

ANALYSIS OF SPATIO-TEMPORAL TRAFFIC PATTERNS BASED ON PEDESTRIAN TRAJECTORIES

S. Busch^a, T. Schindler^a, T. Klinger^b, C. Brenner^a

^a Institute of Cartography and Geoinformatics, Leibniz Universität Hannover, Germany -
(busch, brenner)@ikg.uni-hannover.de, t.schindlermail@gmail.com

^b Institute of Photogrammetry and GeoInformation, Leibniz Universität Hannover, Germany -
klinger@ipi.uni-hannover.de

Commission II, WG II/8

KEY WORDS: Periodic Event Analysis, Walking Path Network, Dynamic Prior Map, Pedestrian Behaviour Prediction, Traffic Pattern

ABSTRACT:

For driver assistance and autonomous driving systems, it is essential to predict the behaviour of other traffic participants. Usually, standard filter approaches are used to this end, however, in many cases, these are not sufficient. For example, pedestrians are able to change their speed or direction instantly. Also, there may be not enough observation data to determine the state of an object reliably, e.g. in case of occlusions. In those cases, it is very useful if a prior model exists, which suggests certain outcomes. For example, it is useful to know that pedestrians are usually crossing the road at a certain location and at certain times. This information can then be stored in a map which then can be used as a prior in scene analysis, or in practical terms to reduce the speed of a vehicle in advance in order to minimize critical situations. In this paper, we present an approach to derive such a spatio-temporal map automatically from the observed behaviour of traffic participants in everyday traffic situations. In our experiments, we use one stationary camera to observe a complex junction, where cars, public transportation and pedestrians interact. We concentrate on the pedestrians trajectories to map traffic patterns. In the first step, we extract trajectory segments from the video data. These segments are then clustered in order to derive a spatial model of the scene, in terms of a spatially embedded graph. In the second step, we analyse the temporal patterns of pedestrian movement on this graph. We are able to derive traffic light sequences as well as the timetables of nearby public transportation. To evaluate our approach, we used a 4 hour video sequence. We show that we are able to derive traffic light sequences as well as time tables of nearby public transportation

1. INTRODUCTION

1.1 General Problem Description

The general issue with human locomotion is that we move faster than we can react. Traffic speed makes it impossible for vehicle drivers to break in time if an unexpected cross traffic appears all of a sudden. The speed limit in which emergency break assistants operate, shows how difficult a reliable situation estimation based on instant observation is. Equally, it is an impossible task for a human driver. Thus, a certain speed makes a priori knowledge necessary. Any human driver adapts his behaviour to different traffic situations, which sometimes do not obey the law, wherefore these situations are particularly difficult to interpret by drivers unfamiliar with the area or by algorithms. Unfortunately, those traffic participants, where the most wrongdoings occur, are also the group with the most dynamic behaviour, namely the pedestrians. They can quickly change the direction of their movement and are very vulnerable traffic participants, which deserve special care. We will present an approach to generate knowledge about pedestrians and to record it in a dynamic map. In our opinion such a map can lower the risk for all traffic participants, raise the speed and comfort of autonomous driving as well as optimise the path planning of navigation systems. The daily traffic of high-tech vehicles equipped with range sensors are a usable data source to generate comprehensive and high frequency data for an a priori map. The highly dynamic behaviour of pedestrians and the wide spread use of cameras in vehicles leads us to the decisions to gain information about pedestrians from images, even if it does not matter where the trajectories are coming from. For the use of a priori knowledge from a map. The

data must be up to data. We followed approaches from road network reconstruction to generate a walking path network. For that purpose we used pedestrian tracking (Klinger et al., 2015) and approaches to generate regions of interest (ROI) (Feuerhake et al., 2011; Shen and Cheng, 2016). Afterwards we analysed nodes and edges for periodic events (Fayazi et al., 2015). Although, there are current effort of analysing events in time and space, like earthquakes with twitter (Sakaki et al., 2010; Weng and Lee, 2011), we still use simple thresholds to separate time point with the count of people to identify periodic events, because of the small amount of data. We decided to use an intuitive heat map extended with schedule tables to present our map instead of spatio-temporal cubes visualisation for time variant events (Kraak, 2003; Nakaya and Yano, 2010; Gatalsky et al., 2004), whose clarity does not scale with the data amount. In Section 2 we give an overview of related work and concepts which we can use with future work if the amount of data rises. This conception is followed by the presentation of our approach in Section 3. We introduce our data set and specify our actions in Section 4 and thereafter we evaluate our results in Section 5. Finally we summarize our results and end with an outlook in Section 6.

