Multi-channel Convolutional Neural Networks for Handling Multi-dimensional Semantic Trajectories and Predicting Future Semantic Locations

Antonios Karatzoglou^(✉)

Robert Bosch GmbH, Chassis Systems Control, Advance Engineering Sector,
Abstatt, Germany

antonios.karatzoglou@de.bosch.com

Abstract. Current location-aware systems rely increasingly on location prediction techniques in order to provide their services in a timely fashion. At the same time, it has been shown that the use of additional context information, that is, elevating the degree of semantic enrichment of movement data, can lead to significant better results both in analyzing as well as in modeling human trajectories and predicting upon them. In this work, we propose a Multi-Channel Convolutional Neural Network (CNN) based approach for capturing all the available context dimensions in our semantic trajectory dataset aiming at achieving a higher prediction accuracy compared to a vanilla locations-only Single-Channel CNN. Moreover, we investigate whether and to what degree time, activity, companionship and the user's emotional state have an impact on the predictive performance of our multi-dimensional CNN. We evaluate our model using a real-world dataset and compare it among others to a probabilistic Markov Chain model and a vanilla CNN at two semantic representation levels. It can be shown that especially for higher level representations, the present approach is able to outperform the baseline models achieving an overall higher accuracy and F1-Score.

Keywords: Semantic locations and trajectories · Location prediction · Purpose of visit · Human activities · Emotional states · Context awareness · Multi-channel convolutional neural networks

1 Introduction

The market of context-aware, and especially, location-aware computing and services (see *Location-Based Services (LBS)*) has gained enormously in importance over the past few decades. In their attempt to provide timely solutions to their users, LBS providers rely more and more on location prediction methods, a fact that additionally strengthened the demand for accurate location prediction algorithms.

Typical location prediction models are usually merely data driven and depend therefore heavily on the size and the quality of the available training datasets. However, recent research has shown that the use of additional semantic information can help overcome, at least to some extent, the aforementioned dependencies and can therefore lead to an overall better predictive performance (see Sect. 2). That is, in the case of modeling and learning human movement patterns, models are usually fed and trained with the users' plain GPS location trajectories. Further context information, such as the location type and the user's activity, may be however used to enrich these semantically and generate so-called *semantic trajectories* (see Sect. 3). This type of extensive input helps the model dive even deeper into the users's movement behaviour and can lead to more accurate predictions.

Common approaches used to model and predict human movement include probabilistic methods, such as Markov Chains [?] [9], Dynamic Bayes Networks [8], Hidden Markov Models [26,31] and Artificial Neural Networks (ANNs). In the latter case, recurrent neural network architectures (RNNs) have generally proved to perform above the average when it comes to learning sequences and for this reason these are commonly found in the location prediction domain as well. Especially memory-based neural network types, like the Long Short-Term Memory network (LSTM), are capable of achieving high prediction rates and tend to outperform the competition [16,28].

While recurrent network types are the preferred choice when it comes to modeling 1-dim movement patterns, recent work shows some promising results on the part of Convolutional Neural Networks (CNNs) [17,25], a model normally used in the 2-dim image classification and object recognition domain. It seems that the locally focused nature of the kernel-based convolution process enables the CNN model to successfully capture existing dependencies between current and future locations found in the data. The presented study builds upon this work and aims at investigating the use of a *multi-channel CNN* based approach with regard to modeling *multi-dimensional semantically enriched location data* and predicting the next semantic location of the user. In particular, our semantic trajectories consist of the following feature dimensions: *semantic location type, time, human activity, emotional state and companionship*. Moreover, this work further explores the impact of the degree of semantic enrichment, that is, whether and to what extent each of the aforementioned dimensions influences the predictive performance of our model. We evaluated our approach using a real-world dataset, which we collected from 21 users by conducting a 2-months long user study. In addition, we selected a 1. Order Markov Chain model and a vanilla CNN to be our baseline.

