Culture-aware Point-of-Interest Category Completion in a Global Location-Based Social Network Database without Access to User Data

Nikolaos Lagos
Naver Labs Europe
Meylan, France
nikolaos.lagos@naverlabs.com

Ioan Calapodescu
Naver Labs Europe
Meylan, France
ioan.calapodescu@naverlabs.com

ABSTRACT

Point of Interest (POI) categories can facilitate a number of services, such as location-based search and place recommendation. However, such information can be incomplete and/or incorrect, especially in crowdsourcing environments. In the literature, automatic category imputation has been suggested to tackle this problem, showing that contextual information is vital for increasing the quality of such predictions. To this end, users' check-in data, and most particularly location and time of visit, is often used as the notion of context. In this work, we propose a method that considers culture as a contextual parameter. Contrary to existing methods, our approach does not require access to user data. We illustrate the feasibility of our method by performing experiments on data from Foursquare, a global location-based social network.

1 INTRODUCTION

Point of Interest (POI) categories can facilitate a number of services, such as location-based search and place recommendation. However, as discussed in recent work, categories, especially in crowdsourcing environments, can be incomplete and/or incorrect. Automatic category prediction has therefore been proposed to remedy this problem and impute missing categories [20].

Recent advances in POI categorisation have shown that contextual information is vital for increasing the quality of automatic POI category prediction [5, 11, 15, 16]. To this end, users' check-in information, and most particularly location and time of visit are often used to define the context. There are two important shortcomings though with these approaches:

(1) Getting relevant data presupposes that users' will allow their check-in data to be shared with the corresponding service. However, recent initiatives and laws, such as the EU General Data Protection Regulation (GDPR), stipulate that users should have more control over their personal data, and potentially disallow the use (and storage) of data, such as their check-in information, from third-parties.
(2) Context is defined in terms of location proximity in existing systems. However, recent advances in recommender systems and information search have shown that integrating information about cultural backgrounds, especially in a global setting, can help in developing high quality systems [7, 17].

The main contributions of this work are as follows.

- To our knowledge this is the first study of a multi-lingual, global location-based social network database related to culture-aware POI category imputation. We formally define the problem and present a corresponding analysis.
- Culture-related categorisation without requiring access to user data, has not been proposed before. To achieve that, we simulate user information, by replacing culture-related training inputs in an appropriate manner, at inference time. For instance, the country where a POI is located can be one of the training inputs. We can replace at inference time the values of this input with the nationality of the user.

The rest of this paper is organised as follows. We review related work in Section 2. We formally define the problem in Section 3 and describe our method in Section 4. Experiments are presented in Section 5. Section 6 includes the conclusions of this work.

1.1 Industrial context

Our company Naver, provides, among other things, location-based services. Good quality POI data is thus of major importance. In this context, in Naver Labs Europe, we have been exploring automatic multi-lingual methods for completing and correcting POI semantic tags found in Foursquare's database, a global crowd-sourced location-based social network[1]. The scope of our work is to eventually support:

- A user that is not familiar with local culture to discover appropriate POIs in the vicinity of her/his position. If POIs are not categorised appropriately, the user can not easily search for them and they will not be included in the search results. In addition, proper POI categorisation could also help in recommendation.

2 RELATED WORK

2.1 Point-of-Interest Category Prediction

Most of the work on Point-of-Interest category prediction has taken place in the context of Location-Based Social Networks.

Ye et al. [16] were the first to show that taking into account the geographical, local context of users can improve POI recommendation and categorisation. From then on, other related work has been systematically using such contextual data, originating from check-ins and/or related mobile sensors [15].

Krumm [11] showed that other personal information (e.g. gender and age range) could also help to improve the results further.

The type of user data explored in the state-of-the-art include the number, frequency, time and duration of user check-ins, and sometimes demographic information (e.g. age range, gender). This data is usually combined with the location of the POI and its proximity to other POIs.

[1] We got access to this data thanks to an agreement between Naver Labs and Foursquare.

He et al. [5] and Zhou et al. [20], have in addition used POI information including user defined semantic tags, and name token embeddings pre-computed on a domain-specific corpus.

Jiang et al. [8] apply machine classification techniques to the problem of fusing different POI databases under a common classification hierarchy, the North American Industry Classification System (NAICS). Their study involves only a few American towns and they use as input features only the categories and their predefined relations, as already manually attributed in the original data sources to the POIs.