2. RELATED WORK

2.1 Graph Derivation

We collected the approaches to derive annotated graphs as navigation map from trajectories. Traditionally, change detection requires manual interpretation of information from field data such

as close-range or airborne imagery. Recent work aims at the recognition of patterns from trajectory data for the automatic derivation of road networks from GPS trajectories. These can be organized, for instance, into three categories (Ahmed et al., 2015): point clustering (Edelkamp and Schrödl, 2003), incremental track insertion (Brüntrup et al., 2005) and intersection linking (Karagiorgou and Pfoser, 2012). Moreover, a fourth strategy can be distinguished; the generative modelling (e.g. Chai et al., 2013; Kuntzsch et al., 2015). The approach presented in this paper is inspired by the point clustering approach.

Alternatively to GPS, trajectories can be derived from image sequences of surveillance cameras using visual object tracking algorithms. In online applications, detection-based approaches for tracking are widely used to generate trajectories (Ess et al., 2008; Schindler et al., 2010; Klinger et al., 2015). Such approaches apply object detection in the individual frames and by solving a data association problem. The detections are used to update a recursive filter that represents the dynamic system state of the tracked object. The algorithms make use of probabilistic graphical models to incorporate different sources of information, e.g. pedestrian detections, motion and depth, to reason about a pedestrian's position. Since the graphic modelling inherently provides a way to incorporate further information such as a path network, we stick to the work of Klinger et al. (2015) and use dynamic Bayesian networks to generate the trajectories. There are some works to derive pedestrian path networks from trajectories. One option to generate nodes of walking path graphs is working with junction and entrance points (Makris and Ellis, 2005). Another option is the use of Gaussian processes to generate movement models with a mean predicted instantaneous velocity grid (Ellis et al., 2009) and derive nodes at the grid point with the lowest uncertainty. Both approaches, Makris and Ellis (2005) and Ellis et al. (2009) work with a static camera. Makris and Ellis (2002) also presented an approach to detect typical behaviour by pedestrian trajectories which works without derivation of nodes. Feuerhake et al. (2011) and Shen and Cheng (2016) calculate regions of interest to analyse pedestrian behaviour, while Feuerhake et al. (2011) form a walking path graph by connecting the ROIs with trajectories.

2.2 Event Identification

There are several approaches to find time points of interest in a time series. Some works detect events in time with a priori knowledge or using a model (Guralnik and Srivastava, 1999; Sakaki et al., 2010), others use time varying Poisson processes (Ihler et al., 2006) or clusters of wavelet-based signals (Weng and Lee, 2011). Each of them find events in time with a procedure which derived satisfactory responses for their data. Furthermore, there are some approaches, which analyse the traffic light sequences and cycles based on GPS trajectories (Fayazi et al., 2015; Kerper et al., 2012). Both approaches determine the period with the smallest error for observed starting events by minimization. While Kerper et al. (2012) minimize the mean square error for each starting point by searching between the minimum and maximum gap of their data, Fayazi et al. (2015) minimize the variance of differences in order to determine the period before they calculate the sequence's starting point via a cyclic mean of different starting points.

2.3 Visualisation

For the presentation of time and space variant trajectories, the use of a space-time cube is introduced as a useful tool for manual exploration of spatio-temporal events (Kraak, 2003; Nakaya and Yano, 2010; Gatalsky et al., 2004). The clarity of this visualization tool decreases with the amount of simultaneously shown

events, as well as the bar chart, which Nakaya and Yano (2010) use during their exploration.