This paper is structured as follows. Section 2 provides a short overview over some of the most related work in the semantic trajectories and location prediction domain. Next, Sect. 3 describes the notion of *semantic trajectories* and *semantic locations* with respect to this work. Section 4 goes briefly through the theory behind Convolutional Neural Networks and discusses in detail the proposed

approach, while Sect. 5 provides the respective evaluation outcomes. Finally, Sect. 6 summarizes the evaluation results and draws some final conclusions.

2 Related Work

There exist many different ways of viewing movement data. Within the scope of mining and analyzing movement patterns, Spaccapietra et al. introduced with [29] one of the first works that make the importance of viewing trajectories of moving objects in a conceptual manner clear. In their work, they highlighted the fact that describing certain aspects of the movement's context by adding semantic information into the available trajectories can significantly support the analysis of the respective movement patterns, as well as the querying process among them. Alvarez et al. came to the same conclusion as they suggested the use of a similar semantic enrichment model to generate *semantic trajectories* for the same reasons [1]. The added value of working on semantically enriched trajectory data in comparison to working on raw data with regard to mining patterns and supporting decision processes has been underpinned by Elragal et al. as well [7]. Bogorny et al.'s work also focuses in mining trajectory patterns and introduced in [2] a sophisticated model, which in contrast to former models is capable of handling complex queries over semantic trajectories, while providing different semantic granularities at the same time. Finally, Karatzoglou et al. showed in [20] that considering multiple context dimensions results in generating more accurate synthetic semantic trajectories.

Due to the aforementioned benefits that accrue from semantic enrichment, a number of location prediction papers have recently emerged presenting algorithms that rely on the notion of semantic trajectories. Ying et al. for instance were one of the first to build upon semantic trajectories generated from the users' raw GPS recordings in order to enhance their location prediction framework [32] with promising results. Some years later, they extended their model by taking, apart from geographic and semantic patterns, temporal patterns into account as well [33].

Karatzoglou et al.'s work explores a big variety of models with respect to modeling human semantic trajectories and predicting the user's next semantic location. In [12] and in [18] they evaluate a multi-dimensional Markov Chain model with respect to predicting among activity-enriched semantic trajectories and show that it is able to outperform Ying et al.'s framework in terms of accuracy. With regard to recall however, they could identify certain limitations on behalf of the model due to its adverse dependency on the small size and the sparsity of the available training dataset. They attempt to solve this issue by combining the probabilistic Markov Chain model with Matrix Factorization in [11], where they were able to raise the recall scores.

In [13, 16, 17, 19], Karatzoglou et al. investigate the performance of Artificial Neural Networks using the probabilistic Markov model as baseline. In addition, they explore the role of the semantic granularity of the considered trajectories in the overall performance of the networks. They show that the higher the semantic level is, the better the modeling quality of the networks. While the findings

in [16] comply with the results of related work showing that Long Short-Term Memory networks are generally able to outperform the vanilla Recurrent (RNN) and the Feed-Forward model, [19] indicates no great advantages towards the attention-based application of Sequence to Sequence learning (Seq2Seq) compared to the standard single-input-single-output LSTM model of [16], a fact that may primarily explained by the limited size of the training dataset. Yao et al. propose in [30] a similar to [16] LSTM-based recurrent approach for predicting next semantic locations using an additional embedding input layer, the benefits of which have been also recently shown by Gao et al. in [10]. Other than in [16] and following a similar direction to the approach proposed in the present paper, Yao et al. used beyond location and time the content of the checkins of the users to enrich the users' semantic trajectories, which describe in a way their activity that we're considering in this work as well (among other). However, in contrast to the Reality Mining dataset [6] used in Karatzoglou et al.'s work, they evaluate their approach on rather long-term dependencies using a Foursquare and a Twitter dataset. In [13], Karatzoglou et al. take a look at a gradient-free optimization method for finding the optimal hyperparameter set of a LSTM model based on an evolutionary algorithm. Their work provides some first preliminary results indicating among others the temporal efficiency on part of the genetic, population-based optimization method, provided the fact that sufficient computational power is available.

