

Linking Streets in OpenStreetMap to Persons in Wikidata

Daria Gurtovoy
Universität Bonn
Germany
s7dagurt@uni-bonn.de

Simon Gottschalk
L3S Research Center, Leibniz Universität Hannover
Germany
gottschalk@L3S.de

ABSTRACT

Geographic web sources such as OpenStreetMap (OSM) and knowledge graphs such as Wikidata are often unconnected. An example connection that can be established between these sources are links between streets in OSM to the persons in Wikidata they were named after. This paper presents *StreetToPerson*, an approach for connecting streets in OSM to persons in a knowledge graph based on relations in the knowledge graph and spatial dependencies. Our evaluation shows that we outperform existing approaches by 26 percentage points. In addition, we apply *StreetToPerson* on all OSM streets in Germany, for which we identify more than 180,000 links between streets and persons.

CCS CONCEPTS

• **Information systems** → *Spatial-temporal systems; Information integration.*

KEYWORDS

Knowledge graphs, OpenStreetMap, Wikidata, Street names

ACM Reference Format:

Daria Gurtovoy and Simon Gottschalk. 2022. Linking Streets in OpenStreetMap to Persons in Wikidata. In *Companion Proceedings of the Web Conference 2022 (WWW '22 Companion)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3487553.3524267>

1 INTRODUCTION

Streets are often named after famous or distinguished individuals who may or may not have a direct connection to the specific location. While geographic data sources such as OpenStreetMap (OSM)¹ contain information about streets across the globe and knowledge graphs such as Wikidata contain information about famous individuals and their relations, these two worlds are often not connected. This makes it challenging to connect streets to whom they are named after.

¹OpenStreetMap, OSM and the OpenStreetMap magnifying glass logo are trademarks of the OpenStreetMap Foundation, and are used with their permission. We are not endorsed by or affiliated with the OpenStreetMap Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW '22 Companion, April 25–29, 2022, Virtual Event, Lyon, France.

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9130-6/22/04...\$15.00

<https://doi.org/10.1145/3487553.3524267>

For example, consider the street named “Wilhelmstraße” in the German capital Berlin. After the removal of the suffix “straße” (German for “street”), only the term “Wilhelm” remains. This term can potentially be linked to a variety of persons, be it via their first name (e.g., Wilhelm Busch), their second name (e.g., Friedrich Wilhelm I.) or their last name (e.g., Paul Wilhelm). With the help of Wikidata, we can infer that Friedrich Wilhelm I. was born in Berlin and was a historically important monarch. Both factors may lead to the correct selection of Friedrich Wilhelm I. for the Wilhelmstraße.

The provision of links between streets in OSM to persons in Wikidata can support several use cases including: (i) Bridging the gap between OSM and knowledge graphs: By bringing OSM and Wikidata together, they will mutually benefit from their strengths [10]. Knowledge graphs such as WorldKG [5] and the Nuremberg Address Knowledge Graph [2] are examples of such efforts. (ii) Providing data for touristic applications: Tourists are often interested in exploring the history behind a place that can often be told by people who have a special significance in that place. An existing example is an application for studying the backgrounds behind the *Stolpersteine* in Germany, which are dedicated to victims of the Nazi regime². (iii) Enabling cultural analyses: With street-to-person links available, analyses are enabled about the motives behind naming streets, potentially revealing discrimination of women [9] and the detection of quarters dedicated to specific groups of persons [3].

In contrast to named entity linking approaches on text, no context information is available additional to the street’s name and location. Therefore, the task of linking streets in OSM to persons in Wikidata lies in identifying the correct person based on features such as the person’s popularity and the person’s relatedness to the street’s location. We propose *StreetToPerson*, which first builds an index for retrieving potential candidate persons given a street name. For each such candidate person, we extract a set of features denoting specific characteristics of that person, such as its popularity and its spatial relatedness to the given street, and classify it as correct or not.

In our evaluation, we first show *StreetToPerson*’s superiority over existing baselines using a ground truth extracted from German streets in Wikidata, and we apply *StreetToPerson* to German streets in OSM. *StreetToPerson* reaches a precision of 0.95 on the ground truth and identifies more than 180,000 street-to-person relations.