A number of works have been carried out on location prediction based in social streams, e.g. Twitter, where the main research interest is using noisy and short text for classification. For instance, Cano et al. [2] uses tweets to infer volatile POI classes according to specific temporary events happening at a specific location. Interested readers may refer to Zheng et al. [19] for a comprehensive survey of the domain. Despite superficial commonalities, this subject is different from the one studied here.

To summarise, as mentioned in the Introduction, all the aforementioned approaches assume access to user data at training time and have not dealt with the notion of user's culture.

2.2 Culture-aware Recommendation

Recent work had highlighted the importance of modelling users' culture in recommendation and information search [7]. Notably, the cultural background of a user was found to play a vital role in how recommended items are judged [14].

To our knowledge, in the area of automatic recommender systems, Zangerle et al. [17] uses culture as a computation parameter. The authors, base their proxy for defining culture on the nationality of the users and use Hofstede et al's. [6] grouping of nationalities in culturally similar clusters as a guide for their computation model[2]. However, in this case as well, the authors assume that they have access to user data.

3 PROBLEM DEFINITION

Our goal is to predict POIs' categories that are appropriate to a specific culture. For instance, a typical place found in the database of Foursquare is "La Table du Ramen"[3], which is located in Paris, France. The category found in Foursquare's database is *Japanese Restaurant*, which may be sufficient for the local culture. However, a Japanese would expect the POI to be categorised at a much more fine-grained level, as e.g. *Ramen Restaurant* or as a *Noodle House*. The objective is to automatically predict such categories, according to one's culture.

We consider that POI category prediction in our context is equivalent to the problem of completing a specific attribute of the dataset, the one that represents categories of POIs, based on data from the remaining attributes. Formally, a POI p should have an attribute that includes an ideal, complete and correct, category labelset i.e. set of relevant labels $L \subset \Lambda$, where $\Lambda = l_1, \ldots, l_m$ is the set of all possible labels, while the rest of the attributes are represented by the set A. However, in practice, the set of labels attributed to p, what we call thereof the observed labelset $L_o \subset \Lambda$, can be incomplete and/or incorrect.

[2]Hofstede et al. [6] notes that culture is always a collective phenomenon, as it is, at least partly, shared with people who live or lived within the same social environment, which is where it was learned. In that sense, context may encompass a lot of different aspects, including the notions of social status, education, and language. Hofstede et al. [6] mentions that "one's country" is an important parameter that defines culture in this sense [6]

[3]English translation:The Table of Ramen.

In this work, we assume that there are several "culture specific" labelsets, such that $\bigcup_{j=1}^{n} L_j \subseteq L$, where n is the maximum number of cultures that are represented by the labels in Λ and $n < \infty$. For instance revisiting our example, if we assume that $L = \{Japanese\ Restaurant, Noodle\ House, Ramen\ Restaurant\}$, and we have at least two cultures French and Japanese with corresponding labelsets L_1 and L_2, then it could be that $L_1 = \{Japanese\ Restaurant, Noodle\ House\}$ and $L_2 = \{Noodle\ House, Ramen\ Restaurant\}$, while $L_o = \{Japanese\ Restaurant\}$.

We denote by $\boldsymbol{y} = (y^1, \ldots, y^m)$ an m-dimensional binary vector where $y^i \in [0, 1]$ such that $y^i = 1$ if and only if $l_i \in L$. We define a variant of it for culture c as $\boldsymbol{y_c} = (y_c^1, \ldots, y_c^m)$ where $y_c^i \in [0, 1]$ such that $y_c^i = 1$ if and only if $l_i \in L_c$ and each $y^i - y_c^i$ is non-negative. Accordingly, the m-dimensional binary vector $\boldsymbol{y_o} = (y_o^1, \ldots, y_o^m)$ with $y^i \in [0, 1]$ has $y_o^i = 1$ if and only if $l_i \in L_o$.

We assume that $C \cup N = A$ where C stands for the set of culture-related attributes, and N the set of culture-independent ones. Considering again our example, the C could be instantiated by the country where the restaurant is located, and N by the name "La Table du Ramen".

