

Understanding dynamics of truck co-driving networks

Gerrit Jan de Bruin^{1,2,3}, Cor J. Veenman^{1,4}, H. Jaap van den Herik², and
Frank W. Takes¹

¹ Leiden Institute of Advanced Computer Science, Leiden University, Leiden, NL,
`g.j.debruin@liacs.leidenuniv.nl`

² Leiden Centre of Data Science, Leiden University, Leiden, NL

³ Human Environment and Transport Inspectorate, Netherlands Ministry of
Infrastructure and Water Management, The Hague, NL

⁴ Data Science Department, TNO, The Hague, NL

Abstract. The goal of this paper is to learn the dynamics of truck co-driving behaviour. Understanding this behaviour is important because co-driving has a potential positive impact on the environment. In the so-called co-driving network, trucks are nodes while links indicate that two trucks frequently drive together. To understand the network’s dynamics, we use a *link prediction* approach employing a machine learning classifier. The features of the classifier can be categorized into spatio-temporal features, neighbourhood features, path features, and node features. The very different types of features allow us to understand the social processes underlying the co-driving behaviour. Our work is based on a spatio-temporal data not studied before. Data is collected from 18 million truck movements in the Netherlands. We find that co-driving behaviour is best described by using neighbourhood features, and to lesser extent by path and spatio-temporal features. Node features are deemed unimportant. Findings suggest that the dynamics of a truck co-driving network has clear social network effects.

Keywords: transport networks, mobility, co-driving behaviour, spatio-temporal networks, link prediction

1 Introduction

Nowadays an increasing volume of published studies concerning social network analysis is combined with spatio-temporal data. Much of the research performed so far used either GPS [5, 11], WiFi [13] or calls from mobile phones [16]. In this study, we analyze 18 million truck movements in the Netherlands.

The goal is to study social phenomena amongst truck drivers, so that we may understand why truck drivers initiate so-called *co-driving behaviour* with other drivers. In simple terms, co-driving is the activity where two trucks drive together, i.e., are frequently at the same place at the same time. We assume a direct and natural relation between a truck and its driver, meaning that a truck

driver only drives one truck. A second assumption is that every truck is always driven by the same driver. To make sure that only *intentional* co-driving activity is investigated, a number of strict selection criteria are used. These criteria are explained in Section 4.2.

Co-driving behaviour is known to have a potential positive impact on the environment through optimizing logistics and consequently reducing fuel use [15]. Hence, an improved understanding of co-driving behaviour may stimulate additional co-driving behaviour. Moreover, innovative forms of transportation, such as autonomous driving, may have major implications for this behaviour.

Without traditional spatio-temporal data mining techniques, we construct a so-called *co-driving network* from the data. The nodes of this network are trucks, while a link is drawn between two nodes when these two trucks frequently show co-driving behaviour. The obtained co-driving network shows properties similar to other networks often analyzed in the field of complex network analysis [1]. We mention three of them. First, we see that the network has a giant component containing the majority of nodes and edges. Second, the relatively low average shortest path length suggests a small-world network structure [17]. Third, the degree distribution appears to follow a power law, indicating that the network may be scale-free [1].

Previous work on similar data focused on communities and static properties of the co-driving network [3]. In contrast, the goal of this work is to learn the *dynamics* of the co-driving network. To this end, we use a link prediction approach [8]. More concretely, we develop a machine learning classifier which predicts for all possible pairs of trucks that are not connected, whether a link is formed in the future. We then investigate the importance of each type of feature that occurs in the link prediction classifier. These allow us to understand what is assessed important by the classifier and hence contributes to co-driving behaviour. The features used can be categorised into four different types of features.

1. Spatio-temporal features, aiming to summarize registrations at different locations over a given period.
2. Neighbourhood features, related to the local embedding of considered trucks in the co-driving network.
3. Path features, describing distance-related properties of truck pairs based on the global structure of the network.
4. Node features, related to static meta-information of trucks.

