

Towards Neural Schema Alignment for OpenStreetMap and Knowledge Graphs

Alishiba Dsouza¹ , Nicolas Tempelmeier² , Elena Demidova¹ 

¹Data Science & Intelligent Systems (DSIS), University of Bonn, Bonn, Germany

²L3S Research Center, Leibniz Universität Hannover, Hannover, Germany

dsouza@cs.uni-bonn.de tempelmeier@L3S.de

elena.demidova@cs.uni-bonn.de

Abstract. OpenStreetMap (OSM) is one of the richest openly available sources of volunteered geographic information. Although OSM includes various geographical entities, their descriptions are highly heterogeneous, incomplete, and do not follow any well-defined ontology. Knowledge graphs can potentially provide valuable semantic information to enrich OSM entities. However, interlinking OSM entities with knowledge graphs is inherently difficult due to the large, heterogeneous, ambiguous and flat OSM schema and the annotation sparsity. This paper tackles the alignment of OSM tags with the corresponding knowledge graph classes holistically by jointly considering the schema and instance layers. We propose a novel neural architecture that capitalizes upon a shared latent space for tag-to-class alignment created using linked entities in OSM and knowledge graphs. Our experiments performed to align OSM datasets for several countries with two of the most prominent openly available knowledge graphs, namely, Wikidata and DBpedia, demonstrate that the proposed approach outperforms the state-of-the-art schema alignment baselines by up to 53 percentage points in terms of F1-score. The resulting alignment facilitates new semantic annotations for over 10 million OSM entities worldwide, which is more than a 400% increase compared to the existing semantic annotations in OSM.

Keywords: OpenStreetMap, Knowledge Graph, Neural Schema Alignment

1 Introduction

OpenStreetMap (OSM) has evolved as a critical source of openly available geographic information globally, including rich data from 188 countries. This information is contributed by a large community, which currently counts more than 1.5 million volunteers. OSM captures a vast, and continuously growing number of geographic entities, currently counting more than 6.8 billion [14]. The descriptions of OSM entities consist of heterogeneous key-value pairs, so-called *tags*, and currently include over 80 thousand distinct keys. As OSM keys and tags do not

possess any machine-readable semantics, OSM data is not directly accessible for semantic applications.

Whereas knowledge graphs (KGs) can provide precise semantics for geographic entities, large publicly available general-purpose knowledge graphs like Wikidata [28], DBpedia [2], YAGO [21], and even specialized KGs like EventKG [9] and LinkedGeoData [24] lack coverage of geographic entities. For instance, in April 2021, 126,894 entities with the tag `highway=road` were present in OSM, whereas Wikidata included only 66,114 entities for the equivalent class “road”. We believe that an alignment of OSM and knowledge graphs at the schema level can make a wide variety of geographic entities in OSM accessible through semantic technologies and applications.

The problem of schema alignment between OSM and knowledge graph is particularly challenging due to several factors, most prominently including the heterogeneous representations of types and properties of geographic entities via OSM tags, unclear tag semantics, the large scale and flatness of the OSM schema, and the sparseness of the existing links. OSM does not limit the usage of keys and tags by any strict schema and provides only a set of guidelines¹. As a result, the types and properties of OSM entities are represented via a variety of tags that do not possess precise semantics. As an example, consider an excerpt from the representations of the entity “Zugspitze” (the highest mountain in Germany) in Wikidata and OSM:

Wikidata			OpenStreetMap	
Subject	Predicate	Object	Key	Value
Q3375	<i>label</i>	<i>Zugspitze</i>	<i>id</i>	27384190
Q3375	<i>coordinate</i>	47°25'N,10°59'E	name	<i>Zugspitze</i>
Q3375	<i>parentpeak</i>	Q15127	natural	<i>peak</i>
Q3375	<i>instance of</i>	<i>mountain</i>	summit:cross	<i>yes</i>

In Wikidata, an entity type is typically represented using the `instance of` property. In this example, the statement “Q3375 `instance of` mountain” indicates the type “mountain” of the entity “Q3375”. In OpenStreetMap, the type “mountain” of the same entity is indicated by the tag `natural=peak`. As OSM lacks a counterpart of the `instance of` property, it is not clear which particular tag represents an entity type and which tags refer to other properties. Furthermore, multiple OSM tags can refer to the same semantic concept. Finally, whereas the OSM schema with over 80 thousand distinct keys is very extensive, the alignment between OSM and knowledge graphs at the schema level is almost nonexistent. For instance, as of April 2021, Wikidata contained 585 alignments between its properties and OSM keys, which corresponds to only 0.7% of the distinct OSM keys. Overall, the flatness, heterogeneity, ambiguity, and the large scale of OSM schema, along with a lack of links, make the alignment particularly challenging.

Existing approaches for schema alignment operate at the schema and instance level and consider the similarity of schema elements, structural similarity, and instance similarity. As OSM schema is flat, ontology alignment methods that utilize hierarchical structures, such as [16,12], are not applicable. A transformation of OSM data into a tabular or relational format leads to highly sparse

tables with a large number of properties. Therefore, approaches to syntactic or instance-based alignment for relational or tabular data, such as, e.g., [30,5], or syntactic matching of schema element names [26] cannot yield good results for matching OSM tags with KG classes.