The remainder of this paper is structured as follows: In Section 2, we present related work. Then, in Section 3, we describe our approach and evaluate it in Section 4. Finally, we provide a conclusion in Section 5.

²<https://stolpersteine.wdr.de/web/en/>

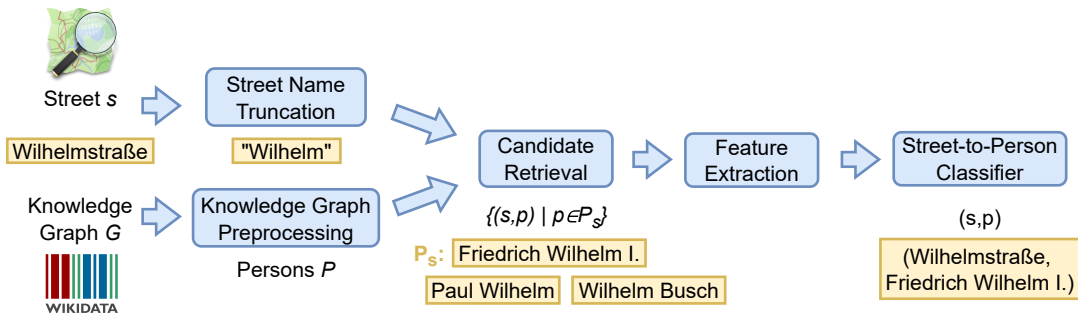


Figure 1: Overview of *StreetToPerson*. Yellow boxes show example values.

2 RELATED WORK

Linking streets to persons is related to named entity linking and connecting geographic data sources with knowledge graphs.

2.1 Named Person Entity Linking

Most related to this paper is the work by Almeida et al. [3], a ranking-based approach to connect streets to the persons they were named after based on a manually defined relevance score. Users then confirm the generated street-to-person pairs in a web interface. Geiss et al. [7] link mentions of persons in a text document to Wikidata using a network of candidate persons. In general, street-to-person linking can be considered a variation of named entity linking to Wikidata [8].

2.2 Connecting OSM with Knowledge Graphs

The task of connecting geographic data sources such as OSM to knowledge graphs has been addressed from different perspectives. Typically, approaches aim at establishing identity links between the different representations of geographic entities and concepts in these sources. For example, [10] proposes a pipeline for link discovery between OSM, Wikidata and DBpedia based on OSM tags, [4] aligns the schema between these sources using an adversarial classifier, and *osm2rdf* converts the whole OpenStreetMap data to RDF triples [1]. In contrast to these approaches, our task deals with persons, a class of entities not present in OSM.

3 APPROACH

The goal of this paper is to link streets to the correct person whom they are named after. Formally, we develop a *street-to-person mapping function* as follows:

Definition 3.1. Street-to-person mapping function. Given a street s and a set of persons P , create a street-to-person mapping function $f(s) \mapsto P$ that identifies the person $p \in P$ after whom the street is named.

Figure 1 illustrates the approach and its components. First, the name of the given street s is truncated to potentially only contain a person’s name. This term is then used in a candidate generation step to retrieve a set of candidate persons P_s out of a set of persons P extracted from Wikidata. For each street-candidate-pair (s, p) , we extract various features and use a street-to-person classifier that determines the candidate with the highest probability as the

correct person for the given street, where the street can be of any type of street in any region. In the example given in Figure 1, the street "Wilhelmstraße" is truncated to "Wilhelm", for which the preprocessed knowledge graph returns the candidates "Paul Wilhelm", "Wilhelm Busch" and "Friedrich Wilhelm I." and others. After extracting all candidate features, the classifier determines the latter to be the best candidate for the given street.

The implementation of this pipeline as well as the generated links are available on GitHub³.