The overall structure of this paper takes the form of the different steps taken to attain our goal and is as follows. In Section 2, relevant work is provided on analyzing dynamics in social networks including spatio-temporal data. Section 3 describes the spatio-temporal truck data. Section 4 reports how a co-driving network is constructed from the data. In this section we also discuss the characteristics of the obtained network. Section 5 provides a formal description of the link prediction approach. It also explains how the different features are constructed from both the data and the obtained network. Section 6 outlines the

experimental setup, demonstrates the performance of the link prediction approach and assesses the feature importance. Finally, in Section 7 we conclude our paper and provide suggestions for future work.

2 Related work

From the substantial body of related work on spatio-temporal data we select three different approaches which have been frequently used to study dynamics in networks at the level of individual nodes. These three different approaches have in common that they all try to understand the underlying social network by studying node attributes available in the data.

First, Sekara *et al.* use Bluetooth sensors to measure proximity of students [14]. The authors show that when high-resolution data is available (both in time and location), groups of interacting nodes can be observed directly. Hence, making sense of individual node attributes using network measures can be performed directly. As an example, the authors show that the students explore new locations in groups during the weekend, while the groups tend to be at the same location.

Secondly, Kossinets and Watts analyze e-mail data gathered from students and employees at a university [7]. Unlike our truck data, e-mail data does not contain spatial information. However, different attributes are collected and analyzed in this work such as professional status, gender and age.

Finally, Wang *et al.* analyze the mobility patterns by tracking both the mobility and interactions of millions of mobile phone users [16]. A social network is constructed from phone calls, where users are connected when they communicate. They provide three findings. First, the authors find that spatial trajectories of two users strongly correlate when they are close in the social network. Second, finding is that mobility features have high predictive power on which nodes will connect, comparable to that of network proximity features. Third, the link prediction performance by using both network proximity and mobility features. We will use a similar link prediction approach in our work. In addition we build on the other works [7, 14] by distinguishing between weekends and weekdays, and using a combination of both network and static attributes.

3 Data

Data collection took place at 18 different locations throughout the Netherlands between 2016 and 2018. Using an automatic number-plate recognition (ANPR) system, every truck passing these locations is registered. The data is obtained by the same systems used in earlier work [3]. At some locations the registration systems faced an unexpected downtime. To ensure a sufficiently valid range of data, only registrations from 6 out of 18 systems have been considered. These systems were placed near the port of Rotterdam. Furthermore, registrations with low quality data are removed, such as invalid characters in license plates and non-existing countries.

The aforementioned quality selections reduce the number of registrations from 18,678,420 to 9,202,764. The number of registrations over time is shown in Figure 1. We observe that the number of registrations after applying the quality selections is more stable over time. In Figure 2 the distribution of the number of registrations per truck is shown (note that both axes have logarithmic scales). The distribution of the number of registrations per truck remains similar.

4 The co-driving network

We start with a description on how the co-driving network is constructed in Section 4.2. Section 4.3 continues then with general statistics of the obtained network.

4.1 Definition of an intentional co-driving event

We will now provide a more formal definition of a co-driving event. Our dataset of all registrations (as discussed in Section 3) is denoted by \mathcal{D} . We use \mathcal{D}_u to refer to all registrations x_i in dataset \mathcal{D} from truck u with license plate $lp_i = u$. More formally, $\mathcal{D}_u = \{x_i \in \mathcal{D} : lp_i = u\}$. We speak of a co-driving event (u, v, t) when two registrations $x_i \in \mathcal{D}_u$ and $x_j \in \mathcal{D}_v$ from trucks u and v exists at the same location $loc_i = loc_j$ at time t_i with at most $\Delta t = |t_i - t_j|$ seconds between them. There are two ways in which these co-driving events can occur: (1) randomly because two trucks just happen to be at the same place around the same time or (2) intentionally because two trucks were involved in *intentional co-driving*. Our goal is to study intentional co-driving behaviour while keeping the random co-driving events to a minimum.

4.2 Network construction

The following two steps are taken to ensure that only intentional co-driving is studied. First, we separate the intentional co-driving events from the random

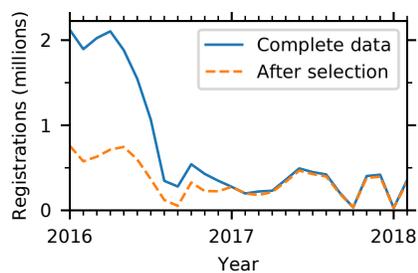


Fig. 1: Number of registrations over time.