This paper takes the first important step to align OSM and knowledge graphs at the schema level using a novel neural method. In particular, we tackle tag-to-class alignment, i.e., we aim to identify OSM tags that convey class information and map them to the corresponding classes in the Wikidata knowledge graph and the DBpedia ontology. We present the Neural Class Alignment (NCA) model - a novel instance-based neural approach that aligns OSM tags with the corresponding semantic classes in a knowledge graph. NCA builds upon a novel shared latent space that aligns OSM tags and KG concepts and facilitates a seamless translation between them. To the best of our knowledge, NCA is the first approach to align OSM and KGs at the schema level with a neural method.

Our contributions are as follows:

- We present NCA - a novel approach for class alignment for OSM and KGs.
- We propose a novel shared latent space that fuses feature spaces from knowledge graphs and OSM in a joint model enabling simultaneous training of the schema alignment model on heterogeneous semantic and geographic sources.
- We develop an effective algorithm to extract tag-to-class alignments from the resulting model.
- The results of our evaluation demonstrate that the proposed NCA approach is highly effective and outperforms the baselines by up to 53 percentage points in terms of F1-score.
- As a result of the proposed NCA alignment method, we provide semantic annotations with Wikidata and DBpedia classes for over 10 million OSM entities. This result corresponds to a more than 400% increase compared to currently existing annotations.
- We make our code and datasets publicly available and provide a manually annotated ground truth for the tag-to-class alignment of OSM tags with Wikidata and DBpedia classes².

2 Problem Statement

In this section, we formalize the problem definition. First, we formally define the concepts of an OSM corpus and a knowledge graph. An OSM corpus contains nodes representing geographic entities. Each node is annotated with an identifier, a location, and a set of key-value pairs known as tags.

Definition 1. *An OSM corpus $\mathcal{C} = (N, T)$ consists of a set of nodes N representing geographic entities, and a set of tags T . Each tag $t \in T$ is represented as a key-value pair with the key $k \in K$ and a value $v \in V: t = \langle k, v \rangle$. A node $n \in N$, $n = \langle i, l, T_n \rangle$ is represented as a tuple containing an identifier i , a geographic location l , and a set of tags $T_n \subset T$.*

A KG contains real-world entities, classes, properties, and relations.

Definition 2. A knowledge graph $\mathcal{KG} = (E, C, P, L, F)$ consists of a set of entities E , a set of classes $C \subset E$, a set of properties P , a set of literals L , and a set of triples $F \subseteq E \times P \times (E \cup L)$.

The entities in E represent real-world entities and semantic classes. The properties in P represent relations connecting two entities, or an entity and a literal value. An entity in a KG can belong to one or multiple classes. An entity is typically linked to its class using the `rdf:type`, or an equivalent property.

Definition 3. A class of the entity $e \in E$ in the knowledge graph $\mathcal{KG} = (E, C, P, L, F)$ is denoted as: $\text{class}(e) = \{c \in C \mid (e, \text{rdf:type}, c) \in F\}$.

An OSM node and a KG entity referring to the same real-world geographic entity and connected via an identity link are denoted linked entities.

Definition 4. A linked entity $(n, e) \in E_L$ is a pair of an OSM node $n = \langle i, l, T_n \rangle$, $n \in N$, and a knowledge graph entity $e \in E$ that correspond to the same real-world entity. In a knowledge graph, a linked entity is typically represented using a $(e, \text{owl:sameAs}, i)$ triple, where i is the node identifier. E_L denotes the set of all linked entities in a knowledge graph.

This paper tackles the alignment of tags that describe types of nodes in an OSM corpus to equivalent classes in a knowledge graph.

Definition 5. Tag-to-class alignment: Given a knowledge graph \mathcal{KG} and an OSM corpus \mathcal{C} , find a set of pairs $\text{tag_class} \subseteq (T \times C)$ of OSM tags T and the corresponding \mathcal{KG} classes such that for each pair $(t, c) \in \text{tag_class}$ OSM nodes with the tag t belong to the class c .

3 Neural Class Alignment Approach

An alignment of an OSM corpus with a knowledge graph can include several dimensions, such as entity linking, node classification (i.e., alignment of OSM nodes with the corresponding semantic classes in a knowledge graph), as well as alignment of schema elements such as keys/tags and the corresponding semantic classes. These dimensions can reinforce each other. For example, linking OSM nodes with knowledge graph entities and classification of OSM nodes into knowledge graph classes can lead to new schema-level alignments and vice versa. Our proposed NCA approach systematically exploits the existing identity links between OSM nodes and knowledge graph entities based on this intuition. NCA builds an auxiliary classification model and utilizes this model to align OSM tags with the corresponding classes in a knowledge graph ontology.

NCA is an unsupervised two-step approach for tag-to-class alignment. Fig. 1 presents an overview of the proposed NCA approach architecture. First, we build an auxiliary neural classification model and train this model using linked entities in OSM and a KG. As a result, the model learns a novel shared latent space that aligns the feature spaces of OSM and a knowledge graph and implicitly captures tag-to-class alignments. Second, we systematically probe the resulting model to identify the captured alignments.

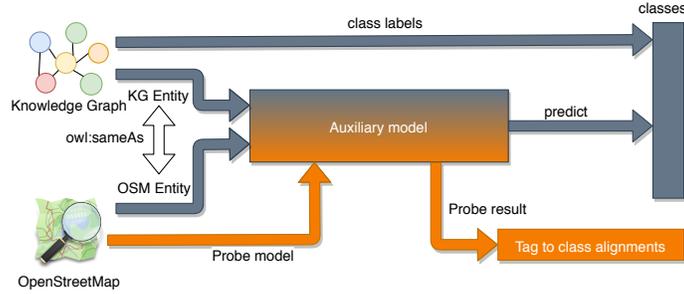


Fig. 1: Overview of the NCA architecture. Grey color indicates the first step (training of the auxiliary classification model). Orange color indicates the second step, i.e., the extraction of tag-to-class alignments.