We denote by $\boldsymbol{x}, \boldsymbol{x_C}, \boldsymbol{x_N}$ the vectors that represent correspondingly A, C, and N, such that the observed p in the dataset, is defined as

$$p = \{\boldsymbol{x}, \boldsymbol{y_o}\} = \{\boldsymbol{x_C}, \boldsymbol{x_N}, \boldsymbol{y_o}\}$$

Our objective is to make culture-specific predictions such that for culture c

$$p = \{\boldsymbol{x}, \boldsymbol{y_c}\}$$

We thus formulate our goal as a multi-label classification problem where we want to find a classifier $b_c : X \to Y$ where X is the input space (all possible attribute vectors) and Y the output space (all possible labelset vectors), such that $\boldsymbol{y_c} = b_c(\boldsymbol{x})$.

4 CATEGORY PREDICTION METHOD

The category of POIs, especially in a location-based social network, is related to the cultural profile of the users that visit it. This has been proven via the inclusion of user profiles in related work (c.f. Section 2). However, instead of accessing user information to discover such profiles, we use the observation that the majority of POIs are categorised in a manner that reflects local culture in location-based social network databases. For instance, as shown in Fig. 1, we find in Foursquare that restaurants selling noodle dishes[4] are usually categorised as *Asian Restaurant* in France, while *Ramen Restaurant* is by a large margin the most popular category in Japan.

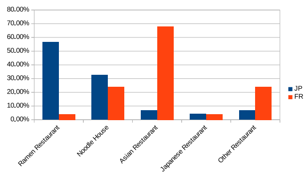

Figure 1: Category distribution of POIs having the token "noodle" in their name in Japan and France. It is obvious that "Ramen Restaurant" is the most popular category in Japan and "Asian Restaurant" in France.

Based on this insight, if at training time we use culture related attributes to learn a latent representation of POI's categories,

[4]The token "noodle" is explicitly mentioned in the POI name.

then at inference time we can replace the corresponding inputs according to the target cultural profile. For instance, if at training time we use the country in which the POI is located as an input parameter to our model, at inference time we can replace the value of this parameter with the target country, simulating what would happen if the same POI was located in the target country instead. We can thus generate culturally-appropriate predictions and complete the database offline. Revisiting the above example, we can assume that some of the Ramen restaurants located in Japan, would be categorised in a different manner, e.g. as *Noodle Houses*, if the value of the country was changed to France[5]. An overview of the proposed method is shown in Fig. 2.

Based on the discussion above, we reformulate our problem, and look for a classifier b_c such that $\boldsymbol{y_c} = b_c(\boldsymbol{x_c})$ where $\boldsymbol{x_c}$ is a culture-specific variant of the input.

To find b_c, we follow a standard approach and transform our problem into finding a real-valued vector function $f : X \to S \in [0, 1]^m$ that allows to indicate the relevance of a label l_i in relation to the input i.e. $f(\boldsymbol{x_c}) = (f(\boldsymbol{x_c}, \boldsymbol{l_1},), f(\boldsymbol{x_c}, \boldsymbol{l_2}), ..., f(\boldsymbol{x_c}, \boldsymbol{l_m}))$ where $f(\boldsymbol{x_c}, \boldsymbol{l_i})$ is the confidence of $l_i \in \Lambda$ being a correct label for $\boldsymbol{x_c}$ and m is the number of labels. Actually this corresponds to an estimation of $p(y_c^i|\boldsymbol{x_c}) : y_c^i \in [0, 1]$. Note that ideally, observed outputs should be completely specified vectors, however in our context the training instances are only partially complete, so of the form $(\boldsymbol{x_c^i}, y_o^i)$. We follow the Binary Relevance method, thus learn m binary models, each specialised into predicting whether one label is correct or not, independently from the other labels. For an unseen $\boldsymbol{x_c}$, the predicted labels are then the union of the predictions of all the binary models.

To learn the binary models we perform the following steps.

- Attribute selection: We use the name and spatial geo-coordinates of the POIs.
- Vectorisation: This step includes transforming the attributes in a form that can be treated by the classifier, as explained in Section 4.1.
- Training: A model that learns to predict the probability of l_i being a correct label given $\boldsymbol{x}$ is computed in this step. As explained in the previous paragraph, our problem is casted as a supervised machine learning problem. Details are provided in Section 4.2.
- Inference: Whether l_i is a correct label for culture-specific variants of $\boldsymbol{x}$ is computed in this step. Details are given in Section 4.3.