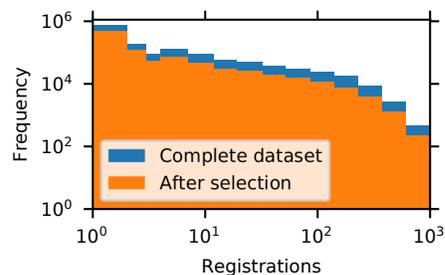


Fig. 2: Histogram of registrations per truck. Note logarithmic axes.

ones by selecting only co-driving events in which trucks u, v at least twice drive at most $\Delta t \leq \Delta t_{\max}$ seconds apart. We will discuss the Δt_{\max} parameter shortly. Second, at least two of those co-driving events should have a gap of more than two hours, i.e., there exists two co-driving events (u, v, t_i) and (u, v, t_j) for which $|t_i - t_j| \geq 2$ h. With the latter requirement, we ensure that the two co-driving events originate from different truck journeys. By applying these criteria, we minimize the probability that a random co-driving event is marked as an intentional co-drive.

The temporal network $G = (V, E)$ is constructed. In this network, the nodes are the trucks $u, v \in V$ that frequently show (intentional) co-driving behaviour. The links of this network consist of the obtained co-driving events $(u, v, t) \in E$ between those trucks. Note that multiple links (u, v, t) exist between two nodes u and v with different t as a result of the first step taken to select only intentional co-driving. We refer to the number of links between u and v as $w_{u,v}$, with $w_{u,v} \geq 2$ as a result of the selection criteria discussed above. When no links exist between u and v , the weight $w_{u,v}$ equals 0.

Then, we need to find the appropriate value for the previously discussed parameter Δt_{\max} . There is a trade-off when setting the value of Δt_{\max} . High values will result in selecting a large share of random co-driving events, while low values will result in the omission of intentional co-driving behaviour. We present three thoughts when determining the value of Δt_{\max} .

First, Figure 3 shows the distribution of the time gap between two co-drivings events. We observe that distinct behaviour is present for random and intentional co-driving events. Two trucks involved in intentional co-driving drive closer together than randomly co-driving trucks. Note that the time gap in intentional co-driving trucks peaks at around 2 s, close to the 1.3 s which is considered a safe driving gap between two trucks [9]. After $\Delta t \approx 8$ s the relative frequency of intentional co-driving trucks becomes similar to that of randomly co-driving trucks. This may indicate that from this value onward only random co-driving events are selected as intentional co-driving.

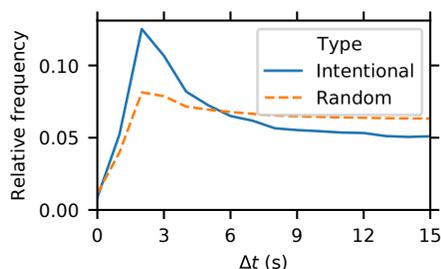


Fig. 3: Frequency distribution of time gap measured between the two trucks in a co-driving event.

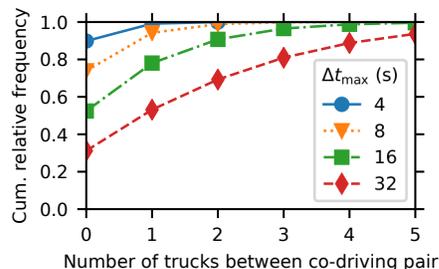


Fig. 4: Frequency distribution of trucks driving between the two trucks in a co-driving event.

Second, Figure 4 shows the distribution of the number of trucks driving between two trucks involved in intentional co-driving for various values of Δt_{\max} . For values between 4 and 8 s, we observe that virtually all trucks are driving with at most one truck in between them. Higher values result in a non-negligible probability that more than two trucks are driving between the two co-driving trucks. We consider it unlikely that trucks are intentionally co-driving when more than two trucks drive between these trucks. This is the case for values of $\Delta t_{\max} \geq 16$ s.