3.1 Auxiliary Neural Classification Model

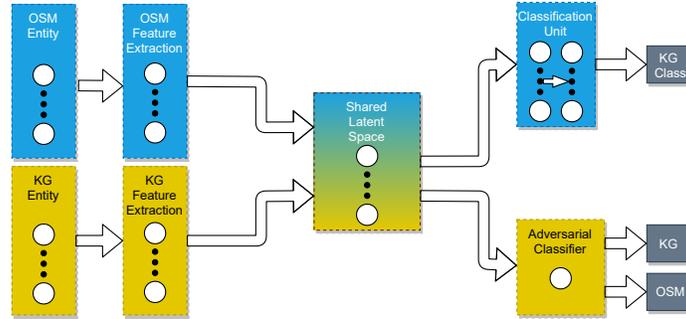


Fig. 2: Neural architecture of the auxiliary classification model. Blue color indicates the KG classification component, yellow marks the adversarial entity discrimination component.

In this step, we build a supervised auxiliary neural classification model for a dummy task of OSM node and KG entity classification. The model resulting from this step is later used for the tag-to-class alignment. Fig. 2 presents the model architecture. The auxiliary classification model architecture consists of several components described below.

OSM node representation. We represent an OSM node as a binary vector in an O -dimensional vector space. The space dimensions correspond to OSM tags or keys, and binary values represent whether the node includes the corresponding tag or key. The vector space dimensions serve as features for the classification model, such that we also refer to this space as the OSM feature space. To select the most descriptive tags to be included as dimensions in the OSM feature space,

we filter out low-quality tags using OSM taginfo³. We include only the tags with an available description in the OSM wiki⁴ having at least 50 occurrences within OSM. For tags with infrequent values (e.g., literals), we include only the keys as dimensions. As we aim to align geographic concepts and not specific entities, we do not include infrequent and node-specific values such as entity names or geographic coordinates in the representation. For instance, the concept of “mountain” is the same across different geographic regions, such that the geographic location of entities is not informative for the schema alignment.

KG entity representation. We represent a KG entity as a binary vector in a V -dimensional vector space. The space dimensions correspond to the KG properties, and binary values represent whether the entity includes the corresponding property. The vector space dimensions serve as features for the classification model, such that we also refer to this space as the KG feature space. To select the most descriptive properties to be included in the KG feature space, we rank the properties based on their selectivity concerning the class and the frequency of property usage (i.e., the number of statements in the KG that assign this property to an entity). Given a property p , we calculate its weight as: $weight(p, c) = n_{p,c} * \log \frac{N}{c_p}$.

Here, $n_{p,c}$ denotes the number of statements in which the property p is assigned to an entity of class c , N denotes the total number of classes in a knowledge graph, and c_p is the number of distinct classes that include the property p . For each class c , we select top-25 properties as features. These properties are included as dimensions in the KG feature space.

OSM & KG feature extraction. The KG and OSM feature representations serve as input to the specific fully connected feature extraction layers: OSM feature extraction and KG feature extraction. The purpose of these layers is to refine the vector representations obtained in the previous step.

Shared latent space & adversarial classifier. We introduce a novel *shared latent space* that fuses the initially disjoint feature spaces of OSM and KG such that entities from both data sources are represented in a joint space similarly. In addition to the training on OSM examples, shared latent space enables us to train our model on the KG examples. These examples provide the properties known to indicate class information [20]. The shared latent space component consists of a fully connected layer that receives the input from the OSM and KG feature extraction layers. Following recent domain adaption techniques [8], we use an adversarial classification layer to align latent representations of KG and OSM entities. The objective of the adversarial classifier is to discriminate whether the current training example is an OSM node or a KG entity, where the classification loss is measured as binary cross-entropy.

$$BinaryCrossEntropy = -\frac{1}{n} \sum_{i=1}^n [y_i \times \log(\hat{y}_i) + (1 - y_i) \times \log(1 - \hat{y}_i)],$$

where n is the total number of examples, y_i is the true class label, and \hat{y}_i is the predicted class label. Intuitively in a shared latent space, the classifier should not be able to distinguish whether a training example originates from OSM or

a KG. To fuse the initially disjoint feature spaces, we reverse the gradients from the adversarial classification loss: $\mathcal{L}_{adverse} = -BinaryCrossEntropy_{adverse}$.

Classification Unit. To train the auxiliary classification model for the OSM nodes, we exploit linked entities. We label OSM nodes with semantic classes of equivalent KG entities. We use these class labels as supervision in the OSM node classification task. More formally, given a linked entity $(n, e) \in E_L$ the training objective of the model is to predict $class(e)$ from n . Analogously, the training objective for a KG entity e is to predict the class label $class(e)$ of this entity.

We utilize a 2-layer feed-forward network as a classification model. In the last prediction layer of this network, each neuron corresponds to a class. As an entity can be assigned to multiple classes, we use a sigmoid activation function and a binary cross-entropy loss to achieve multi-label classification: $\mathcal{L}_{classification} = BinaryCrossEntropy_{classification}$. Finally, the joint loss function \mathcal{L} of the network is given by $\mathcal{L} = \mathcal{L}_{classification} + \mathcal{L}_{adverse}$. In the training process, we alternate OSM and KG instances to avoid bias towards one data source.

