in remote Scotland. For example in the ‘Highland’ LAD, 27% of place names are ambiguous, compared with 40% in Manchester, meaning toponym disambiguation is likely more accurate in the Highlands.

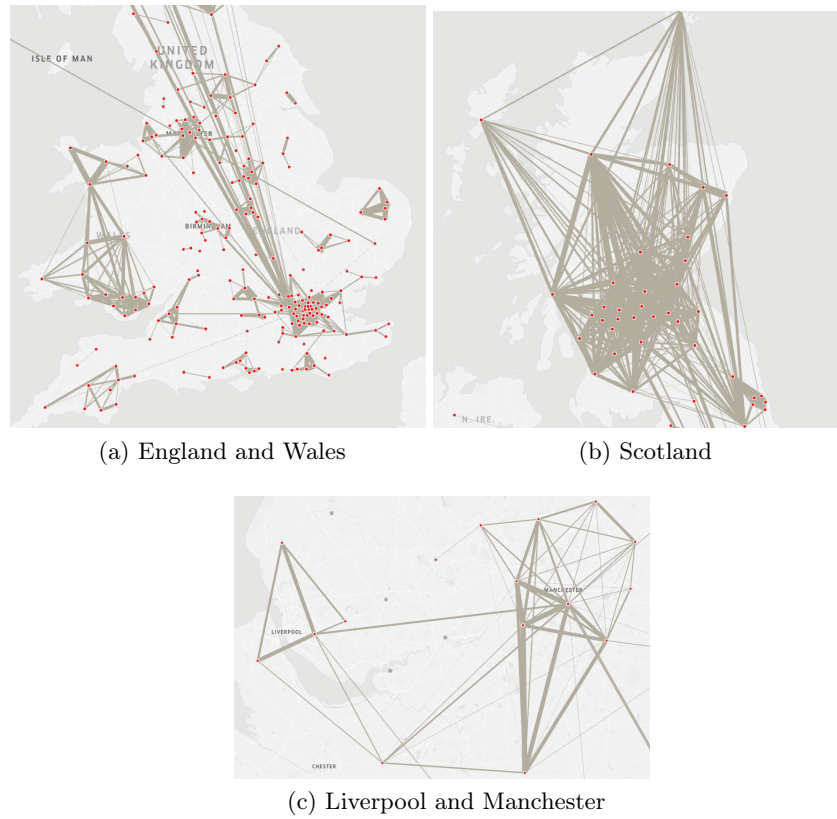


Fig. 1: Place connections aggregated to Local Authority District, edge size weighted by PCI values, full interactive map available through [Unfolded](#)




Fig. 2: Correlation between PCI and target (a) IMD, (b) Sentiment, and (c) Distance for unaggregated connections between locations. Connections with fewer than 100 shared users have been excluded. Values aggregated into 200 bins for readability.

Figure 1 (c) shows a highly urbanised area of England, with two major cities; Liverpool and Manchester. While Manchester links directly with Liverpool, this figure does not appear to reflect the contiguous urban area that links these two cities [30], given intermediate LADs are not connected. The link between these two cities appears direct from a mental perspective, meaning intermediate urban zones are perceived as less connected through our index.

While there are a multitude of factors that affect this mental place connectivity, typical connectivity models explore the effects of distance decay [1, 3, 10, 16, 18]. As past work has found, with both social and physical connections between places, our mental PCI does experience distance decay, with a medium strength negative correlation between PCI and the log of distance ($R = -0.55; p = 0.0$; Figure 2 (c)). Additionally we explore how sentiment and deprivation may influence PCI values. Figure 2 (a) shows a significant weak correlation between PCI and sentiment ($R = 0.18; p = 0.0$), while Figure 2 (b) shows a very weak negative correlation between PCI and the IMD Score ($R = -0.04; p = 0.0$).

Lower connectivity for places with more negative sentiment may be influenced by factors like ‘fear of crime’ [31], and given deprivation appears less influential, these perceptions likely capture information that is not directly quantifiable through traditional data sources. Certain areas in the United Kingdom have been shown to have ‘reputations’ [32], meaning they are known to be viewed negatively both inside and outside their occupants, particularly between the North and South of England [13].

4 Conclusion

Our paper demonstrates the use of Reddit to explore mental place connectivity from a novel source of data, without the reliance on explicit geographic information. Despite Reddit being both more pseudo-anonymous than Twitter (users more frequently use pseudonyms), and without explicit social connections between users (unlike Twitter, followers are not emphasised), mental place connectivity derived through Reddit still exhibits the distance decay that is expected when observing place connectivity [1, 10, 16]. While Reddit users are not completely representative of the UK population, the volume of unique users that contribute to this corpus are expected to better reflect a more accurate, general view of mental connectivity, compared to alternative text sources like news articles [18].