3.1 Knowledge Graph Preprocessing

The data required to retrieve candidates and their features is taken from Wikidata. For efficient access to relevant information, we create the following data sets from Wikidata in a preprocessing step⁴:

- The *ground truth* for training and evaluation contains known streets in Germany that are named after a person. We provide details about this dataset in Section 4.1.
- The *person index* takes a term as input and returns the Wikidata IDs of all people that match it. This person index contains over 4 million names of over 9 million persons.
- The *person occupation index* contains all occupations of a person (e.g., monarch or writer).
- The *person location index* contains relevant locations such as a person’s birthplace.
- The *spatial dependency database* denotes containment relations between locations (e.g., Berlin is located in Germany).

Figure 2 shows relations of Friedrich Wilhelm I. retrieved from the person occupation index and the person location index.⁵

3.2 Street Name Truncation

When analysing a street name, it is essential to differentiate between the part that maps to a person and street affixes (i.e., prefixes and suffixes), such as "street", "road", and "avenue". Given an extensive list of common street affixes, they can be removed from the street name so that only the person name remains, which can then be used to look up matches in the person index database. To achieve this, we gathered the streets of the ground truth and removed all names

³<https://github.com/d-gurtovoy/streetnameLinks>

⁴To extract these data sets, we process the Wikidata dump using the *qwikidata* library (<https://qwikidata.readthedocs.io/>).

⁵Over time, Friedrich Wilhelm I. has been buried in different churches.

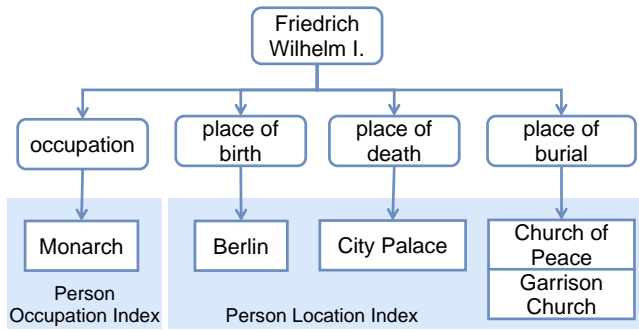


Figure 2: Relations of Friedrich Wilhelm I. from Wikidata.

and aliases from the persons associated with them. The remaining parts of the street names are then manually split into prefixes and suffixes and corrected when necessary. The result is a set of 80 German suffixes and 34 prefixes which we make available⁶. Using this set of affixes, the name of a street s is truncated so that street affixes are removed and the part that potentially maps to a person name remains.

3.3 Candidate Retrieval

The truncated street name is used to query the person index, which returns the Wikidata IDs of matching candidates $p \in P_s$. For instance, the term “Wilhelm” would return the ID of “Friedrich Wilhelm I.” but also of “Wilhelm Busch” among others.

3.4 Feature Extraction

After truncating the street name and retrieving candidates, we extract characteristics of the person p and the spatial relations between p and the street s that serve as features for training a binary classifier. The following 30 features are extracted for every street-candidate-pair (s, p) :

- *Link count*: The number of links pointing to p in the German Wikipedia. If the candidate does not exist in Wikipedia but only in Wikidata, this feature value is set to 0.
- *Name*: These four binary features show which part of the person’s name is contained in the name of s . It can be its full, first or last name or an alias.
- *Occupations*: From the person occupation index, we gather 20 of the most common occupations of people that German streets were named after and introduce 20 binary features to mark if the person p held any of these positions.
- *Spatial relations*: We introduce five numerical features representing whether there is a spatial relation between p and s : “born”, “died”, “buried”, “educated at”, and “work location”.

The *spatial relations* feature values represent to which extent the street and the person’s related location are contained within each other. These feature values range from 0 to 1, with 0 meaning the given location is either unknown or lies outside of Germany and 1 if the location is the street itself. If a person is related to multiple locations for the same relation, we take the location with the highest score.

⁶<https://github.com/d-gurtovoy/streetnameLinks/tree/master/data/affixes>




Figure 3: Example of the spatial dependencies between the street “Wilhelmstraße” and the city Berlin. Arrows denote “located in” relations (e.g., Berlin is located in Germany).