Third, we rationalize that following a truck intentionally is only possible when there is at most a few hundred of meters between the two trucks. Provided that trucks in our data drive typically at a speed around 20 m/s, this means that reasonable values for Δt_{\max} should be at most 20–30 seconds.

In summary, we are conjectured that the robustness checks above enable us to properly select intentional co-driving behaviour for further analysis in the remained of this paper.

4.3 Network statistics

We continue by providing general statistics of the obtained network in Table 1. For definitions of these statistics, see [2]. Note that multiple links are present between nodes. This is caused by the first measure taken to select intentional co-driving events in Section 4.2. In Figure 5a the distribution of the number of neighbours of each node is shown. We show in Figure 5b the distribution of node strengths. For a node, this value is equal to the sum of the weights of the nodes connected to it. Note that both distributions appear to have power-law behaviour. Together with the presence of a giant component and a relatively low average shortest path length the power-law behaviour suggests that the network is remarkably similar to other social networks and scale-free networks commonly observed in real-world settings [2, 17].

Table 1: General network statistics

Property	Value
Number of nodes	25,553
Number of links	73,059
Number of connected node pairs	27,986
Fraction nodes in giant component	62%
Fraction links in giant component	79%
Density	2.2×10^{-4}
Power law exponent γ	3.3
Average shortest path length	7.8
Diameter	24

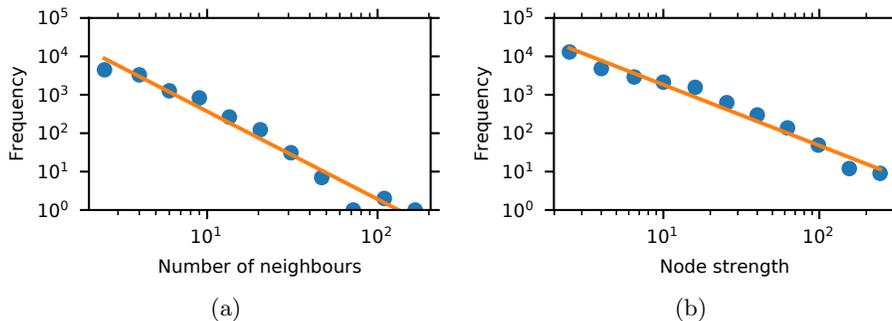


Fig. 5: The number of neighbours (a) and node strength (b) distribution of the network. Note the logarithmic axes.

5 Approach

This section presents our approach to the analysis of the dynamics of the co-driving network. We start with a description of the proposed link prediction approach in Section 5.1. The feature construction steps are outlined in Section 5.2. Finally, we provide the measures taken to reduce the observed class imbalance in Section 5.3.

5.1 Link prediction

We start with a formal description of the link prediction problem. Given a snapshot of the network at time τ , the link prediction classifier needs to predict newly formed links in the evolved network after time τ . In doing so, the classifier is able to use present information to predict future links. The input of this classifier is a feature matrix X , which is based on a snapshot of the network at time τ : $G_\tau = (V_\tau, E_\tau)$ with $E_\tau = \{(u, v, t) \in E : t \leq \tau\}$ and $V_\tau = V$. The feature vector is calculated for each candidate node pair which is not linked (yet) in G_τ : $\{(u, v) \in V_\tau \times V_\tau : u \neq v, (u, v, t) \notin E_\tau\}$. The target of the classifier, y , denotes for a node pair whether a link is present in the evolved network:

$$y_{u,v} = \begin{cases} 0 & \text{if } (u, v, t) \notin E \\ 1 & \text{if } (u, v, t) \in E \end{cases} \quad \text{for some } t > \tau$$

Note that only link formation is predicted; we do not predict the weight of the link. Accordingly, this classifier can be seen as a supervised binary classifier.