The most striking findings come rather from [17], where a Convolutional Neural Network based approach in combination with an embedding layer as its input is capable of achieving higher prediction scores than the FFNN, the RNN and the LSTM. To our knowledge this represents the only work that explores the use of CNNs with respect to modeling and predicting upon 1-dim semantic trajectories. The closest work to [17] would be the work of Lv et al. in [25], which evaluates the use of a CNN for modeling and predicting large-scale taxi trajectories. Other than in [17] and the present paper, Lv et al. work with raw GPS data without using any semantic information and map past trajectory data onto 2-dim images before feeding them into the CNN model, transforming in this way the trajectory modelling task into an image classification task. In the present paper, following the example of [17], we skip this kind of 1-dim to 2-dim intermediate transformation step and apply our CNN model on the 1-dim semantic trajectory as it is. As in [17], we build our approach upon similar CNN-based work on 1-dim data, work, that comes mostly from the Natural Language Processing (NLP) domain, such as the framework described in [4] and the multi-channel CNN model of [22].

3 Semantic Trajectories

The term *trajectory* refers to a sequence of consecutive location points traversed by a moving object within a certain time interval. Equation 1 describes a typical GPS trajectory with each location point being represented by a tuple containing its coordinates ($long_i$, lat_i) and the corresponding point of time t_i at which this was visited.

$$Traj_{GPS} = (long_1, lat_1, t_1), (long_2, lat_2, t_2), \dots, (long_i, lat_i, t_i), \quad (1)$$

As already mentioned in Sect. 2, in order to better understand the moving behaviour of moving objects and create more accurate models, Spaccapietra et al. [29] and Alvares et al. [1] went beyond this kind of numerical sequences by adding a semantic view upon them and introduced the so-called *semantic trajectories*. Starting initially with the simple notion of “stops” and “moves”, a (human) semantic trajectory can nowadays be defined generally as a sequence of semantically significant locations (*semantic locations*, e.g., “home”, “burger joint”, etc.) as follows:

$$Traj_{Sem} = (SemLoc_1, t_1), (SemLoc_2, t_2), \dots, (SemLoc_i, t_i) \quad (2)$$

A significant location in this case is usually defined by a location within a certain radius (e.g., 200 m) where a user stays longer than a pre-defined temporal threshold, e.g. 20 min (see [?]). Some researchers add further thresholds, like the loss of the GPS signal due to entering into a building, the GPS recording stop [3] or the popularity, in order to extract the most significant common or public locations [32]. In this work, we evaluate our method using a dataset in which the users annotated their longest visits (>15 min) by themselves (see Sect. 5).

Depending on the number of the considered semantic features, a semantic trajectory can have multiple dimensions. Thus, we could say that the number of dimensions expresses the *degree of semantic enrichment* of the respective semantic trajectory. In this work, we follow the concept of the Location-Specific Cognitive Frames introduced by Karatzoglou et al. in [14, 15] and we consider each stop at a semantic location to be a tuple encapsulating the current location type, the current time, the current activity, as well as the user’s current emotional state and whether he is alone or not (companionship). Beyond that, locations can be described differently depending on the semantic representation level, e.g., “restaurant” → “fast food restaurant” → “burger joint”. In this work, we evaluate the modeling performance of a CNN at two different semantic levels. That is, we evaluate two different models, one that is trained for handling and predicting low level trajectories and one for handling higher level ones.

4 Multi-channel Convolutional Neural Networks on Semantic Trajectories

This section consists of two parts. The first part gives a brief insight into the theory behind Convolutional Neural Networks and goes briefly through some of the most common CNN steps and layers using the example of image classification. Then, the second and last part describes in detail the architecture of the multi-channel CNN model proposed in this paper for handling multi-dimensional semantic trajectories.