3.2 Tag-to-Class Alignment

In this step, we systematically probe the trained auxiliary classification model to extract the tag-to-class alignment. The goal of this step is to obtain the corresponding KG class for a given OSM tag. Algorithm 1 details the extraction process. First, we load the pre-trained auxiliary model m (line 1) and initialize the result set (line 2). We then probe the model with a given list of OSM tags \mathcal{T} (line 3). For a single tag $t \in \mathcal{T}$, we feed t to the OSM input layer of the auxiliary model and compute the complete forward propagation of t within m (line 4). We then extract the activation of the neurons of the last layer of the classification model before the sigmoid nonlinearity (line 5). As the individual neurons in this layer directly correspond to KG classes we expect that the activation of the specific neurons quantifies the likeliness that the tag t corresponds to the respective class. For each activation of a specific neuron a that is above the alignment threshold th_a (line 6-7), we extract the corresponding class c and add this class to the set of alignments (line 8). We determine the threshold value experimentally, as described later in Section 5.3. As an OSM tag can have multiple corresponding classes, we opt for all matches above the threshold value. Finally, the resulting set *align* constitutes the inferred tag-to-class alignments.

3.3 Illustrative Example

We illustrate the proposed NCA approach at the example of the ‘‘Zugspitze’’ mountain introduced in Section 1. We create the representation of the Wikidata object ‘‘Q3375’’ in the KG feature space by creating a binary vector that has ones in the dimensions that correspond to the properties this entity contains such as, `label`, `coordinate`, `parentpeak` and zeros otherwise. Note that `instance of predicate` is not included in the feature space as it represents the class label. Similarly, we encode the OSM node with the id ‘‘27384190’’ in the OSM feature space by creating a vector that includes `name`, `natural=peak`, `summit:cross`

Algorithm 1 Extract Tag-to-Class Alignment

Input: m Trained auxiliary model
 \mathcal{T} List of OSM tags
 th_a Alignment threshold
Output: $align \subseteq (T \times C)$ Extracted alignment of tags and classes

```

1: load( $m$ )
2:  $align \leftarrow \emptyset$ 
3: for all  $t \in \mathcal{T}$  do
4:   forward_propagation( $t, m$ )
5:    $activations \leftarrow extract\_activations(m)$ 
6:   for all  $a \in activations$  do
7:     if  $a > th_a$  then
8:        $align \leftarrow align \cup \{(t, class(a))\}$ 
9:     end if
10:  end for
11: end for
12: return  $align$ 

```

as ones and has zeros in all other dimensions. As described above, we use frequent key-value pairs such as `natural=peak` as features, whereas for the infrequent key-value pairs, such as `name=Zugspitze`, we use only the key (i.e., `name`) as a feature. The KG and OSM features spaces are then aligned in the shared latent space. To form this space, we train the auxiliary classification model that learns to output the correct class labels, such as “mountain”. In the last prediction layer of this model, each neuron corresponds to a class. After the training is completed, we can prob the classification model with a single tag such as `natural=peak`. The activation of the neurons in the prediction layer correspond to the predicted tag-to-class mapping. We output all classes with the activation values above the threshold th_a (here: “mountain”).

4 Evaluation Setup

This section introduces the evaluation setup in terms of datasets, ground truth generation, baselines, and evaluation metrics. All experiments were conducted on an AMD Opteron 8439 SE processor @ 2.7 GHz and 252 GB of memory.

4.1 Datasets

We carry out our experiments on OSM, Wikidata [28], and DBpedia [2] datasets.

Knowledge graphs: A sufficient number of linked entities and distinct classes is essential to train the proposed neural model and to achieve a meaningful schema alignment. We systematically rank European countries according to the number of linked entities between OSM and knowledge graphs and choose the top-4 countries having at least ten distinct classes in the linked entity set.

Based on these criteria, we select the Wikidata datasets for France, Germany, Great Britain, and Russia and the DBpedia datasets for France, Germany, Great Britain, and Spain. Although over 100,000 entity links between Russian DBpedia and OSM exist, most of these entities belong to only two classes. Hence, we omit Russian DBpedia from our analysis. In our experiments, we consider Wikidata and DBpedia snapshots from March 2021. We only consider geographic entities, i.e., the entities with valid geographic coordinates.

OpenStreetMap: We extract OSM data for France, Germany, Great Britain, Spain, and Russia. To facilitate evaluation, we only consider OSM nodes which include links to knowledge graphs. The number of entities assigned to specific knowledge graph classes follows a power-law distribution. We select the classes with more than 100 entities (i.e., 3% of classes in Wikidata) to facilitate model training. Note that some KG entities are linked to more than one OSM node, such that the number of OSM nodes and KG entities in the dataset differ.

4.2 Ground Truth Creation

For Wikidata, we start the creation of our ground truth based on the ‘‘OpenStreetMap tag or key’’ Wikidata property⁵. This property provides a link between a Wikidata class and the corresponding OSM tag. However, this dataset is incomplete and lacks some language-specific classes as well as superclass and subclass relationships based on our manual analysis. We manually extended the ground truth by checking all possible matches obtained by the proposed NCA approach and all baseline models used in the evaluation and added all correct matches to our ground truth. For DBpedia, we constructed the ground truth manually by labeling all combinations ($T \times C$) of OSM tags t and \mathcal{KG} classes C in our dataset. For both KGs, we take region-specific matches (‘‘Ortsteil’’ vs. ‘‘District’’) as well as subclass and superclass relations (e.g., ‘‘human settlement’’ vs. ‘‘city/village’’) into consideration.