A random forest classifier is used. We choose this classifier because random forests are known to generalize well on unseen data and allows one to determine the importance of each feature [4, 6]. The random forest classifier contains 128 decision trees, since larger values usually bring no significant performance gain [10]. Each individual decision tree is trained on a randomly drawn selection of variables. The number of randomly drawn features is equal to the square root

of the total number of variables, which is a common value used in classification [12]. By sampling randomly with replacement from the data, randomness is increased for the individual decision trees. The splitting criteria of the nodes are determined by considering the reduction in Gini impurity [6]. The random forest classifier allows to obtain the importance of each feature by determining the reduction in the Gini impurity for splitting nodes with a certain variable [6]. Recall that this is important, as it enables us to understand the dynamics of truck co-driving behaviour.

Because only a sample of the data is used for each tree, we can calculate an estimate of the performance for the decision forest on the remaining part of the data [6, 12]. We use this estimate to assess the optimal value for the depth of the decision trees in the random forest. The performance of the classifier is calculated on the test set, which is a 10% random sample of the data. This data is not used in the training of the classifier nor in finding the optimal value for the depth of the decision trees.

5.2 Feature construction

We continue with the explanation of the feature vector which is used for each candidate truck pair (a, b) . Table 3 provides all 52 features used by the link prediction classifier. These features can be categorised into four different types, each described in more detail below.

1. Node features, constructed from information available about the trucks.
2. Spatio-temporal features.
3. Neighbourhood features, which consider relevant operations related to micro-level properties of the nodes of the candidate pair. The neighbourhood of a node is defined by $\Gamma(a) = \{v \in V : (a, v, t) \in E \text{ for some } t\}$. The strength of a node is the total number of links connected to a node, $s_a = \sum_{\{u \in V\}} w_{a,u}$.
4. Path features, that consider macro-level properties of the network. We consider only the shortest path length in this work.

Truck properties The ANPR-system determines the license plate and country (*country_u*) of each truck u passing by. We use \mathcal{D}_u to denote all registrations x_i available of truck u , as explained in Section 4.2. The registration systems are also equipped with sensors to measure the length (*length_i*), mass (*mass_i*) and number of vehicle axes (*axes_i*) of each truck. These measurements may slightly differ between registrations. Therefore, we calculate the averages shown in Table 2 for each truck in the network. The *driving_hours* and *weekend_driver* features are calculated because they are known to vary between trucks operating in different industrial sectors.

Spatio-temporal information The goal of this type of feature is to capture both spatial and temporal behaviour for the truck pair under consideration. We do so by counting the number of registrations in different time periods. As an

Table 2: Information available about truck u , collected from its registrations \mathcal{D}_u .

Property		Description	Type
$truck_country_u$		country of registration	string
$truck_axes_u$	Median $axes_i$ $\{x_i \in \mathcal{D}_u\}$	number of axes	number
$truck_length_u$	Median $length_i$ $\{x_i \in \mathcal{D}_u\}$	length	number
$truck_mass_u$	Median $mass_i$ $\{x_i \in \mathcal{D}_u\}$	mass	number
$driving_hours_u$	Mean $ t_i(\text{h}) - 12\text{h} $ $\{x_i \in \mathcal{D}_u\}$	usual driving hours	number (0–12)
$weekend_driver_u$	Mean $\begin{cases} 0 & \text{if } t_i = \text{weekday} \\ 1 & \text{if } t_i = \text{weekend} \end{cases}$ $\{x_i \in \mathcal{D}_u\}$	fraction driving in weekend	number (0–1)

example, for feature $last_day_l$ ($a + b$) registrations are counted for trucks a and b at location l in the last day before the considered time. We consider time periods of one week, one month and one year. These periods are chosen in such a way that they cover a broad window of possible relevant time periods.

5.3 Class imbalance

It is well-known that in real-world networks the link prediction classifier comes with a large class imbalance [16]. The performance of the employed random forest classifier may drop if there is a large class imbalance. To overcome this limitation, we use the following three measures. First, both classes are given a weight, such that in total both classes have equal weight. Second, we consider only truck pairs where both trucks are involved in co-driving events in the last two months before time τ . This will both reduce the number of considered truck pairs and reduce class imbalance. Third, we consider only node pairs with both trucks in the giant component of the network. An additional advantage of this measure is that the shortest path feature is well-defined.