Convolutional Neural Networks (CNN) constitute the state of the art choice in the image classification and object recognition domain [24]. However, this doesn't mean that it is the only domain in which we can apply them expecting reasonable results as we saw in Sect. 2 and can also be seen in [23]. Figure 1 illustrates a typical CNN pipeline used for classifying images.

Fig. 1. Typical CNN architecture for the image classification task (source: [27]).

A typical CNN consists of many different layers starting usually with the (first) convolutional layer. This layer is responsible for convolving the input image and generating the so-called *feature maps*. This is done by sliding a group of small-sized *filters* (also called *kernels*) with each containing a certain number of learnable weights over the input image and performing element-wise multiplication at each possible position. The generated feature map from each kernel is a new layer and contains the findings of the particular kernel in the input image, ideally with respect to a specific and distinguished single feature. The number of kernels defines the number of the generated feature maps (so-called *depth* of the convolutional layer) and represents a CNN hyperparameter which needs to be selected appropriately based on the available data and task. In the next step, this resulting group of layers undergoes a so-called *pooling* process. Pooling refers to a downsampling operation, in which sets of elements in the feature maps are combined and reduced to a single value based on some criterion (e.g., take the maximum value: *max pooling*) or on some type of calculation (e.g., take the average over all values: *average pooling*). The two aforementioned layers can be repeated multiple times using different kernels of different size and depth. This supports the successive extraction of higher level features and represents one of the strengths of CNNs. Finally, the last pooled layer can be *flattened* into a single vector containing all its weights and connected to a *fully connected layer*, which is further connected to the output layer that contains a field for every possible class and provides us with the classification estimation for the given input.

The multi-channel approach introduced in this paper builds upon the aforementioned typical CNN architecture and extends it by adding a further *embedding layer* into the model and by *raising the number of channels* matching

Fig. 2. Multi-channel CNN for modeling multi-dimensional semantic trajectories.

the degree of semantic enrichment of our data (see Sect. 3). Figure 2 illustrates the architecture of our approach. Our framework takes as input a part of a semantic trajectory, which is in our case a sequence of tuples in the form of $(locationtype, purposeofvisit, time, emotionalstate, companionship)$ according to a predefined temporal horizon t_n that determines how far backwards in the movement history of the user should the model consider for providing an estimation about her next semantic location. In a first step, every single feature type is encoded as a one-hot vector. The additional embedding layer between the one-hot encoded input and the convolutional layer maps the sparse asymmetrical one-hot encoded binary vectors into dense vector representations in a continuous vector space. This fact contributes to a more efficient training and helps improving the prediction accuracy while keeping the model consistent at the same time. The number of dimensions of the vector space is selected based on the properties of the available data, e.g., the number of unique classes of the corresponding feature. In our case, each semantic feature is encoded separately, and therefore the generated vectors may have different number of dimensions.

Raising the number of channels according to the semantic enrichment degree of our trajectories represents an intuitive way of viewing upon them. Each channel handles solely the corresponding semantic dimension. For example the first channel is responsible for the location type, the second channel for the purpose of visiting that location (activity), the third one for covering temporal information and so on. At the end, all channels are merged into a single representation, flattened and forwarded to the output layer in order to provide a final prediction about the next semantic location of the user. It should be noted here that the kernels' depth should match the channel dimension (5 in our case).

Our CNN has one convolutional, one pooling, a flattening, a fully connected and a Softmax output layer. A deeper architecture, that is, adding more layers led in most of the cases to overfitting and reduced the overall performance of our model due to the small size of our dataset compared to the higher parameter number. Other than the CNN model in Fig. 1, our model executes a 1-dim convolution operation instead of the typical 2-dim operation conducted in the image classification task. Each kernel convolves each semantic dimension in one direction only, namely according the chronological order found in the input tuple sequence which is fed into the model. Thus, the width of each kernel covers the whole row of the CNN input matrices while its height can vary, constituting a further hyperparameter of our model. A higher height indicates a kernel, able to observe a higher number of consecutive locations at the same time, a fact that can be useful when aiming at capturing long-term dependencies in our data, and vice versa.