Future work may consider a further exploration of sentiment between subreddits, London for example may be viewed differently based on subreddit communities focussed on North of England, compared with the South [13, 33]. There is also the opportunity to explore relationships between locations or communities through their semantic typology; clustering through Latent Dirichlet Allocation (LDA) [34], or finding the cosine similarity between derived lexicons [8]. Locations identified from within these communities also likely represent urban areas of interest which may be derived based on their frequency of mentions [35], or semantic regions that reflect mental perceptions of places [36].

References

1. Yang, Y., Li, D., Li, X.(Robert).: Public Transport Connectivity and Intercity Tourist Flows. *Journal of Travel Research*. 58, 25–41 (2019). <https://doi.org/10.1177/0047287517741997>.
2. Allard, R.F., Moura, F.: The Incorporation of Passenger Connectivity and Intermodal Considerations in Intercity Transport Planning. *Transport Reviews*. 36, 251–277 (2016). <https://doi.org/10.1080/01441647.2015.1059379>.
3. Gong, J., Li, S., Ye, X., Peng, Q., Kudva, S.: Modelling impacts of high-speed rail on urban interaction with social media in China’s mainland. *Geo-spatial Information Science*. 24, 638–653 (2021). <https://doi.org/10.1080/10095020.2021.1972771>.
4. Farber, S., Li, X.: Urban sprawl and social interaction potential: An empirical analysis of large metropolitan regions in the United States. *Journal of Transport Geography*. 31, 267–277 (2013). <https://doi.org/10.1016/j.jtrangeo.2013.03.002>.
5. Lin, J., Wu, Z., Li, X.: Measuring inter-city connectivity in an urban agglomeration based on multi-source data. *International Journal of Geographical Information Science*. 33, 1062–1081 (2019). <https://doi.org/10.1080/13658816.2018.1563302>.
6. SafeGraph: Places Data Curated for Accurate Geospatial Analytics | SafeGraph, (2022).
7. Goodchild, M.F.: Citizens as sensors: The world of volunteered geography. *GeoJournal*. 69, 211–221 (2007). <https://doi.org/10.1007/s10708-007-9111-y>.
8. Arthur, R., Williams, H.T.P.: The human geography of Twitter: Quantifying regional identity and inter-region communication in England and Wales. *PLOS ONE*. 14, e0214466 (2019). <https://doi.org/10.1371/journal.pone.0214466>.
9. Ostermann, F.O., Huang, H., Andrienko, G., Andrienko, N., Capineri, C., Farkas, K., Purves, R.S.: EXTRACTING AND COMPARING PLACES USING GEO-SOCIAL MEDIA. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. II-3/W5, 311–316 (2015). <https://doi.org/10.5194/isprsannals-II-3-W5-311-2015>.
10. Li, Z., Huang, X., Ye, X., Jiang, Y., Martin, Y., Ning, H., Hodgson, M.E., Li, X.: Measuring global multi-scale place connectivity using geotagged social media data. *Scientific Reports*. 11, 14694 (2021). <https://doi.org/10.1038/s41598-021-94300-7>.
11. Kuchler, T., Russel, D., Stroebel, J.: The Geographic Spread of COVID-19 Correlates with the Structure of Social Networks as Measured by Facebook. National Bureau of Economic Research, Cambridge, MA (2020). <https://doi.org/10.3386/w26990>.