3. OUR APPROACH

We enhance a map of a junction with information about pedestrian behaviour, for which we use the nodes of a walking path graph as starting points to determine the frequency of pedestrian traffic along the edges of our graph. These periodic events are mapped in schedules at nodes and as mean periods at edges. Our approach is split up into two parts. First, we generate a walking path network from pedestrian trajectories tracked by a static camera. Even though the approaches of Makris and Ellis (2005) and Ellis et al. (2009) work with static cameras, we do not pursue either way because our work is aiming to gather information using dynamic cameras from vehicles in daily traffic. Hence, we do not work with entrance points in a static scene but used ROIs (Feuerhake et al., 2011; Shen and Cheng, 2016) as nodes for a walking path graph. Secondly, we analyse the nodes and edges of that graph for periodic events to gain some prior information about the movements and intentions of pedestrians. Therefore, we separate event points with a static threshold and classify these events as either starting points of green phases (Fayazi et al., 2015) or subway arrivals. Finally, we map the frequencies of such periodic events and present a visualisation approach. We use a heat map to clearly show increasing pedestrian traffic at nodes over time. On the other hand, time tables of events are added. For the prediction of periodic events the combination of both is exemplary shown in figure 1.




Figure 1: Snapshot of dynamic map visualisation

3.1 Graph of Walking Path

We use a probabilistic multi-person tracker (Klinger et al., 2015) to generate pedestrian trajectories following the assumption that pedestrians walk slower or even stop at ROIs (Feuerhake et al., 2011; Shen and Cheng, 2016). First we derive the walking speed for each sample point of trajectories and subsequently cluster them with a spatio-temporal density based clustering (ST-DBSCAN) following Shen and Cheng (2016) and Birant and Kut (2007). For the ST-DBSCAN we only use points with a low speed according to determined ROIs. Therefore, a ROI is identified by a set of points with a specified density in time and space. These derived ROIs are assigned to a position in space and scope in time by the

underlying cluster points. These time variant ROIs present the nodes of our walking path network and defined edges between this graph following Feuerhake et al. (2011). For the determination of edges we take a look at the trajectories again. If a trajectory passes through several ROIs, these ROIs become connected with a directed edge or if already present in the graph, the edge weight is incremented. Passing through a ROI means getting into a specified radius and scope of a ROI, which are defined by the underlying cluster. The weight of the edges will be used in future work for the prediction of pedestrian movements. We use the scope of ROIs to adapt the graph structure over time. For further event analysis we focus on a snapshot of the spacial and time variant graph, as shown in figure 6.

3.2 Event Analysis

As an example for useful additional map information, derived from pedestrian trajectories, we determine traffic light sequences and subway schedules from pedestrian movements. This periodic events can be used for prediction of future traffic behaviour.

3.2.1 Traffic Light Sequence (TLS): The determination of traffic light sequences from vehicle trajectories barley varies from the pedestrian trajectories. However, pedestrians do not queue and their acceleration phase is negligible. Thus, assumption 1 applies.

Assumption 1: A node flushes instantly at the start of a green phase.

We successively check every observed time point t_i for being a potential starting point of a green phase, which is the case if the current count of pedestrians is zero and the previous one was greater than zero.

$$t_{SG} = \{t_i | c_{i-1} > 0 \ \& \ c_i = 0\} \quad (1)$$

where c_i is the count of pedestrians at the time i at a ROI. We minimize the variance of all $t_{SG} \in TSG$ over discrete periods between 1 and 120 in 1 second interval steps. Therefore, we calculate the mean traffic light sequence following Fayazi et al. (2015).

$$mod_P(\delta_j) = \delta_j - round(\delta_j/P)P \quad (2)$$

where $\delta_j = t_{SG_j} - t_{SG_{j-1}}$. $mod_P(\delta_j)$ represents the difference from δ to the assumed period P , it holds: $-(P/2) < mod_P(.) < P/2$. Then the mean traffic light sequence is determined by the period with the smallest sum of deviations solving:

$$P_{min} = \arg \min_P \sum_{j=1}^n \left(\frac{mod_P(\delta_j)}{P/2} \right)^2 \quad (3)$$

where the division of $P/2$ normalize the variances from different periods for a relative comparability. A determination of the starting point for the TLS is not useful in our case because of the variation of data, as we will show in section 4.1.