4.3 Baselines

As OSM has a flat schema, ontology alignment methods that utilize structural features are not applicable. Therefore, we utilize schema alignment methods based on the schema element names and methods developed for tabular data as baselines. To fit our data to the baselines, we convert our OSM (source) data and KG (target) data into a tabular format. For OSM, we use the tags and keys as columns and convert each node into a row. Similarly, for KGs, the properties and classes are converted into columns, and the entities form the rows. Overall, we evaluate our proposed method against the following baselines:

Cupid: Cupid [12] matches schema elements based on element names, structure, and data types. Cupid is a 2-phase approach. The first phase calculates the lexicographic similarity of names and data types. The second phase matches elements using the structural similarity based on the element proximity in the ontology and the hierarchical position. As the OSM schema is flat, we consider a

flat hierarchy, where the OSM table is the root and all columns are child nodes. The final Cupid score is the average similarity between the two phases.

Levenshtein Distance (LD): The Levenshtein distance (edit distance) is a string-based similarity measure used to match ontology elements lexicographically. The Levenshtein distance between two element names is calculated as the minimal number of edits needed to transform one element name to obtain the other. The modifications include addition, deletion, or replacement of characters [26]. We calculate the Levenshtein distance between all pairs of class names and tags and accept all pairs with a distance lower than a matching threshold $th_l \in [0, 1]$. We apply an exhaustive grid search to optimize the value of th_l for each dataset and report the highest resulting F1-scores.

EmbDi: EmbDi [4] is an algorithm for schema alignment and entity resolution. The algorithm maps table rows to a directed graph based on rows, columns, and cell values. EmbDi infers column embeddings by performing random walks on the graph. The random walks form sentences that constitute an input to a Word2Vec model. Finally, the similarity of the two columns is measured as the cosine similarity of the respective embeddings.

Similarity Flooding (SF): Similarity Flooding [13] transforms a data table into a directed labeled graph in which the nodes represent table columns. The weights of graph edges represent the node similarity, initialized using string similarity of the column names. The algorithm refines the weights by iteratively propagating similarity values along the edges. Each pair of nodes connected with a similarity value above the matching threshold forms an alignment. The algorithm filters out specific matches such as data types and schema elements. We experimentally select a threshold value of 0 to increase recall, as at higher threshold values, the algorithm obtains no matches.

4.4 Metrics

The standard evaluation metrics for schema alignment are precision, recall, and F1-score computed against a reference alignment (i.e., ground truth). We evaluate the kg_{class} mappings as pairs, where each pair consists of one tag and one class (tag-to-class alignment). **Precision** is the fraction of correctly identified pairs among all identified pairs. **Recall** is the fraction of correctly identified pairs among all pairs in the reference alignment. **F1-score** is the harmonic mean of recall and precision. We consider the F1-score to be the most relevant metric since it reflects both recall and precision.

5 Evaluation

The evaluation aims to assess the performance of the proposed NCA approach for tag-to-class alignment in terms of precision, recall, and F1-score. Furthermore, we aim to analyze the influence of the confidence threshold and the impact of the shared latent space on the alignment performance. Note that we do not evaluate the artificial auxiliary classification task. Instead, we evaluate the utility of the auxiliary model in the overall schema alignment task.

Table 1: Tag-to-class alignment performance for OSM tags to Wikidata classes.

Approach	France			Germany			Great Britain			Russia			Average		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
CUPID	0.06	1	0.12	0.03	0.70	0.06	0.07	1	0.14	0.08	0.80	0.15	0.08	0.88	0.12
LD	0.45	0.28	0.35	0.65	0.34	0.44	0.54	0.37	0.44	0.64	0.34	0.45	0.58	0.33	0.42
EMBDi	0.03	1	0.06	0.02	1	0.03	0.04	1	0.06	0.02	1	0.03	0.03	1	0.05
SF	0.03	1	0.06	0.02	1	0.03	0.01	1	0.03	0.02	1	0.03	0.02	1	0.04
NCA	0.63	0.66	0.65	0.59	0.65	0.61	0.71	0.56	0.63	0.64	0.51	0.58	0.64	0.60	0.62

Table 2: Tag-to-class alignment performance for OSM tags to DBpedia classes.

Approach	France			Germany			Great Britain			Spain			Average		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
CUPID	0.32	1	0.48	0.18	1	0.31	0.41	1	0.58	0.44	1	0.63	0.33	1	0.50
LD	0.31	0.57	0.41	0.32	0.37	0.34	0.73	0.46	0.57	0.34	0.94	0.50	0.42	0.58	0.45
EMBDi	0.16	1	0.28	0.09	1	0.17	0.29	1	0.45	0.24	1	0.38	0.20	1	0.32
SF	0.14	1	0.27	0.10	1	0.18	0.27	1	0.42	0.24	1	0.39	0.19	1	0.31
NCA	0.95	0.90	0.92	0.96	0.79	0.87	0.81	0.84	0.83	1	0.84	0.91	0.93	0.84	0.88

5.1 Tag-to-Class Alignment Performance

Table 1 and Table 2 summarize the performance results of the baselines as well as our proposed NCA approach with respect to precision, recall and F1-score for tag-to-class alignment of OSM tags to Wikidata and DBpedia classes, respectively.