4.1 Vectorisation

Categorical variables. We represent them with one-hot encoded embeddings, as usually reported in the literature.

Sequential variables. Biessmann et al. [1] report that character-based representations are more robust for a similar setting to ours (i.e. sparse data and multiple languages). In addition, Joulin et al. [9] and Biessmann et al. [1] mention that character n-grams can perform better than simple, unigram, character-based LSTMs. After experimentation we have adopted trigram character based LSTMs for POI names.

Spatial variables. Geographical coordinates are the most important spatial attributes that characterise a POI. For instance, latitude and longitude are two of the most frequently used geographical coordinates. The predominant way of modelling coordinates is to discretise the input space [13, 18]. This could take the form of a grid separated into a fixed number of cells. Usually in this case the form and granularity of the cells has to be selected appropriately. We use countries as a proxy of different cultures. We represent countries as categorical variables, as this granularity can be related to different cultures, as explained in Section 2.2. However, other representations could also be suitable, such as regions etc.

Note: Other cultural variables. This is an optional category, which can include other parameters related to culture. For instance, to determine socio-cultural context, opening hours and the price range of corresponding services may also be important. Both of these variables could be discretised and considered categorical variables.

4.2 Training

Once we vectorise our attributes, as explained in the previous section, we use a concatenation layer to combine them.

If a is a POI attribute such that $a \in A$ and $\phi_a(x_a) \in \Re^{D_a}$ is the attribute specific vectorisation function, where D_a denotes the dimensionality associated with the attribute a, then the final input vector is a concatenation of all vectorised individual attributes:

$$\tilde{\boldsymbol{x}} = [\phi_1(a_1), \phi_2(a_2), ..., \phi_n(a_n)]$$

where n denotes the number of attributes. We feed this to a dense layer

$$\boldsymbol{h} = relu[\boldsymbol{W^h}\tilde{\boldsymbol{x}} + \boldsymbol{b^h}]$$

After applying a dropout layer, we then calculate

$$p(\boldsymbol{y}|\boldsymbol{h}, \boldsymbol{\theta}) = sigmoid[\boldsymbol{W}\boldsymbol{h} + \boldsymbol{b}]$$

where $\boldsymbol{\theta} = (\boldsymbol{W}, \boldsymbol{b}, \boldsymbol{W^h}, \boldsymbol{b^h})$ are learned parameters of the model. $sigmoid(s)$ denotes the logistic function $f(s_i) = \frac{1}{1+e^{s_i}}$. The parameters $\boldsymbol{\theta}$ are learned by minimising the binary cross-entropy loss function.

4.3 Inference

The multi-label model learned is given culture-specific inputs at this step. It then generates for each label a probability score. To get from that the corresponding set of accepted labels, a constant can be applied as threshold (usually this is 0.5) [4].

5 EXPERIMENTS

5.1 Set up

5.1.1 Data. We perform our experiments on 2.4M POIs extracted from a large database provided by Foursquare to illustrate the feasibility of our method. Details are provided below.

Categorisation hierarchy. Our dataset includes 808 POI categories from the categorisation hierarchy of Foursquare[6], as shown in Table 1. As the classification hierarchy is based on crowdsourced data, the parts of the dataset that include more POI instances are represented with more categories, resulting in it being heavily imbalanced. For instance, the most well developed category that is located at the root of the hierarchy is *Food*,

[5] In our experiments, 7% of them were categorised as "Noodle House" c.f. Section 5.

[6] https://developer.foursquare.com/docs/resources/categories as of 3rd October 2018