6 Experiments and results

To determine important factors that govern the dynamics in the truck co-driving network, the approach as set out in Section 5 is applied to the network discussed in Section 4.

The value of τ is chosen such that 50% of the links are formed. In this way, the number of considered truck pairs and class imbalance are reduced, while ensuring that at least 1,000 truck pairs are present that will make a link. For this value of τ we find a class imbalance of 1:61,000. By taking the measures mentioned in Section 5.3, the class imbalance is reduced to 1:15,000. The random forest is used as implemented in Python sci-kit learn 0.21.2. We find an optimal value of three for the maximum depth of the decision trees in the random forest.

We report the trade-off between true positives and false positives to assess the accuracy of the classifier. The relation between these two values can be

Table 3: Features for truck pair (a, b) used in the link prediction model. The rightmost column lists the feature importance calculated using Gini importance as provided by the random forest classifier.

Index	Feature	Type	Importance
X_1	$truck_country(a) = truck_country(b)$	node	0.005
X_2	$truck_ax(a) + truck_ax(b)$	node	0.006
X_3	$ truck_ax(a) - truck_ax(b) $	node	0.008
X_4	$truck_len(a) + truck_len(b)$	node	0.017
X_5	$ truck_len(a) - truck_len(b) $	node	0.040
X_6	$truck_mass(a) + truck_mass(b)$	node	0.016
X_7	$ truck_mass(a) - truck_mass(b) $	node	0.030
X_8	$driving_hours(a) + driving_hours(b)$	node	0.016
X_9	$ driving_hours(a) - driving_hours(b) $	node	0.030
X_{10}	$weekend_driver(a) + weekend_driver(b)$	node	0.014
X_{11}	$ weekend_driver(a) - weekend_driver(b) $	node	0.019
$X_{12}-X_{19}$	$last_week_l(a+b)$ for $l = 1, \dots, 8$	spatio-temporal	0-0.027
$X_{20}-X_{27}$	$last_month_l(a+b)$ for $l = 1, \dots, 8$	spatio-temporal	0-0.057
$X_{28}-X_{45}$	$last_year_l(a+b)$ for $l = 1, \dots, 8$	spatio-temporal	0.010-0.060
X_{46}	$ \Gamma(a) + \Gamma(b) $	neighbourhood	0.117
X_{47}	$ \Gamma(a) - \Gamma(b) $	neighbourhood	0.013
X_{48}	$ \Gamma(a) \cup \Gamma(b) $	neighbourhood	0.093
X_{49}	$ \Gamma(a) \cap \Gamma(b) $	neighbourhood	0.021
X_{50}	$s_a + s_b$	neighbourhood	0.056
X_{51}	$ s_a - s_b $	neighbourhood	0.017
X_{52}	shortest path length in G	path	0.111

shown using the well-known Receiver Operator Characteristic (ROC)-curve [10], shown in Figure 6. The area under the ROC curve (AUROC) is 0.84. This value indicates that the classifier is able to accurately select the links that will appear. This allows us to analyze the feature importance.

The feature importance of each individual feature is provided in Table 3. In Figure 7 the features are shown for each of the aforementioned four types. We observe that the neighbourhood features score highest, closely followed by the single path feature. The two neighbourhood features with the highest scores are X_{46} and X_{48} . These features provide the sum of the degrees of the node pairs and of the union of their neighbourhoods, respectively. Both the spatiotemporal and node features score lower.

Since the features based on network metrics have a higher feature importance, from these experiments we conclude that the network perspective on this data is fruitful. A next step is to interpret these findings in the infrastructure domain; we leave this step for future work.

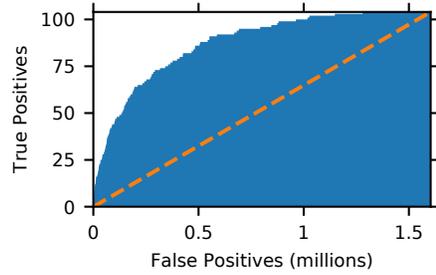


Fig. 6: Receiver Operator Characteristic curve of the random forest link prediction classifier. The AUROC is 0.84.