The proposed NCA approach outperforms the baselines in terms of precision and F1-score on all datasets. On Wikidata, we achieve up to 30 percent points F1-score improvement and 19.75 percent points on average compared to the best baseline. On DBpedia, we achieve up to 53 percent points F1-score improvement and 37.5 percent points on average. As OSM lacks a hierarchical structure, limiting structural comparison, most of the applicable baselines build on the name comparison. Here, the heterogeneity of OSM tags limits the precision of the baselines substantially. Compared to the other baselines, Levenshtein distance and Cupid obtain the highest F1-scores for Wikidata and DBpedia, respectively, even though the absolute values achieved by these baselines are relatively low. SF and EmbDI obtain only low similarity values, resulting in low precision. An increase of the confidence threshold for these baselines leads to zero matches.

We observe performance variations across countries and knowledge graphs, with French Wikidata and DBpedia achieving the highest F1-scores in comparison to other countries. These variations can be explained by the differences in the dataset characteristics, such as the number of links, the number of entities per class, and the number of unique tags and classes per country. These characteristics vary significantly across the datasets. Furthermore, the number of classes per entity varies. On average, Wikidata indicates one class per entity (i.e., the most specific class), whereas DBpedia indicates three classes per entity on average (i.e., the specialized as well as more generic classes at the higher levels of the DBpedia ontology). This property makes the model trained on the DBpedia knowledge graph more confident regarding the generic classes compared

Table 3: Tag-to-class alignments obtained using NCA approach.

	France	Germany	Great Britain	Russia
Wikidata	amenity=bicycle_rental: bicycle-sharing station	amenity=cinema: movie theater	railway=station: railway station	station=subway: metro station
	France	Germany	Great Britain	Spain
DBpedia	railway=station: Place	place=municipality: Place	place=hamlet: Place	railway=station: ArchitecturalStructure

Table 4: Tag-to-class alignment performance for Wikidata and DBpedia.

Approach	Avg. Wikidata			Avg. DBpedia		
	Precision	Recall	F1	Precision	Recall	F1
NCA w/o shared latent space	0.60	0.33	0.40	0.69	0.87	0.76
NCA	0.64	0.60	0.62	0.93	0.84	0.88

to the specialized classes, such that generic classes obtain higher F1-scores. Our observations indicate that it is desirable to obtain more training examples that align entities with more specific classes, such as in the Wikidata dataset.

Table 3 illustrates the most confident tag-to-class alignments in terms of the obtained model activations using the NCA approach. As discussed above, Wikidata alignments with high confidence scores are more specific than those obtained with the DBpedia datasets.

5.2 Influence of the Shared Latent Space

Table 4 summarizes the performance of the proposed NCA approach and NCA without the shared latent space for tag-to-class alignment of OSM with Wikidata and DBpedia, respectively.

We observe that the shared latent space helps to achieve an increase in F1-score of 22 percentage points and 12 percentage points for Wikidata and DBpedia, respectively.

Compared to the Wikidata datasets, we observe smaller improvements on DBpedia datasets. DBpedia has an imbalance between the tags and classes that results in many-to-one alignments between tags and classes, where one class corresponds to several different tags. For example, in all DBpedia datasets, the *place* and *populatedPlace* are frequently occurring classes for various tags such as *tourism=museum*, *place=village*, *place=town*. In such a case, DBpedia properties add less focused information to the matching process. Furthermore, we observe high F1-score of the proposed NCA model without the shared latent space on DBpedia dataset. Intuitively, further improving these high scores is more difficult than improving the comparable low scores observed on Wikidata (e.g., 0.4 F1-score on Wikidata). In summary, the shared latent space improves the matching performance. We observe the highest improvements on Wikidata.

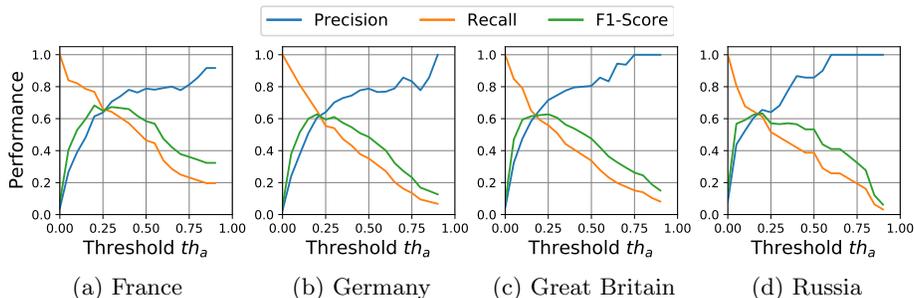


Fig. 3: Precision, recall, and F1-score vs. the confidence threshold for Wikidata.

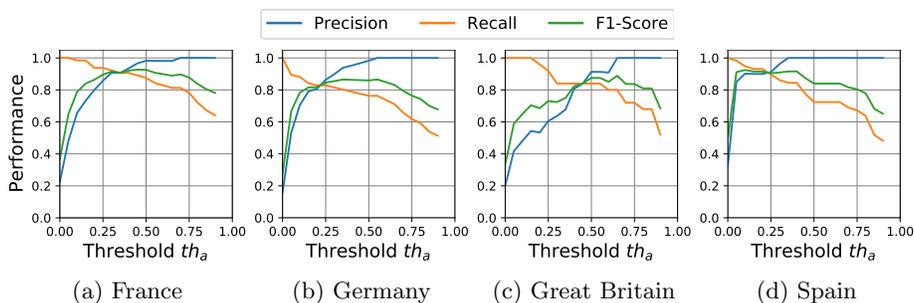


Fig. 4: Precision, recall, and F1-score vs. the confidence threshold for DBpedia.