3.2.2 Subway Schedule: The identification of subway schedules could be solved with an approach identical to TLS determination if there were only one route. Commonly there is more than one route at a platform or arrival times of the opposed directions differ. These different periods in the signal exclude the presented

TLS approaches for solving the problem of period determination. Even though Fayazi et al. (2015) detect the change of signal offset by Gaussian mixture models, their approach is not useful for determining several periods. For the detection of subway arrivals we presume assumption 2.

Assumption 2: At subway arrival the ROI will fill over-average.

We successively check every observable time point if it is a possible subway arrival point. The time point which have a count over average while the previous count is below-average form the set SA of possible subway arrivals.

$$SA = \{t_i | c_{i-1} > c_{avg} \ \& \ c_i < c_{avg}\} \quad (4)$$

where c_i again is the count of pedestrians at a ROI. We use linear regression models $sa'_x = p*x + s$ to find several frequencies with fixed periods. The problem is solved through a multi periodic event sample consensus (MPE-SAC) algorithm. This algorithm looks for possible periods (P) and in addition starting points (S) to determine the best fitting periods, analogue to the TLS algorithm previously introduced. Therefore, we rate the regression models SA' in a convoluted manner. Each subway arrival is assigned to a period if the difference to regression model is less than the threshold θ . The problem of ambiguous sub-sequences is solved by using a quadratic score function for the matches and the punishment of theoretically observable points with no observation by 0.5, see equation 5. The quadratic score function weight the matches between 1 and 0 depending on their distance. Thus, we avoid running into a wrong maximum at compound periods. Hence, we derive the following score function and use negative value as termination condition for the MPE-SAC regression:

$$Score = \sum_{i=1}^{i=n} \max_{P,S} \frac{\theta^2 - mod_P(MSA_{p_i,s_i})^2}{\theta^2} - \frac{obe_{p_i,s_i} - |MSA_{p_i,s_i}|}{2} \quad (5)$$

with

$$obe_{p_i,s_i} = \frac{\max(SA_{p_i,s_i}) - \min(SA_{p_i,s_i})}{p} \quad (6)$$

where n is the number of periods being searched and SA_{p_i,s_i} is the subset of subway arrival time points which do not match for the best matching period p with starting point s calculated for frequency $i-1$.

$$MSA_{p,s} = \{t \in SA_{p,s} | mod_p(t-s) < \theta\} \quad (7)$$

where the number of frequency i describes the level of the subset, it is:

$$SA_{p_{i+1},s_{i+1}} = SA_{p_i,s_i} \setminus MSA_{p_i,s_i} \quad (8)$$

where $SA_{p_0,s_0} = SA$ the set of all subway arrivals.

4. EXPERIMENT

4.1 Database

We used a static camera to track pedestrians at a T-junction where cars, public transportation and pedestrians interacted over a time

period of 4 hours. We analysed the video manually for ground truth and automatically to compare the results in section 5. During



Figure 2: Snapshot of trajectories within the observed scene

the manual analysis we tracked both the times a subway train stopped and the time at which traffic lights switched. In addition we counted the number of people exiting a subway train and waiting at a traffic light aggregated over a 10 second interval. For the automatic analysis a probabilistic multi-person tracker using dynamic Bayes networks from Klinger et al. (2015) supplied trajectories exemplarily shown in figure 2. We obtained 2766 trajectories with a mean length of about 14 meters consisting of 352225 data points, which corresponds to a mean density of 10 cm between trajectory points. In the captured scene the main road




Figure 3: Schematic schedule of two subway lines L1 and L2

separates the subway platform. Pedestrians as well as vehicles on the side road have the possibility to trigger the traffic light. The length of the traffic light's green phases for the side road varies between 8 and 20 seconds depending on whether a pedestrian, car or both trigger the traffic light. The red phase of the side road lasts at least 28 seconds because of the different triggering possibilities. Furthermore, two subway routes stop on the platform, according to the timetable the opposite directions meet at the platform. Both routes have the same period. Thus, by schedule a 4 minute and a 6 minute gap arises, see figure 3. Unfortunately, the