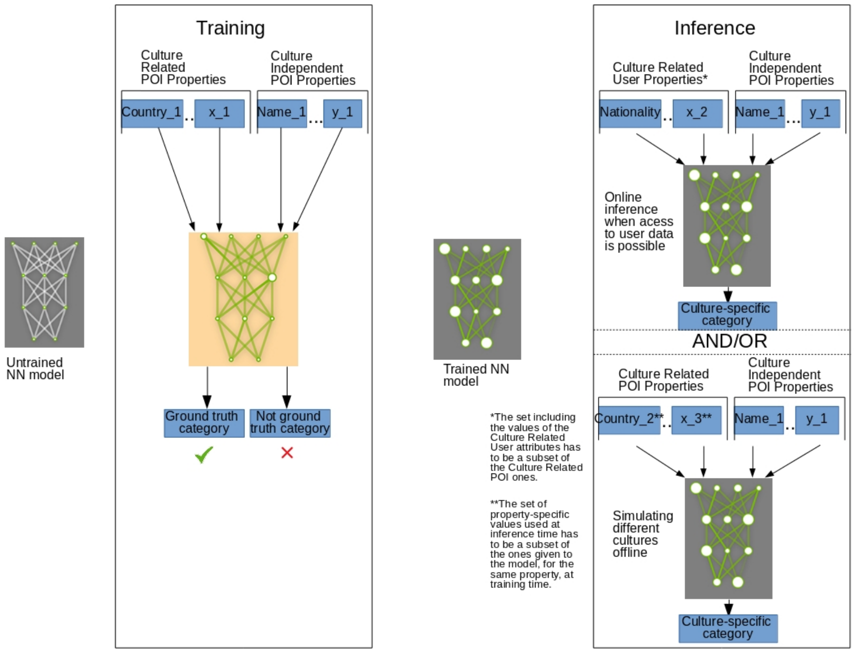

Figure 2: Overview of proposed method. Images of NN models are adapted from [3]

Table 1: Category Distribution in the dataset

Root Category	Levels	Categories in Path
Food	5	337
Shop & Service	3	144
Outdoors & Recreation	4	83
Professional & Other Places	3	77
Arts & Entertainment	3	53
Travel & Transport	3	44
College & University	3	32
Nightlife Spot	3	25
Event	2	8
Residence	2	5

Table 2: Percentage & distribution of semantic tags

Category	POIs (%)
Café	9.83
Restaurant	5.88
Pizza Place	4.49
Coffee Shop	4.48
Bakery	3.8
Fast Food Restaurant	3.23
...	...
Chinese Restaurant	2.98
Japanese Restaurant	2.5
Asian Restaurant	2.41
...	...
Noodle House	1.02
...	...
Ramen Restaurant	0.65
...	...

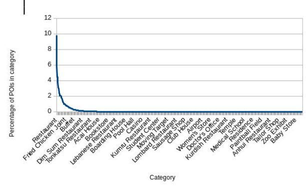

with 336 categories distributed in 5 levels. The least developed one is *Residence* having only 4 subcategories, over 2 levels[7].

Label distribution. We have extracted a sample dataset from the existing Foursquare database for our experiments. We used the most well developed root category, *Food*, as seed, and took only POIs with a high reality index. The distribution of the categories is similar to the one found in the original hierarchy. To better understand the dataset we provide the overall distribution of POIs and the one of the top categories in Table 2. The dataset is skewed in terms of the POI instances attributed to each category, with the first 10 top categories having more than 40% of the POIs attributed to them.

Point-of-Interest attributes. . POI attributes include its name, and the latitude and longitude, transformed into the representation discussed in Section 4.1. It is important to note that contrary to completely freely crowdsourced POI databases such as OpenStreetMaps [12], the format is normalised. Latitude and longtitude are written in the standard form, normally with >10-decimal point precision (e.g. latitude:55.76942424341726, longtitude: 44.948036880105064). Comments are also available for some of the POIs but as they are relatively sparse we chose not to use them for the current experiments.

Dataset creation. To generate training, development, and test data, we used approximate stratified sampling. The goal was to maintain the distribution of positive and negative examples of each label by considering each label independently. Consequently, we allocated the POIs proportionally into 40% for training, 10% for development, and 50% for testing purposes[8].

[7]Even if we used *Food* as seed, we find also the rest of the root categories in the dataset. The reason is that POIs can be categorised using multiple labels, although at least one of the labels must have as seed *Food*.

[8]We kept a relatively large percentage of the data for testing pusrpose, in order to have a large enough sample of POIs belonging to long tail categories, where cultural differences may be more obvious.