5.3 Confidence Threshold Tuning

We evaluate the influence of the confidence threshold value th_a on the precision, recall, and F1-score. The threshold th_a indicates the minimum similarity at which we align a tag to a class. Fig. 3 and 4 present the alignment performance with respect to th_a for Wikidata and DBpedia. As expected, we observe a general trade-off between precision and recall, whereas higher values of th_a result in higher precision and lower recall. In our experiments, we select the confidence threshold of $th_a = 0.25$ and $th_a = 0.4$ for Wikidata and DBpedia, respectively, as these values allow us to balance precision and recall. The threshold can be further tuned for specific geographic regions.

5.4 Alignment Impact

To assess the impact of NCA, we compare the number of OSM entities that can be annotated with semantic classes using the alignment discovery by NCA with the number of entities that are linked to a KG in the currently existing datasets. For Wikidata, we observe 2,004,510 linked OSM entities and 10,163,762 entities that can be annotated with semantic classes using NCA. This result corresponds to an increase of 407.04% of entities with semantic class information. For DBpedia, we observe 1,396,378 linked OSM entities and 8,301,450 entities that

can be annotated with semantic classes. This result corresponds to an increase of 494.5% of entities with semantic class information. We make the resulting dataset available to facilitate reproducibility and further research.

6 Related Work

This work is related to ontology alignment, alignment of tabular data, feature space alignment, and link discovery.

Ontology Alignment. Ontology alignment (also ontology matching) aims to establish correspondences between the elements of different ontologies. The efforts to interlink open semantic datasets and benchmark ontology alignment approaches have been driven by the W3C SWEO Linking Open Data community project⁶ and the Ontology Alignment Evaluation Initiative (OAEI)⁷ [1]. Ontology alignment is conducted at the element-level, and structure-level [19]. The element-level alignment typically uses natural language descriptions of the ontology elements, such as labels and definitions. Element-level alignment adopts string similarity metrics such as, e.g., edit distance. Structure-level alignment exploits the similarity of the neighboring ontology elements, including the taxonomy structure, as well as shared instances [16]. Element-level and structure-level alignment have also been adopted to align ontologies with relational data [5] and tabular data [30]. Jiménez-Ruiz et al. [10] divided the alignment task into independent, smaller sub-tasks, aiming to scale up to very large ontologies. Machine learning has been widely adopted for ontology alignment. In the GLUE architecture [6], semantic mappings are learned in a semi-automatic way while [18] proposed a matching system that integrates string-based and semantic similarity features. More recently, more complex approaches using deep neural networks have been used for ontology alignment and schema matching. The proposed architectures include convolutional neural networks [3] and stacked autoencoders [29]. The lack of a well-defined ontology of OSM hinders the application of ontology alignment approaches to OSM data. In contrast, our instance-based NCA approach presented in this paper enables an effective alignment of tags to classes.

Tabular Data Alignment. Another branch of research investigated the schema alignment of tabular data [22]. Cappuzzo et al. [4] proposed EmbDi approach, which generates the graph structure for the tabular data, builds sentences from the graph, and generates embeddings to find similarity between the schema elements. Cupid [12] matches schema elements based on element names, structure, and data types. Similarity Flooding [13] transforms a table into a directed labeled graph in which nodes represent columns to compute similarity values iteratively. We employ the EmbDi, Cupid, and Similarity Flooding algorithms as baselines for our evaluation.

Feature Space Alignment. Recently, various studies investigated the alignment of feature spaces extracted from different data sources. Application domains include computer vision [7] and machine translation [11]. Ganin et al. [8] proposed a neural domain adaptation algorithm that considers labeled data from a source domain and unlabeled data from a target domain. While this approach

was originally used to align similar but different distributions of feature spaces, we adopt the gradient reversal layer proposed in [8] to fuse information from the disjoint features spaces of OSM and KGs, not attempted previously.

Link Discovery. Link Discovery is the task of identifying semantically equivalent resources in different data sources [15]. Nentwig et al. [15] provide a recent survey of link discovery frameworks with prominent examples, including Silk [27] and LIMES [17]. In particular, the Wombat algorithm, integrated within the LIMES framework [23], is a state-of-the-art approach for link discovery in knowledge graphs. Specialized approaches [25] focus on the discovery links between OSM and knowledge graphs. We build on existing links between OSM and knowledge graphs to align knowledge graph classes to OSM tags in this work.

7 Conclusion

In this paper, we presented NCA - the first neural approach for tag-to-class alignment between OpenStreetMap and knowledge graphs. We proposed a novel shared latent space that seamlessly fuses features from knowledge graphs and OSM in a joint model and makes them simultaneously accessible for the schema alignment. Our model builds this space as the core part of a neural architecture incorporating an auxiliary classification model and an adversarial component. Furthermore, we proposed an effective algorithm that extracts tag-to-class alignments from the resulting shared latent space with high precision. Our evaluation results demonstrate that NCA is highly effective and outperforms the baselines by up to 53 percentage points in F1-score. We make our code and manually annotated ground truth data publicly available to facilitate further research.

Acknowledgements This work was partially funded by DFG, German Research Foundation (“WorldKG”, 424985896, DE 2299/2-1), BMBF, Germany (“Simple-ML”, 01IS18054) and BMWi, Germany (“d-E-mand”, 01ME19009B).