5 Evaluation

In this section, we evaluate a multi-user version of our approach, which is trained on location data coming from multiple users. For this purpose, we first concatenated the trajectories of all users to a single trajectory, ordered by the user ID. Then, we randomly split the resulting trajectory into a training and a test dataset with a ratio of 80% to 20% while maintaining the user order at the same time (i.e., without breaking a user’s trajectory into 2 parts). All in all, we randomly split the data 3 times and the findings in this section refer to the average over these 3 runs.

In order to evaluate our approach, we conducted a 8-week long user study tracking 21 users via a tracking and annotation app. The participants of the study were asked to semantically label each significant stop (location type) during their movement, as well as to note the purpose of visiting the certain location (e.g., activity), their companion (if any) and their emotional state by selecting among the following states: happy, hungry, neutral, sleepy, energetic, frustrated, stressed, bored, adventurous, ill, sad, angry and shocked. At the end of the study we end up with approximately 1400 annotated locations covering around 70 unique location types, 53 unique activities, and all 13 emotional states. A thorough description of the user study can be found in [21].

In order to take time into account, we defined $24 \times 7 = 168$ hourly slots, which similar to the other input signals were one-hot encoded first and transformed into an embedding vector in a next step. However, our evaluation results showed that taking time into account had a severe negative impact on the prediction outcome of our model. We saw a similar behaviour in the work of Karatzoglou et al. in [16] and in [18]. This can be mainly attributed to the small size of our dataset which makes it extremely hard for the model to find temporal patterns in this 168-slot temporal granularity. The use of wider time slots, e.g., the use of just daily slots, couldn’t yield significantly higher scores either, due to the fact that our 8-week long evaluation dataset contains solely 8 recordings from each day, that is, there

exist solely 8 unique Mondays, 8 unique Tuesdays, etc. For this reason, and due to space reasons, this evaluation section neglects to refer thoroughly to the individual results with respect to time. In addition, our users provided very little information regarding their type of companionship (e.g., relative, friend, etc.). Solely the fact whether a user was alone or not can be reliably extracted from our dataset. Therefore, instead of handling the companionship in a separate channel, we integrated the particular information into the emotional state one-hot vector by extending it to a further dimension ('0', when the user is alone and '1' when he is not). Finally and as already mentioned previously in this work, we evaluated our approach at two semantic representation levels, which will be referred to as *low* and *high* level, with the latter being more abstract and subsuming the first one. In order to generate these two layers, we built a corresponding location taxonomy based on the Foursquare venue categorization¹. Lastly, a Grid Search helped us to determine the following optimal hyperparameter configuration for our model listed in Table 1.

Table 1. Optimal hyperparameter set determined via Grid Search.

Kernel size	Number of kernels	Embedding dimension	Dropout probability	Batch size	Learning rate	Number of epochs	Sequence length	Pooling size
6	64	100	0.6	16	0.001	100	10	2

Fig. 3. Accuracy and F1-Scores at the higher semantic representation level.

Figure 3 compares the result from 5 different models at the higher representation level, a standard 1-channel CNN (*Location*) that takes just the current semantic location as input, a 2-channel CNN that considers the location

¹ <https://developer.foursquare.com/docs/resources/categories>.

Fig. 4. Accuracy and F1-Scores at the lower semantic representation level.

Fig. 5. Training accuracy and loss curves at the higher representation level. (a): Location, (b): Location&Companion&Mood, (c): Location&Purposes, (d): Location&Purposes&Companion&Mood.

type and the purpose of visit (*Location&Purpose*), a 2-channel CNN that considers the location type and the emotional state as well as the companionship status of the user (*Location&Companion&Mood*), a 3-channel CNN that takes location type, purpose of visit, emotional state and companionship (*Location&Purposes&Companion&Mood*) into account, and a probabilistic Markov

Fig. 6. Training accuracy and loss curves at the lower representation level. (a): Location, (b): Location&Companion&Mood, (c): Location&Purposes, (d): Location&Purposes&Companion&Mood.