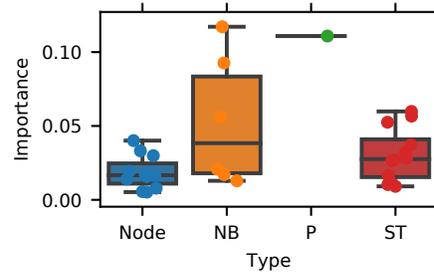


Fig. 7: Gini feature importance. NB, P and ST are the neighbourhood, path and spatiotemporal features resp.

7 Conclusions and outlook

In this work we assessed to what extent the dynamics of co-driving behaviour are stimulated by the characteristics of the nodes involved. The research has shown that features based on network measures are able to explain the dynamics of the studied co-driving network. This means that the network perspective on this spatio-temporal dataset of truck driving in the Netherlands is meaningful and worth exploring further. Our findings may also suggest that, in general, the link prediction approach is suitable to analyze spatiotemporal datasets containing social behaviour.

In future work, we will extend the devised link prediction approach to other datasets in the infrastructure domain. This will allow us to investigate how different type of features perform in different contexts. A second interesting angle is to use a similar approach to predict which nodes will turn inactive, i.e., will not form any new links. This will result in a substantially smaller set of candidate nodes for the link prediction algorithm. Finally, future work will focus on interpreting and applying the knowledge observed from the micro-level dynamics in the infrastructure domain.

References

1. Barabási, A.L.: Network Science. Cambridge University Press (2016)
2. Barabási, A.L., Albert, R.: Emergence of scaling in random networks. *Science* 286(5439), 509–512 (1999)
3. de Bruin, G.J., Veenman, C.J., van den Herik, H.J., Takes, F.W.: Understanding behavioral patterns in truck co-driving networks. In: *International Conference on Complex Networks and their Applications*. pp. 223–235. Springer (2018)
4. Criminisi, A., Shotton, J., Konukoglu, E.: Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning. *Foundations and Trends in Computer Graphics and Vision* 7(2–3), 81–227 (2012)

5. Cuttone, A., Lehmann, S., González, M.C.: Understanding predictability and exploration in human mobility. *EPJ Data Science* 7(1), 2 (2018)
6. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning*. Springer Series in Statistics, Springer-Verlag (2009)
7. Kossinets, G., Watts, D.J.: Empirical analysis of an evolving social network. *Science* 311(5757), 88–90 (2006)
8. Liben-Nowell, D., Kleinberg, J.: The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology* 58(7), 1019–1031 (2007)
9. Mazurek, U., van Hattem, J.: Rewards for safe driving behavior: Influence on following distance and speed. *Transportation Research Record* 1980(1), 31–38 (2006)
10. Oshiro, T.M., Perez, P.S., Baranauskas, J.A.: How many trees in a random forest? In: *International Workshop on Machine Learning and Data Mining in Pattern Recognition*. pp. 154–168. Springer (2012)
11. Pappalardo, L., Rinzivillo, S., Qu, Z., Pedreschi, D., Giannotti, F.: Understanding the patterns of car travel. *The European Physical Journal Special Topics* 215(1), 61–73 (2013)
12. Probst, P., Wright, M.N., Boulesteix, A.L.: Hyperparameters and tuning strategies for random forest. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 9(3), e1301 (2019)
13. Sapiezynski, P., Stopczynski, A., Gatej, R., Lehmann, S.: Tracking human mobility using wifi signals. *PloS one* 10(7), e0130824 (2015)
14. Sekara, V., Stopczynski, A., Lehmann, S.: Fundamental structures of dynamic social networks. *Proceedings of the National Academy of Sciences* 113(36), 9977–9982 (2016)
15. Tsugawa, S., Kato, S.: Energy ITS: Another application of vehicular communications. *IEEE Communications Magazine* 48(11) (2010)
16. Wang, D., Pedreschi, D., Song, C., Giannotti, F., Barabási, A.L.: Human mobility, social ties, and link prediction. In: *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1100–1108. ACM (2011)
17. Watts, D.J., Strogatz, S.H.: Collective dynamics of small-world networks. *Nature* 393(6684), 440 (1998)