References

1. Algergawy, A., et al.: Results of the ontology alignment evaluation initiative 2019. In: Proc. of the OM-2019. CEUR Workshop Proceedings, vol. 2536 (2019)
2. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: DBpedia: A Nucleus for a Web of Open Data. In: ISWC 2007. Springer (2007)
3. Bento, A., Zouaq, A., Gagnon, M.: Ontology matching using convolutional neural networks. In: LREC (2020)
4. Cappuzzo, R., Papotti, P., Thirumuruganathan, S.: Creating embeddings of heterogeneous relational datasets for data integration tasks. In: SIGMOD’20 (2020)
5. Demidova, E., Oelze, I., Nejd, W.: Aligning freebase with the YAGO ontology. In: Proceedings of CIKM 2013. pp. 579–588. ACM (2013)
6. Doan, A., Madhavan, J., Domingos, P., Halevy, A.: Ontology matching: A machine learning approach. Handbook on Ontologies in Information Systems (07 2003)
7. Fernando, B., Habrard, A., Sebban, M., Tuytelaars, T.: Unsupervised visual domain adaptation using subspace alignment. In: 2013 IEEE ICCV (2013)

8. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* (2016)
9. Gottschalk, S., Demidova, E.: EventKG - the hub of event knowledge on the web - and biographical timeline generation. *Semantic Web* **10**(6), 1039–1070 (2019)
10. Jiménez-Ruiz, E., Agibetov, A., Chen, J., Samwald, M., Cross, V.: Dividing the ontology alignment task with semantic embeddings and logic-based modules. In: *ECAI 2020* (2020)
11. Lample, G., Conneau, A., Ranzato, M., Denoyer, L., Jégou, H.: Word translation without parallel data. In: *ICLR 2018* (2018)
12. Madhavan, J., Bernstein, P., Rahm, E.: Generic schema matching with cupid. Tech. rep. (August 2001)
13. Melnik, S., Garcia-Molina, H., Rahm, E.: Similarity flooding: a versatile graph matching algorithm and its application to schema matching. In: *ICDE 2002* (2002)
14. Neis, P.: OSMstats. <https://osmstats.neis-one.org/>, [Online; accessed 10-April-2021]
15. Nentwig, M., Hartung, M., Ngomo, A.N., Rahm, E.: A survey of current link discovery frameworks. *Semantic Web* **8**(3), 419–436 (2017)
16. Ngo, D., Bellahsene, Z., Todorov, K.: Opening the black box of ontology matching. In: *Proc. of the ESWC 2013* (2013)
17. Ngomo, A.N., Auer, S.: LIMES - A time-efficient approach for large-scale link discovery on the web of data. In: *IJCAI 2011*. pp. 2312–2317 (2011)
18. Nkisi-Orji, I., Wiratunga, N., Massie, S., Hui, K., Heaven, R.: Ontology alignment based on word embedding and random forest classification. In: *ECML PKDD 2018*
19. Otero-Cerdeira, L., Rodríguez-Martínez, F.J., Gómez-Rodríguez, A.: Ontology matching: A literature review. *Expert Syst. Appl.* **42**(2), 949–971 (2015)
20. Paulheim, H., Bizer, C.: Type inference on noisy RDF data. In: *ISWC 2013* (2013)
21. Pellissier Tanon, T., Weikum, G., Suchanek, F.: YAGO 4: A reason-able knowledge base. In: *ESWC 2020*. Springer (2020)
22. Rahm, E., Bernstein, P.A.: A survey of approaches to automatic schema matching. *VLDB J.* **10**(4), 334–350 (2001)
23. Sherif, M.A., Ngonga Ngomo, A., Lehmann, J.: Wombat - A generalization approach for automatic link discovery. *LNCS*, vol. 10249, pp. 103–119. Springer (2017)
24. Stadler, C., Lehmann, J., Höffner, K., Auer, S.: Linkedgedata: A core for a web of spatial open data. *Semant. Web* **3**(4), 333–354 (Oct 2012)
25. Tempelmeier, N., Demidova, E.: Linking openstreetmap with knowledge graphs - link discovery for schema-agnostic volunteered geographic information. *Future Gener. Comput. Syst.* **116**, 349–364 (2021)
26. Unal, O., Afsarmanesh, H.: Using linguistic techniques for schema matching. pp. 115–120 (01 2006)
27. Volz, J., Bizer, C., Gaedke, M., Kobilarov, G.: Silk - A link discovery framework for the web of data. In: *Proceedings of the WWW2009 Workshop on Linked Data on the Web, LDOW 2009*. vol. 538
28. Vrandečić, D., Krötzsch, M.: Wikidata: A free collaborative knowledgebase. *Communications of the ACM* **57**, 78–85 (09 2014)
29. Xiang, C., Jiang, T., Chang, B., Sui, Z.: ERSOM: A structural ontology matching approach using automatically learned entity representation. In: *EMNLP* (2015)
30. Zhang, S., Balog, K.: Web table extraction, retrieval, and augmentation: A survey. *ACM Trans. Intell. Syst. Technol.* **11**(2), 13:1–13:35 (2020)

Notes

¹OSM “How to map a”: https://wiki.openstreetmap.org/wiki/How_to_map_a

²GitHub repository: <https://github.com/alishiba14/NCA-OSM-to-KGs>

³OSM taginfo: <https://taginfo.openstreetmap.org/tags>

⁴OSM wiki: <https://wiki.openstreetmap.org/wiki/>

⁵Wikidata “OpenStreetMap tag or key” property: <https://www.wikidata.org/wiki/Property:P1282>

⁶<https://www.w3.org/wiki/SweoIG/TaskForces/CommunityProjects/LinkingOpenData>

⁷OAEI evaluation campaigns: <http://oaei.ontologymatching.org>