5.1.2 Model. We use a neural architecture with one hidden dense layer, followed by a dropout layer, and the output layer. We use the Rectified Linear Unit as the activation function of the hidden layer. The dropout rate is set to 0.3. The loss we use is binary crossentropy. We have set an early stopping criterion for the training based on a pre-defined threshold that takes into account the delta of the loss between two consecutive epochs. For all sequential features we applied a length of 50. For the LSTM layer we set the dimensions of the embedding layer vector space to 128 and the number of the LSTM hidden units to 128. The LSTM has a recurrent droupout rate of 0.3. Experiments were run on a single GPU instance (1 GPU with 16GB VRAM, 4 CPUs, with 256GB RAM). Training was performed with a batch size of 32. We used the Adam optimiser with the default parameters recommended in the paper [10]. All the results reported below are on the test dataset, which has not been used for training or validation purposes.

5.2 Results

Table 3 shows shows an excerpt of how POIs are categorised in different cultures. As mentioned in previous sections, we use as a proxy to represent culture, one's country. For each category included in Table 3, we provide observations related to the prediction differences between different cultures. Based on the results we can note the following.

- The model learned that some categories are by design (as defined by Foursquare) allowed only in specific countries e.g. *Acai House* and *Churrascaria* only in Brazil.
- The predictions reflect, to the best of our knowledge, local culture reasonably well e.g. *Bistro* is popular in France and *Souvlaki Shop* in Greece.
- Semantically similar categories are found by the classifier between different cultures, even when the original category is specific to one culture only. For instance *Pastelaria*, a typical Brazilian and Portuguese POI category, is predicted in other cultures as *Bakery* and/or *Snack Place*, which is correct. Similarly, *Churrascaria* is classified as *BBQ Joint* in other cultures.

Latent similarities discovered by the model are questionable though in some cases. For instance, *Souvlaki Shop*, a prominent category in the Greek culture, is correlated to the category *Kebab Shop* in other cultures. In reality, in a number of aspects there are strong similarities: for instance the form of the corresponding sandwiches (pita-based and circular), the meat that is cooked on a spit, and the "fast food" type of food delivery. However, the former is made traditionally out of pork and the latter does not have pork, which is an important difference. In future work, we plan to tackle this aspect further.

6 CONCLUSIONS

We have presented a new method to predict, in a culture-specific manner, POI categories, without requiring access to user information. To achieve that, we simulate user information, by replacing culture-related training inputs in an appropriate manner, at inference time. For instance, the country where a POI is located can be replaced at inference time the by the nationality of the user. We have performed preliminary experiments on data of a global location-based social network, Foursquare, that give us promising results. In future work, these results will be further verified with user studies.

REFERENCES

[1] Felix Biessmann, David Salinas, Sebastian Schelter, Philipp Schmidt, and Dustin Lange. 2018. "Deep" Learning for Missing Value Imputationin Tables with Non-Numerical Data. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. ACM, New York, NY, USA, 2017–2025. https://doi.org/10.1145/3269206.3272005

[2] Amparo Elizabeth Cano, Andrea Varga, and Fabio Ciravegna. 2011. Volatile Classification of Point of Interests based on Social Activity Streams. In *Proceedings of the 4th International Workshop on Social Data on the Web, SDoW@ISWC 2011, Bonn, Germany, October 23, 2011 (CEUR Workshop Proceedings)*, Alexandre Passant, Sergio Fernández, John G. Breslin, and Uldis Bojars (Eds.), Vol. 830. CEUR-WS.org. http://ceur-ws.org/Vol-830/sdow2011_paper_7.pdf

[3] Exxact. 2017. Discover the Difference Between Deep Learning Training and Inference. https://blog.exxactcorp.com/discover-difference-deep-learning-training-inference/. Accessed: 2019-11-06.

[4] Ouadie Gharroudi. 2017. *Ensemble multi-label learning in supervised and semi-supervised settings*. Ph.D. Dissertation. Université de Lyon.

[5] Tieke He, Hongzhi Yin, Zhenyu Chen, Xiaofang Zhou, Shazia Sadiq, and Bin Luo. 2016. A Spatial-Temporal Topic Model for the Semantic Annotation of POIs in LBSNs. *ACM Trans. Intell. Syst. Technol.* 8, 1, Article 12 (July 2016), 24 pages. https://doi.org/10.1145/2905373

[6] Hofstede Gert J. Minkov Michael Hofstede, Geert. 2010. *Cultures and organizations: Software of the mind*. McGraw-Hill, New York, 6, 18.