Figure 3 illustrates an example of computing the containment of the street “Wilhelmstraße” and Berlin, the birthplace of Friedrich Wilhelm I. We create a containment chain for both locations from the spatial dependency database and check their overlap. In this example, the “born” relation feature is set to $2/4 = 0.5$.

4 EVALUATION

This section introduces the data used for training and testing and evaluates *StreetToPerson* against three baselines.

4.1 Data

We require a dataset of known person-to-street links to train our approach and evaluate. For a few cases, Wikidata and OSM provide this information. Wikidata contains 4,799 German streets connected to a person via Wikidata’s “named after” property⁷. We use these street-to-person relations as positive examples for training our model. As negative examples, we retrieve up to 50 candidate persons of these streets with the highest link count each.

OSM provides a tag called “name:etymology”⁸, which links streets to persons in Wikipedia⁹. This key provides 19,592 street-to-person relations involving German streets.

4.2 Baselines

We compare *StreetToPerson* to three different baselines:

- **TagMe** [6] (**TagMe**): **TagMe** is a traditional entity linking approach on short text fragments. As it may link a street name to the actual street entity and not the person it is named after, we feed **TagMe** with the street name after applying street name truncation.
- **Popularity Ranking (PopRank)**: In a simpler version of *StreetToPerson*, we replace the street-to-person classifier by taking the person with the highest link count.
- **Relevance Ranking** [3] (**RelRank**): The approach by Almeida et al. described in Section 2 is, to the best of our knowledge, the only existing approach for street-to-person linking. As **RelRank** does not limit its results to person entities, we consider two configurations: **RelRank** (all entities) and **RelRank** (person entities), where we only consider links to person entities.

⁷<https://www.wikidata.org/wiki/Property:P138>

⁸<https://wiki.openstreetmap.org/wiki/Key:name:etymology>

⁹We use Wikidata *sitelinks* to link persons in Wikipedia to Wikidata

4.3 Evaluation of the Classification

We evaluate *StreetToPerson*'s classification performance on the Wikidata ground truth described in Section 4.1 using 10-fold cross-validation. Results are shown in Table 1. *StreetToPerson* clearly outperforms the other approaches and reaches a precision of 0.95, 26 percentage points more than the second-best baseline, **PopRank**. The low recall of **RelRank** can be explained by its limited number of German affixes ("straße" and "weg"), the low amount of features and the manually created relevance criterion.

	Precision	Recall	F1 Score
TagMe	0.49	0.45	0.47
PopRank	0.69	0.66	0.67
RelRank (all entities)	0.08	0.08	0.08
RelRank (person entities)	0.35	0.11	0.17
<i>StreetToPerson</i>	0.95	0.91	0.93

Table 1: Evaluation of the classification for *StreetToPerson* and the selected baselines.

4.4 Application on OpenStreetMap

In a second step of the evaluation, we demonstrate how *StreetToPerson* can be used for identifying street-to-person relations between OpenStreetMap and Wikidata that are not yet contained in these sources. To this end, we apply *StreetToPerson* to the whole of Germany and two selected German states – the most populated state, North Rhine-Westphalia (NRW), and the least populated state, Bremen. From OSM, we select all streets contained in these regions and then apply *StreetToPerson* as depicted in Figure 1. Table 2 shows the results: For more than half of the 219,768 streets in NRW (Bremen: 2,504 of 6,733 streets), the person index returned at least one candidate person. After applying the street-to-person classifier, 28,857 street-to-person relations were identified (Bremen: 896). In the whole of Germany, more than 180,000 street-to-person relations are identified. This table also emphasizes the difficulty of identifying the correct person candidate: For 669,304 streets in Germany, more than 16 million candidate persons are found. An example of a false classification is the street "Grevesmühlweg" in Bremen which is wrongly assigned to Maria Grevesmühl but should be assigned to her father Hermann Grevesmühl – both musicians who were born and died in Bremen, thus highly similar.

Number of	Bremen	NRW	Germany
Streets	6,733	219,768	1,321,464
with candidate persons	2,504	110,968	669,304
Candidate persons	47,659	2,675,761	16,165,454
Street-to-person relations	896	28,857	183,022

Table 2: Number of street-to-person relations identified for OSM streets in Germany and two of its states.