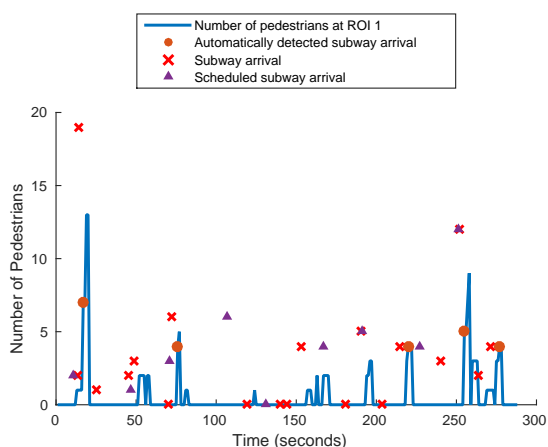


Figure 4: Sub-section of manually collected aggregated number of pedestrians for ROI #1 with automatic and manually detected subway arrivals and schedule.

opposite directions do not always meet due to delay, see figure 4, which shows the number of pedestrian at ROI #1 (see figure 1) with the subway arrivals. It becomes obvious that people exiting an arriving subway need a few seconds to come up to ROI #1 and trigger the automatic detection. Furthermore, the subway arrivals indicate a common delay in accordance with the schedule. Additionally people at the platform cannot be tracked until they reach the junction so that different walking speeds and delay from schedule generate deviations. Furthermore, assumption 2 fails to detect the subway arrivals with only few people exiting, as well as too many people who chose another way or the crowd arriving successively. The problem of a second exit could be weakened by a comprehensive walking path network. The problem of too few people exiting a subway could not be solved with the individual pedestrian detection and the clustering of successive arrivals will shift the problem to false detection. Figure 5 shows the tracked




Figure 5: Sub-section of automatically detected number of pedestrians for ROI #1 with automatically and manually detected subway arrivals.

persons at ROI#1 with subway information. Consequently, we have additional corruption mismatches, particularly bad at large crowds. We show that the use of the mean allows a reliable schedule detection, despite noise and corruption. Assumption 1 is not corrupted by the number of waiting pedestrians.

4.2 Graph Determination

The graph determination is inspired by point clustering, because our pedestrian trajectories have a higher resolution in space and time as well as shorter segments in contrast to GPS tracks, shown in figure 2. For the identification of slow walking or waiting peo-




Figure 6: Snapshot of the dynamic graph derived from pedestrian trajectories

ple we use a threshold of $0 \frac{m}{s}$ for walking speed. We perform

ST-DBSCAN clustering following Shen and Cheng (2016) with a threshold of 50 data points which have to be closer together than one meter in space and lay within a one hour interval. We present the result in part as a snapshot, shown in figure 6. For our data set the ROI number 0 and 1 are significant because they are so frequently visited that they stay over the 4 hour period. Our further analysis focus on these two regions ROI#0 and ROI#1.

4.3 Periodic Time Events

With the assumption 1 and 2 automatically detected subway arrivals and starting point of green phases exemplary shown in figure 7 and figure 5 were identified. With the algorithm to deter-

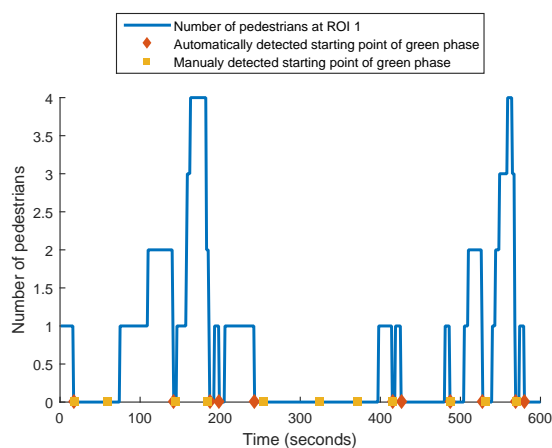


Figure 7: Sub-section of automatically tracked pedestrians at ROI#1 and ROI#2 with starting points of green phase considering one second reaction time.

mine the TLS declared in section 3.2.1 a mean traffic light sequence of 44 seconds with a minimum sum of variance was calculated over a four hour period. Figure 8 shows the histogram of minimal variance for the automatically detected 44 second TLS. The high mean, in contrast to an explicit peak, maps the variation of sequences due to the switching capacity of the traffic lights. The trend over all analysed periods shown in figure 10 points to the best fitting mean TLS of 44 seconds.