Chain model of 1. order that serves as our reference. It can be seen that all the CNN-based approaches are able to outperform the Markov model both in terms of accuracy and F1-Score. What also stands out in the same figure is that the 2-channel CNN approach that considers the activity of the user (purpose of visit) can clearly outperform the competition. However, this doesn't hold for the other 2-channel CNN model. On the contrary, it seems that taking the user's emotional state into account affects negatively the predictive performance. Apparently, our model wasn't able to establish an association between the users' movement behaviour and their mood, a fact that could be partly attributed once again to the small size of our dataset. The more "sophisticated" 3-channel CNN achieves a similar accuracy to the standard CNN, but a lower F1-Score and therefore can't really compete with the *Locations&Purposes* model. Its results are likely to be related to the aforementioned negative impact of the emotional state when this is taken explicitly into account.

Figure 4 presents the results for the lower semantic representation level. It is apparent that all models perform worse than in the higher level shown in Fig. 3. This can be mainly attributed to the fact that the lower semantic representation carries a higher number of unique classes to predict, which makes the learning process of the models much harder. At the same time, another possible

Fig. 7. Training F1-Score and loss curves at the higher representation level. (a): Location, (b): Location&Companion&Mood, (c): Location&Purposes, (d): Location&Purposes&Companion&Mood.

explanation for this might be the fact that human movement shows stronger regularities at rather higher levels, e.g., a user may often visit a food location after going to gym, regardless whether this location is an Italian or a Greek restaurant, a pizza house or a burger joint. Similar to the high level case, the CNN models outperform in most of the cases the probabilistic baseline. However, this time, other than at the higher level, it seems that the additional channels result in a deterioration of our prediction models. The more channels, the worse the predictive behaviour seems to become. In general, due to the small size of our dataset and its imbalance, all of our models had to deal with massive overfitting issues. Adding a dropout layer while making our model simpler by reducing the size of our layers could improve significantly our models, but only to a certain extent.

Figures 5, 6, 7 and 8 illustrate the training behaviour of our 4 CNN models. We can see that the greater the number of channels and thus, the greater the semantic enrichment degree of the trajectory, the faster and smoother the training of the CNN model becomes. Taking additional context dimensions into account seems to contribute to shorter convergence times and results in a more efficient training. The 3-channel CNN is characterized by the shortest convergence, while the vanilla 1-channel CNN straggles with the loss reduction along the whole training process to the 100th epoch. The benefits of the multiple chan-

Fig. 8. Training F1-Score and loss curves at the lower representation level. (a): Location, (b): Location&Companion&Mood, (c): Location&Purposes, (d): Location&Purposes&Companion&Mood.

nel approach can be more clearly seen during the harder learning task, namely at the lower semantic representation level (see Fig. 6). However, on the other hand, this comes with a certain overfitting effect, as mentioned previously, that grows with the number of the CNN input channels.

The models presented in this work use 1-dim fixed-sized kernels of size 6. This number was determined by applying a Grid Search. The problem when using fixed-sized kernels is that these are able to capture only data dependencies of a certain length. Varied-sized kernels as in the work of Kim et al. in [22] could help overcome this issue and capture the individual properties of each semantic dimension in our data.

6 Conclusion

In this work, we explore the performance of a Multi-Channel Convolutional Neural Network (CNN) based approach with respect to its capability of modeling semantic trajectories at different semantic representation levels and predicting the next semantic location of a user. Moreover, we investigate whether and to what extent the degree of semantic enrichment, that is, the number of the context feature dimensions considered in the semantic trajectories, affects the predictive performance of our model. We considered 5 different semantic enrichment