[7] Hannu Jaakkola and Bernhard Thalheim. 2017. *Supporting culture-aware information search*. Frontiers in Artificial Intelligence and Applications, Vol. 292. IOS Press, Netherlands, 161–181. https://doi.org/10.3233/978-1-61499-720-7-161 JUFOID=56381.

[8] Shan Jiang, Ana Alves, Filipe Rodrigues, Joseph Ferreira, and Francisco C. Pereira. 2015. Mining point-of-interest data from social networks for urban land use classification and disaggregation. *Computers, Environment and Urban Systems* 53 (2015), 36 – 46. https://doi.org/10.1016/j.compenvurbsys.2014.12.001 Special Issue on Volunteered Geographic Information.

[9] Armand Joulin, Edouard Grave, Piotr Bojanowski, Maximilian Nickel, and Tomas Mikolov. 2017. Fast Linear Model for Knowledge Graph Embeddings. *arXiv e-prints*, Article arXiv:1710.10881 (Oct 2017), arXiv:1710.10881 pages. arXiv:stat.ML/1710.10881

[10] Diederik Kingma and Jimmy Ba. 2015. Adam: a method for stochastic optimization (2014). *arXiv preprint arXiv:1412.6980* 15 (2015).

[11] John Krumm and Dany Rouhana. 2013. Placer: Semantic Place Labels from Diary Data. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, New York, NY, USA, 163–172. https://doi.org/10.1145/2493432.2493504

[12] OpenStreetMap contributors. 2017. Planet dump retrieved from https://planet.osm.org . https://www.openstreetmap.org.

[13] Tsangaratos Paraskevas, Rozos Dimitrios, and Benardos Andreas. 2014. Use of artificial neural network for spatial rainfall analysis. *Journal of Earth System Science* 123, 3 (01 Apr 2014), 457–465. https://doi.org/10.1007/s12040-014-0417-0

[14] Markus Schedl and Dominik Schnitzer. 2014. Location-Aware Music Artist Recommendation. In *Proceedings of the 20th Anniversary International Conference on MultiMedia Modeling - Volume 8326 (MMM 2014)*. Springer-Verlag New York, Inc., New York, NY, USA, 205–213. https://doi.org/10.1007/978-3-319-04117-9_19

[15] Yan Wang, Zongxu Qin, Jun Pang, Yang Zhang, and Jin Xin. 2017. Semantic Annotation for Places in LBSN Through Graph Embedding. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '17)*. ACM, New York, NY, USA, 2343–2346. https://doi.org/10.1145/3132847.3133075

[16] Mao Ye, Dong Shou, Wang-Chien Lee, Peifeng Yin, and Krzysztof Janowicz. 2011. On the Semantic Annotation of Places in Location-based Social Networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '11)*. ACM, New York, NY, USA, 520–528. https://doi.org/10.1145/2020408.2020491

[17] Eva Zangerle, Martin Pichl, and Markus Schedl. 2018. Culture-Aware Music Recommendation. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization (UMAP '18)*. ACM, New York, NY, USA, 357–358. https://doi.org/10.1145/3209219.3209258

[18] Junbo Zhang, Yu Zheng, Dekang Qi, Ruiyuan Li, and Xiuwen Yi. 2016. DNN-based Prediction Model for Spatio-temporal Data. In *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPACIAL '16)*. ACM, New York, NY, USA, Article 92, 4 pages. https://doi.org/10.1145/2996913.2997016

[19] Xin Zheng, Jialong Han, and Aixin Sun. 2018. A Survey of Location Prediction on Twitter. *IEEE Transactions on Knowledge and Data Engineering* 30, 9 (Sep. 2018), 1652–1671. https://doi.org/10.1109/TKDE.2018.2807840

[20] Jingbo Zhou, Shan Gou, Renjun Hu, Dongxiang Zhang, Jin Xu, Xuehui Wu, Airong Jiang, and Hui Xiong. 2019. A Collaborative Learning Framework to Tag Refinement for Points of Interest. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. ACM, New York, NY, USA, 1752–1761. https://doi.org/10.1145/3292500.3330698

Table 3: Culture-specific prediction results for different POI categories. Green coloured values are significantly higher than in the the rest of the cultures, and red significantly lower, indicating a notable culture-specific influence.