12. Lynch, K.: *The Image of the City*. MIT Press (1964).
13. Gould, P., White, R.: *Mental Maps*. Taylor and Francis, Hoboken (1986).
14. Miller, H.J.: Time geography. *Handbook of Behavioral and Cognitive Geography*. 74–94 (2018). <https://doi.org/10.4337/9781784717544.00011>.
15. Patterson, Z., Farber, S.: Potential Path Areas and Activity Spaces in Application: A Review. *Transport Reviews*. 35, 679–700 (2015). <https://doi.org/10.1080/01441647.2015.1042944>.
16. Bailey, M., Cao, R., Kuchler, T., Stroebel, J., Wong, A.: Social Connectedness: Measurement, Determinants, and Effects. *Journal of Economic Perspectives*. 32, 259–280 (2018). <https://doi.org/10.1257/jep.32.3.259>.
17. Ballatore, A., Bertolotto, M., Wilson, D.C.: An evaluative baseline for geo-semantic relatedness and similarity. *GeoInformatica*. 18, 747–767 (2014). <https://doi.org/10.1007/s10707-013-0197-8>.
18. Hu, Y., Ye, X., Shaw, S.-L.: Extracting and analyzing semantic relatedness between cities using news articles. *International Journal of Geographical Information Science*. 31, 2427–2451 (2017). <https://doi.org/10.1080/13658816.2017.1367797>.
19. Ye, X., Gong, J., Li, S.: Analyzing Asymmetric City Connectivity by Toponym on Social Media in China. *Chinese Geographical Science*. 31, 14–26 (2021). <https://doi.org/10.1007/s11769-020-1172-6>.
20. Liu, Y., Wang, F., Kang, C., Gao, Y., Lu, Y.: Analyzing Relatedness by Toponym Co-Occurrences on Web Pages: Analyzing Relatedness by Toponym Co-Occurrences on Web Pages. *Transactions in GIS*. 18, 89–107 (2014). <https://doi.org/10.1111/tgis.12023>.
21. Meijers, E., Peris, A.: Using toponym co-occurrences to measure relationships between places: Review, application and evaluation. *International Journal of Urban Sciences*. 23, 246–268 (2019). <https://doi.org/10.1080/12265934.2018.1497526>.
22. Purves, R.S., Clough, P., Jones, C.B., Hall, M.H., Murdock, V.: *Geographic Information Retrieval: Progress and Challenges in Spatial Search of Text*. now (2018). <https://doi.org/10.1561/1500000034>.
23. Murphy, N.: *Reddit’s 2019 Year in Review - Upvoted*, (2019).
24. Statista: *Most popular social networks worldwide as of January 2022, ranked by number of monthly active users*, (2022).
25. Baumgartner, J., Zannettou, S., Keegan, B., Squire, M., Blackburn, J.: *The Pushshift Reddit Dataset*, <https://arxiv.org/abs/2001.08435>, (2020).

26. Derczynski, L., Nichols, E., van Erp, M., Limsopatham, N.: Results of the WNUT2017 Shared Task on Novel and Emerging Entity Recognition. In: Proceedings of the 3rd Workshop on Noisy User-generated Text. pp. 140–147. Association for Computational Linguistics, Copenhagen, Denmark (2017). <https://doi.org/10.18653/v1/W17-4418>.
27. Barbieri, F., Camacho-Collados, J., Neves, L., Espinosa-Anke, L.: **Tweet-Eval: Unified Benchmark and Comparative Evaluation for Tweet Classification**. arXiv:2010.12421 [cs]. (2020).
28. Kamaloo, E., Rafiei, D.: A Coherent Unsupervised Model for Toponym Resolution. In: Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18. pp. 1287–1296. ACM Press, Lyon, France (2018). <https://doi.org/10.1145/3178876.3186027>.
29. Loureiro, D., Barbieri, F., Neves, L., Anke, L.E., Camacho-Collados, J.: TimeLMs: Diachronic Language Models from Twitter, <https://arxiv.org/abs/2202.03829>, (2022).
30. Dembski, S.: Structure and imagination of changing cities: Manchester, Liverpool and the spatial in-between. *Urban Studies*. 52, 1647–1664 (2015). <https://doi.org/10.1177/0042098014539021>.
31. Solymosi, R., Buil-Gil, D., Vozmediano, L., Guedes, I.S.: Towards a Place-based Measure of Fear of Crime: A Systematic Review of App-based and Crowdsourcing Approaches. *Environment and Behavior*. 53, 1013–1044 (2021). <https://doi.org/10.1177/0013916520947114>.
32. Kearns, A., Kearns, O., Lawson, L.: Notorious Places: Image, Reputation, Stigma. The Role of Newspapers in Area Reputations for Social Housing Estates. *Housing Studies*. 28, 579–598 (2013). <https://doi.org/10.1080/02673037.2013.759546>.
33. Jewell, H.M.: *The North-south Divide: The Origins of Northern Consciousness in England*. Manchester University Press (1994).
34. Gao, S., Janowicz, K., Couclelis, H.: Extracting urban functional regions from points of interest and human activities on location-based social networks. *Transactions in GIS*. 21, 446–467 (2017). <https://doi.org/10.1111/tgis.12289>.
35. Chen, M.: Understanding the dynamics of urban areas of interest through volunteered geographic information. 21 (2019). <https://doi.org/10.1007/s10109-018-0284-3>.
36. Gao, S., Janowicz, K., Montello, D.R., Hu, Y., Yang, J.-A., McKenzie, G., Ju, Y., Gong, L., Adams, B., Yan, B.: A data-synthesis-driven method for detecting and extracting vague cognitive regions. *International Journal of Geographical Information Science*. 31, 1245–1271 (2017). <https://doi.org/10.1080/13658816.2016.1273357>.