dimensions for our trajectories, the location type, the purpose of visit (e.g., activity), the time, the user's mental and emotional state, and his companionship. We evaluated our model using a 8-week long real-world dataset from 21 users and compared it to a vanilla Single-Channel CNN and a probabilistic Markov Chain model that served among other as our baseline. We could show that raising the semantic enrichment degree of our trajectory data while increasing the corresponding number of channels at the same time can indeed lead to an improvement in terms of prediction accuracy and F1-Score. This could be particularly seen when we attempted to model and predict upon semantic trajectories at a higher representation level. Furthermore, the results of this work indicate a strong correlation between the degree of semantic enrichment, the number of CNN channels and the training behaviour, with our multi-channel based approach being characterized by generally much smoother and faster converging learning curves. However, our evaluation also identified some limitations leaning mostly on certain overfitting effects, which could be mainly attributed to data-specific properties, such as the small size of our dataset and its imbalance. This is also the reason why a certain uncertainty about the generalizability and the representativity of the findings in this work arises. Nevertheless, the presented study still establishes a solid basis for further work and investigations. In our future work, we plan to further explore the use of CNNs in the semantic location prediction scenario. In particular, we plan to investigate the use of varied-sized kernels and depthwise separable convolution layers aiming at improving both the predictive performance as well as the computational efficiency of our model. Furthermore, we would like to experiment with further types of context information, such as the personality of the user, the weather and the transportation mode; features, that have led to promising results in existing studies. However, the gathering of context information and especially of personal information has become increasingly difficult in recent years, among others, due to stricter data privacy regulations. One solution for overcoming this issue would be to rely on privacy preserving methods such as the semantic obfuscation techniques of [5].

References

1. Alvares, L.O., et al.: A model for enriching trajectories with semantic geographical information. In: *Proceedings of the 15th Annual ACM International Symposium on Advances in Geographic Information Systems*, p. 22. ACM (2007)
2. Bogorny, V., et al.: Constant-a conceptual data model for semantic trajectories of moving objects. *Trans. GIS* **18**(1), 66–88 (2014)
3. Cao, X., et al.: Mining significant semantic locations from GPS data. *Proc. VLDB Endow.* **3**(1–2), 1009–1020 (2010)
4. Collobert, R., et al.: Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* **12**, 2493–2537 (2011)
5. Damiani, M.L., Bertino, E., Silvestri, C.: Protecting location privacy through semantics-aware obfuscation techniques. In: Karabulut, Y., Mitchell, J., Herrmann, P., Jensen, C.D. (eds.) *IFIPTM 2008. ITIFIP*, vol. 263, pp. 231–245. Springer, Boston, MA (2008). https://doi.org/10.1007/978-0-387-09428-1_15

6. Eagle, N., Pentland, A.S.: Reality mining: sensing complex social systems. *Pers. Ubiquitous Comput.* **10**(4), 255–268 (2006)
7. Elragal, A., El-Gendy, N.: Trajectory data mining: integrating semantics. *J. Enterp. Inf. Manag.* **26**(5), 516–535 (2013)
8. Etter, V., et al.: Been there, done that: what your mobility traces reveal about your behavior (2012)
9. Gambs, S., et al.: Next place prediction using mobility Markov chains. In: *Proceedings of the First Workshop on Measurement, Privacy, and Mobility, MPM 2012*, pp. 3:1–3:6. ACM, New York (2012)
10. Gao, Q., et al.: Identifying human mobility via trajectory embeddings. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pp. 1689–1695. AAAI Press (2017)
11. Karatzoglou, A., et al.: Matrix factorization on semantic trajectories for predicting future semantic locations. In: *2017 IEEE 13th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pp. 1–7, October 2017
12. Karatzoglou, A., et al.: Purpose-of-visit-driven semantic similarity analysis on semantic trajectories for enhancing the future location prediction. In: *2018 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pp. 100–106, March 2018
13. Karatzoglou, A.: Evolutionary optimization on artificial neural networks for predicting the user’s future semantic location. In: Macintyre, J., Iliadis, L., Maglogiannis, I., Jayne, C. (eds.) *EANN 2019. CCIS*, vol. 1000, pp. 379–390. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20257-6_32
14. Karatzoglou, A., Beigl, M.: Applying situation-person-driven semantic similarity on location-specific cognitive frames for improving the location prediction. In: *8th International Conference on Knowledge Engineering and Semantic Web (KESW)* (2017)
15. Karatzoglou, A., Beigl, M.: Enhancing the affective sensitivity of location based services using situation-person-dependent semantic similarity. In: *Proceedings of the Eleventh International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies, UBIComm 2017*, pp. 95–100 (2017)
16. Karatzoglou, A., Sentürk, H., Jablonski, A., Beigl, M.: Applying artificial neural networks on two-layer semantic trajectories for predicting the next semantic location. In: Lintas, A., Rovetta, S., Verschure, P.F.M.J., Villa, A.E.P. (eds.) *ICANN 2017. LNCS*, vol. 10614, pp. 233–241. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68612-7_27
17. Karatzoglou, A., Schnell, N., Beigl, M.: A convolutional neural network approach for modeling semantic trajectories and predicting future locations. In: Kůrková, V., Manolopoulos, Y., Hammer, B., Iliadis, L., Maglogiannis, I. (eds.) *ICANN 2018. LNCS*, vol. 11139, pp. 61–72. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01418-6_7
18. Karatzoglou, A., et al.: Semantic-enhanced multi-dimensional markov chains on semantic trajectories for predicting future locations. *Sensors* **18**(10), 3582 (2018)
19. Karatzoglou, A., et al.: A Seq2Seq learning approach for modeling semantic trajectories and predicting the next location. In: *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2018*, pp. 528–531. ACM, New York (2018)