Category	Original data	Culture					
		KR	FR	US	BR	TR	GR
Acai House	1488	43	29	0	1534	0	0
Note:Except for Brazil in the rest of the cultures the same POIs are categorised as Snack Place, Juice Bar, Dessert Shop. According to Foursquare's documentation Acai house is a category only supported in Brazil.							
Bistro	879	470	3310	351	1162	95	10
Note: Bistros predicted using the French culture are tagged in the original data as: Café, Bar, Gastropub, Diner. Corresponding predictions in other cultures are: Café, Wine Bar, Bar, Gastropub.							
Brasserie	18	0	91	0	3	0	0
Note: In other cultures Café is the main predicted category for the same POIs (or there is no prediction at all). In the silver standard the POIs are also categorised as Bistro or Café.							
Café	126665	108908	87381	63223	108837	148050	236436
Note: In the US culture a lot of Cafes seem to be categorised as Coffee Shops instead. In the Greek culture Coffee Shop, Breakfast Place, Dessert Shop, Bar, Tea Room POIs are categorised as Café (which is actually representative of the culture).							
Coffee Shop	54291	49069	45769	68102	50513	52742	23629
Note: As explained in the previous row.							
Churrascaria	557	0	0	0	674	0	0
Note: Churrascaria is a Portuguese/Brazilian BBQ. In other cultures the majority of the same POIs are classified as BBQ Joint and a small percentage as Steakhouse (especially in the US).							
Creperie	978	775	2968	932	1138	917	1740
Note: Creperies are obviously common place in the French culture. In the US and KR ones the same POIs are rather categorised as Dessert Shop or Breakfast Shop. In the BR one in addition to Dessert Shop we find also Pastelaria[9].							
Dessert Shop	15164	19422	4747	9651	15460	21941	13736
Note: In the French culture Dessert Shop POIs are rather classified as Café, Bakery, Creperie, Pastry Shop, Chocolate Shop. In the US one we have to note the large number of POIs categorised as Ice Cream Shop, Frozen Yogurt Shop, Candy Store.							
Diner	2590	2289	2388	3980	1539	1593	100
Note: Some of the POIs categorised in the US culture as Diner, are mainly categorised in other cultures as Café or Breakfast Spot or there is no prediction (the difference is really big with Greece where almost all of them are categorised as Café).							
Friterie	656	7	4864	2	419	4	62
Note: The model has learned that Friterie is a culture-specific category (supported in FR, BE, NL in the Foursquare database). It is interesting to note that in the US culture the same POIs are categorised as Burrito Place, Taco Place, Food Truck, in the TR one as Kofte Place and in the GR one as Snack Place. They all seem to share a fast food aspect.							
Meyhane	854	0	0	0	0	2157	0
Note: The model has learned that Meyhane is TR culture-specific category. In other cultures Meyane POIs usually do not get any prediction.							
Pastelaria	340	1	0	0	898	57	0
Note: The model has learned that Pastelaria is BR (and Portuguese) culture-specific category. In FR culture categorised as Creperie, Snack Place, Bakery, US:Bakery,Snack Place, KR:Bakery, Dessert Shop, Snack Place.							
Pastry Shop	60	5	307	0	0	88	0
Note: The model has learned that Pastry Shop is more frequent in FR (i.e. Patisserie). In other cultures we mainly find Dessert Shop or Bakery.							
Souvlaki Shop	94	0	0	0	0	0	2434
Note: The model learned that Souvlaki Shop is specific to GR culture. In other cultures and in the silver standard, the same POIs are categorised mainly as Fried Chicken Joint and BBQ joint. There is a strong correlation to Steakhouse as well, or more precisely to Kebab shops that are also classified as Steakhouse in the silver standard.							
Sports Bar	67	9	2	100	4	5	0
Note: The model learned that Sports Bar is more frequent in the US culture. In the silver standard, the same POIs are categorised also just as Bar and/or Wing Joint. Furthermore, the model learned a strong correlation between the category Wings Joint and sports Bar - the two categories are predicted together 77% of the time.							
Takoyaki Place	186	130	45	34	22	44	30
Note: Takoyaki Place POIs do not get any prediction most of the times in other cultures except for KR. 86% of the predicted POIs are correct according to the silver standard.							