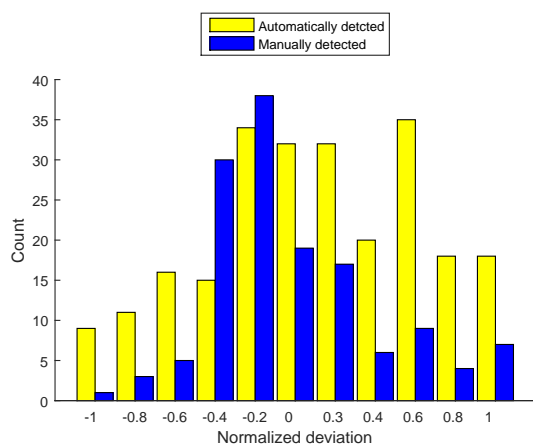


Figure 8: Histogram of normalized deviation automatically manually determined mean TLS of 44/49 seconds follow Fayazi et al. (2015).

Despite the corrupted data the MPE-SAC was clearly able to determine the schedule of the subways with a threshold of 120 seconds as tolerance for delays to ensure detection of subway arrivals as shown in figure 9. Figure 9 shows the score for the best peri-




Figure 9: Score of MPE-SAC over first periods of best (periods, starting points) tuple.

ods and step pairs. These pairs are identified by the first value for the periods on the x-axis and the composed score, see equation 5, on the y-axis. The highest peak at 600 seconds shows the clarity of our score function. The other peaks at multiples of the period point out the need of an adapted score function. We analysed periods from 1 to 3000 seconds and starting points between 1 and 3000 seconds. The upper borders for the periods are caused by the amount of data, which also applies for useful starting points. MPE-SAC found two frequencies, both with a period of 600 seconds and with starting points at 0 seconds and 240 seconds. The period and the gap between the starting points matches the schedule of the subway as shown in the next section 5.

5. EVALUATION

We evaluated the presented algorithm by comparing the output to ground truth. Analogous to the sections before we split the evaluation into two parts. First, we take a closer look at the generated graph network and secondly we evaluate the automatically derived parameter for periodic events with the given circumstances.

5.1 Graph Generation

Pedestrians do not follow such strict conditions regarding their movement as vehicles. Also there are no walking path network graphs which map real pedestrian behaviour and could be used as ground truth. We found ROIs next to traffic lights and the determined edges follow the pavement and the traffic light crossings. Thus, even if the edges of the graph do not perfectly match the edges of a map like OSM, (compare figure 1) we assume the derived graph to be valid.

5.2 Periodic Events

During the evaluation of periodic events we compare the algorithm result of the automatically detected events to the ground truth.

5.2.1 Evaluation of TLS The match and the false detection of starting points of green phases are exemplarily shown in figure 7. The automatically detected mean TLS over the 4 hour period

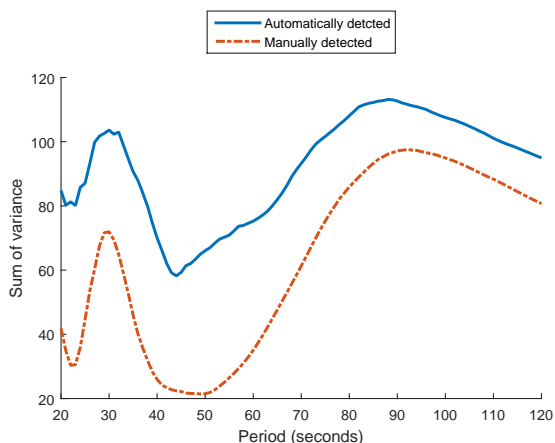


Figure 10: Trend of the minimal sum of variance for TLS.