Finally, we measure the correctness of the identified street-to-person pairs based on a subset of the streets in Table 2, which are

assigned to a person in OSM through the name:etymology key. Table 3 shows the results on this subset of 45 (Bremen) and 1,352 streets (NRW). The results show a precision of 0.90 for NRW and 0.94 for Bremen which confirms the generalizability of our approach from Wikidata to OSM.

	Bremen	NRW
Precision	0.94	0.90
Recall	0.64	0.61
F1 Score	0.76	0.73

Table 3: Evaluation of *StreetToPerson* on the OSM ground truth in two German states.

5 CONCLUSION AND FUTURE WORK

In this paper, we have presented *StreetToPerson*, an approach for connecting streets in OpenStreetMap to those persons in Wikidata whom they were named after. Through a combination of knowledge graph features and spatial features, *StreetToPerson* precisely identifies new relations. In future work, we plan to enrich the feature space through the utilisation of graph-based embeddings and to extend our approach to further countries and other relations between geographic entities in OpenStreetMap and knowledge graph entities.

ACKNOWLEDGMENTS

This work was partially funded by the Federal Ministry of Education and Research (BMBF), Germany under "Simple-ML" (01IS18054) and the DFG, German Research Foundation, under "WorldKG" (424985896).

REFERENCES

- [1] Hannah Bast, Patrick Brosi, Johannes Kalmbach, and Axel Lehmann. 2021. An Efficient RDF Converter and SPARQL Endpoint for the Complete OpenStreetMap Data. In *The 29th International Conference on Advances in Geographic Information Systems*. 536–539.
- [2] Oleksandra Bruns, Tabea Tietz, Mehdi Ben Chaabane, Manuel Portz, Felix Xiong, and Harald Sack. 2021. The Nuremberg Address Knowledge Graph. In *Extended Semantic Web Conference (ESWC) (Lecture Notes in Computer Science, Vol. 12739)*. Springer, 115–119.
- [3] Paulo Dias Almeida, Jorge Rocha, Andrea Ballatore, and Alexander Zipf. 2016. Where the Streets Have Known Names. In *International Conference on Computational Science and Its Applications (ICCSA '16)*, Vol. 9789. 1–12. https://doi.org/10.1007/978-3-319-42089-9_1
- [4] Alishiba Dsouza, Nicolas Tempelmeier, and Elena Demidova. 2021. Towards Neural Schema Alignment for OpenStreetMap and Knowledge Graphs. In *International Semantic Web Conference (ISWC)*. Springer, 56–73.
- [5] Alishiba Dsouza, Nicolas Tempelmeier, Ran Yu, Simon Gottschalk, and Elena Demidova. 2021. WorldKG: A World-Scale Geographic Knowledge Graph. In *Conference on Information and Knowledge Management (CIKM)*. ACM, 4475–4484.
- [6] Paolo Ferragina and Ugo Scaiella. 2010. TAGME: On-the-fly Annotation of Short Text Fragments (by Wikipedia Entities). In *Conference on Information and Knowledge Management (CIKM)*. 1625–1628.
- [7] Johanna Geiß and Michael Gertz. 2016. With a Little Help from my Neighbors: Person Name Linking Using the Wikipedia Social Network. In *International Conference on World Wide Web (WWW)*. ACM, 985–990.
- [8] Cedric Möller, Jens Lehmann, and Ricardo Usbeck. 2022. Survey on English Entity Linking on Wikidata. *Semantic Web (2022)*.
- [9] Nouria Ouali, Pierre Lannoy, Virginie Desamoury, Sandrine Guillaume, Fanny Mayné, Sophie Merville, Charlotte Odier, and Adèle Thébaux. 2021. Women in Brussels Street Names. Topography of a Minoritisation. *Brussels Studies (2021)*.
- [10] Nicolas Tempelmeier and Elena Demidova. 2021. Linking OpenStreetMap with Knowledge Graphs - Link Discovery for Schema-Agnostic Volunteered Geographic Information. *Future Generation Computer Systems* 116 (2021), 349–364.