20. Karatzoglou, A., et al.: Towards an affective semantic trajectory generator (ASTG). In: 14th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications, WiMob 2018, Limassol, Cyprus, 15–17 October 2018, pp. 1–10 (2018)
21. Karatzoglou, A., et al.: Sentient destination prediction (under review). User Modeling and User-Adapted Interaction (UMUAI) (2019)
22. Kim, Y.: Convolutional neural networks for sentence classification. arXiv preprint [arXiv:1408.5882](https://arxiv.org/abs/1408.5882) (2014)
23. LeCun, Y., et al.: Convolutional networks for images, speech, and time series. In: The Handbook of Brain Theory and Neural Networks, vol. 3361, p. 10 (1995)
24. LeCun, Y., et al.: Convolutional networks and applications in vision. In: Proceedings of 2010 IEEE International Symposium on Circuits and Systems (ISCAS), pp. 253–256. IEEE (2010)
25. Lv, J., et al.: T-CONV: a convolutional neural network for multi-scale taxi trajectory prediction. arXiv preprint [arXiv:1611.07635](https://arxiv.org/abs/1611.07635) (2016)
26. Mathew, W., et al.: Predicting future locations with hidden Markov models. In: Proceedings of the 2012 ACM Conference on Ubiquitous Computing, UbiComp 2012, pp. 911–918. ACM, New York (2012)
27. Mathworks: Convolutional neural network (2018). <https://www.mathworks.com/discovery/convolutional-neural-network.html>. Accessed 19 Feb 2018
28. Song, X., et al.: DeepTransport: prediction and simulation of human mobility and transportation mode at a citywide level. IJCAI **16**, 2618–2624 (2016)
29. Spaccapietra, S., et al.: A conceptual view on trajectories. Data Knowl. Eng. **65**(1), 126–146 (2008)
30. Yao, D., et al.: SERM: a recurrent model for next location prediction in semantic trajectories. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, pp. 2411–2414. ACM, New York (2017)
31. Ye, J., et al.: What’s Your Next Move: User Activity Prediction in Location-based Social Networks, pp. 171–179 (2013)
32. Ying, J.J.C., et al.: Semantic trajectory mining for location prediction. In: Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, pp. 34–43. ACM (2011)
33. Ying, J.J.C., et al.: Mining geographic-temporal-semantic patterns in trajectories for location prediction. ACM Trans. Intell. Syst. Technol. **5**(1), 2:1–2:33 (2014)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.