is compared to the ground truth in figure 10. The minimum of automatic detection is 44 seconds whereas the ground truth is 49 seconds, the automatic false detections are the reason for the higher sum of variance. The more evenly spread of variance and the 5 seconds shorter TLS is explainable through the pedestrian behaviour. As shown in figure 7 sometimes pedestrians crossed the road at a red light, consequently the automatically detected interval of 30 to 59 seconds determined by the mean derivation and TSL has is lower as the ground truth interval of 40 to 59 seconds. Additionally, we analyse the mean deviation of TLS for different time windows instead of one mean for the entire 4 hour period. With a time window of 50 minutes we were able to reduce the mean difference compare to the ground truth to two seconds shown in figure 11. This smaller variance and the perfect match to ground truth for single time windows show that a more precise estimation for TLS with defined time windows by more data is possible.




Figure 11: Mean difference between automatically and manually detected mean TLS for different time windows for the automatic detection

5.2.2 Evaluation of Subway Schedule We evaluate the automatically derived subway schedule by comparing the differences between subway arrival as well as detected events to the calculated and the original schedule, see figures 12 and 13. Both figures show subway delays to automatically detected and original schedule, each subway arrival is assigned to a absolute delay bin

with 10 seconds width. The smaller maximum difference of the detected schedule to the detected events, shown in figure 12, as well as to the manually detected subway arrival, shown in figure 13 indicate a higher reliability of the detected subway schedule in contrast to the actual schedule. Moreover, the automatically detected schedule shows less delay to the ground truth subway arrivals in contrast to the detected ones, this shows the robustness of the MPE-SAC against false detection. Furthermore, the mean squared error of the automatically detected schedule is only half the one of the predefined schedule. This is reflected by standard deviations of 60 seconds compared to 90 seconds and shows the greater consistence of the automatically generated schedule.




Figure 12: Differences of schedules to the automatically detected subway arrival events




Figure 13: Differences of schedules to the manually detected subway arrival events

We observe a shift of our schedule of about two minutes to the given one. A certain delay is inevitable because of the observation of pedestrians only starting at the exit. We manually determined a mean offset of about 30 seconds due to the time necessary to exit the subway. The other offset seems to be a condition of the subway average delay, which we determined to 90 seconds.

6. SUMMARY AND OUTLOOK

In this paper, we showed that the ST-DBSCAN can be used to determine ROIs of pedestrians. In addition, we used these regions to generate a time variant graph of walking paths. Furthermore, we analysed the stable subset of this graph for periodic

events for which we used an algorithm to derive traffic light sequences to analyse the edges. The nodes were analysed for periodic events with a MPE-SAC algorithm to determine periods and starting points. We showed that the automatically derived schedule is a better match to the current traffic situation than the given schedule. Thus, we were able to enhance the graph with this information to build a time variant prediction map. With the gained information we were able to enhance a map shown in figure 1 with additional information for pedestrian behaviour and information about frequent behaviours, which manifested e.g. subway arrivals or traffic light phases. This information can be used for more efficient path planning and to warn drivers to focus their attention in dangerous situations. Furthermore, automatically derived path networks and schedules from videos have the potential to correct given maps and produce more realistic schedules. In the future, we plan to develop and evaluate a prediction model of pedestrian behaviour based on the presented dynamic map. Next we will collect further data to verify our algorithms to general validity. Thus, we will analyse the improvement by using more data from the prediction model of periodic events. Furthermore, with a bigger walking path network we will expand the graph analysis beyond edges and nodes to more complex time and space dependencies of several nodes. With the larger amount of data, new event identification is necessary. Therefore, we will utilise big data methods to distinguish different kinds of nodes in the walking path network. Moreover we will analyse different detectors for periodic events, e.g. complexer latent mixed models.

References

- Ahmed, M., Karagiorgou, S., Pfoser, D. and Wenk, C., 2015. A comparison and evaluation of map construction algorithms using vehicle tracking data. *GeoInformatica* 19(3), pp. 601–632.
- Birant, D. and Kut, A., 2007. St-dbscan: An algorithm for clustering spatial-temporal data. *Data & Knowledge Engineering* 60(1), pp. 208–221.
- Brüntrup, R., Edelkamp, S., Jabbar, S. and Scholz, B., 2005. Incremental map generation with gps traces. In: *Intelligent Transportation Systems, 2005. Proceedings. 2005 IEEE, IEEE*, pp. 574–579.
- Chai, D., Forstner, W. and Lafarge, F., 2013. Recovering line-networks in images by junction-point processes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1894–1901.
- Edelkamp, S. and Schrödl, S., 2003. Route planning and map inference with global positioning traces. In: *Computer Science in Perspective*, Springer, pp. 128–151.
- Ellis, D., Sommerlade, E. and Reid, I., 2009. Modelling pedestrian trajectory patterns with gaussian processes. In: *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pp. 1229–1234.
- Ess, A., Leibe, B., Schindler, K. and van Gool, L., 2008. A mobile vision system for robust multi-person tracking. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–8.
- Fayazi, S. A., Vahidi, A., Mahler, G. and Winckler, A., 2015. Traffic signal phase and timing estimation from low-frequency transit bus data. *Intelligent Transportation Systems, IEEE Transactions on* 16(1), pp. 19–28.
- Feuerhake, U., Kuntzsch, C. and Sester, M., 2011. Finding interesting places and characteristic patterns in spatio-temporal trajectories. In: *8th LBS symposium*.
- Gatalsky, P., Andrienko, N. and Andrienko, G., 2004. Interactive analysis of event data using space-time cube. In: *Information Visualisation, 2004. IV 2004. Proceedings. Eighth International Conference on, IEEE*, pp. 145–152.
- Guralnik, V. and Srivastava, J., 1999. Event detection from time series data. In: *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM*, pp. 33–42.
- Ihler, A., Hutchins, J. and Smyth, P., 2006. Adaptive event detection with time-varying poisson processes. In: *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM*, pp. 207–216.
- Karagiorgou, S. and Pfoser, D., 2012. On vehicle tracking data-based road network generation. In: *Proceedings of the 20th international conference on advances in geographic information systems, ACM*, pp. 89–98.
- Kerper, M., Wewetzer, C., Sasse, A. and Mauve, M., 2012. Learning traffic light phase schedules from velocity profiles in the cloud. In: *New Technologies, Mobility and Security (NTMS), 2012 5th International Conference on*, pp. 1–5.
- Klinger, T., Rottensteiner, F. and Heipke, C., 2015. Probabilistic multi-person tracking using dynamic bayes networks. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences II-3/W5*, pp. 435–442.
- Kraak, M.-J., 2003. The space-time cube revisited from a geovisualization perspective. In: *Proc. 21st International Cartographic Conference*, pp. 1988–1996.
- Kuntzsch, C., Sester, M. and Brenner, C., 2015. Generative models for road network reconstruction. *International Journal of Geographical Information Science* pp. 1–28.
- Makris, D. and Ellis, T., 2002. Spatial and probabilistic modelling of pedestrian behaviour. In: *BMVC, Citeseer*, pp. 1–10.
- Makris, D. and Ellis, T., 2005. Learning semantic scene models from observing activity in visual surveillance. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 35(3), pp. 397–408.
- Nakaya, T. and Yano, K., 2010. Visualising crime clusters in a space-time cube: An exploratory data-analysis approach using space-time kernel density estimation and scan statistics. *Transactions in GIS* 14(3), pp. 223–239.
- Sakaki, T., Okazaki, M. and Matsuo, Y., 2010. Earthquake shakes twitter users: real-time event detection by social sensors. In: *Proceedings of the 19th international conference on World wide web, ACM*, pp. 851–860.
- Schindler, K., Ess, A., Leibe, B. and Van Gool, L., 2010. Automatic detection and tracking of pedestrians from a moving stereo rig. *ISPRS Journal of Photogrammetry and Remote Sensing*, 65 (6) pp. 523–537.
- Shen, J. and Cheng, T., 2016. A framework for identifying activity groups from individual space-time profiles. *International Journal of Geographical Information Science* pp. 1–21.
- Weng, J. and Lee, B.-S., 2011. Event detection in twitter. *ICWSM* 11, pp. 401–408.