Map Matching for Vehicle Guidance *(Draft)*

Britta Hummel
University of Karlsruhe, Germany

1 Introduction

The past years have revealed a dramatically increasing interest in the use of mobile Geographical Information Systems (GIS) in various automotive applications. Car navigation systems employ digital maps to guide the driver to the desired destination. Next generation driver assistance systems will use enhanced maps in order to present precise navigation hints (including speed limits, locations of gas stations, restaurants, etc.), and for assisted vehicle control. Furthermore, methods for autonomous enhancements of existing maps using video and lidar sensors are currently under development. An exhaustive overview of potential applications for upcoming mobile GI systems can be found in chapter 2.

All applications share the need for a robust assignment of the measured vehicle position to a road segment in the digital map. This process is called *map matching*. Since the emergence of the field in the 1970s (French (1989)), considerable progress has been reported (cf. the surveying articles of Bernstein and Kornhauser (1998), White *et al.* (2000), Quddus *et al.* (2003) and Lakakis *et al.* (2004)). However, users of navigation systems still encounter some erroneous map matching results.

On the contrary, manufacturers increasingly aim for a simple architecture of navigation systems. Hence, a growing number of navigation systems do not rely on multiple vehicle sensors - such as a combination of DGPS, odometer and gyro - but instead restrict themselves to GPS only. This is necessarily true for the newly evolving PDA (Personal Digital Assistant) navigation systems, which are identified in chapter 1 as key devices for next generation GIS, as well as for low cost in-car navigation solutions.

In this chapter a robust map matching algorithm is presented which exclusively relies upon information from a standard GPS receiver (however, the integration of data from an integrated GPS dead-reckoning unit is straightforward). Mobile phone based location systems can equally serve as input. The achievable accuracy of mobile phone location is examined in detail in chapter 11.

In contrast to standard map matching techniques the whole *vehicle path* is estimated for each time step within an iterative, statistically optimal Bayesian estimator (a Bayesian estimator using a different formulation has been developed by Scott and Drane (1994)) .

The algorithm is suitable for all maps using the standard road segment representation as piecewise linear links. An introduction into digital maps and GPS is given in NCHRP (2002), for example. The map database used by the authors is off-the-shelf

and frequently used in today's navigation systems. Errors of up to 40 metres with respect to ground truth data have been encountered.

Section 2 derives the Bayesian classifier for matching one single GPS position and orientation datum to the map. Incorporation of the classifier into an Hidden Markov Model in section 3 accounts for the vehicle position and orientation history, information about the road network topology, driving restrictions and the assumed driving direction.

The vehicle path estimation proves to be robust even for challenging inner city scenarios, some of which are shown in section 5. A further improvement of the quality of current navigation systems for platforms without access to in-vehicle sensors (i.e. odometer or gyro) is anticipated.

The reliable and accurate determination of the current user position is considered a prerequisite not only for automotive applications but for a wide variety of mobile Geographical Information Systems. The described algorithm is not particularly tailored to automotive applications and can thus be integrated into any mobile GIS requiring positioning information.

2   Bayesian classification of GPS data

Map Matching can be formulated as a stochastic classification task: The measured GPS position and orientation vector $\mathbf{x} = (x, y, \phi)^T$ is to be assigned to the road element $\mathbf{k}_i$ with highest a posteriori probability:

$$\hat{i} = \text{argmax}_i \ p(\mathbf{k}_i|\mathbf{x}) \ . \tag{1}$$

The map represents a road element as a line segment defined by its start and end vertex. Figure 1(a) illustrates the classification task.

A situation that clearly justifies the use of both position and orientation information is given in Figure 1(b). While standard position-based map matching procedures would erroneously assign the encircled position data to road element $\mathbf{k}_3$, the orientation data assist to their correct assignment.

The assumption of uniformly distributed a priori probabilities for the road elements $p(\mathbf{k}_i)$, together with Bayes formula (cf. Duda *et al.* (2001), for example), yields:

$$\text{argmax}_i \ p(\mathbf{k}_i|\mathbf{x}) \quad \overset{\text{Bayes Theorem}}{=} \quad \text{argmax}_i \ \frac{p(\mathbf{x}|\mathbf{k}_i) \ p(\mathbf{k}_i)}{p(\mathbf{x})}$$

$$\overset{p(\mathbf{X})\,\text{const.}}{=} \quad \text{argmax}_i \ p(\mathbf{x}|\mathbf{k}_i) \ p(\mathbf{k}_i)$$

$$\overset{p(\mathbf{k}_i)\,\text{uniformly distributed}}{=} \quad \text{argmax}_i \ p(\mathbf{x}|\mathbf{k}_i) \quad . \tag{2}$$

The class conditional probability $p(\mathbf{x}|\mathbf{k}_i)$ of the vehicle state measurement $\mathbf{x}$ when traversing road element $\mathbf{k}_i$ is modelled by two, statistically independent, random variables:

- The Euclidean distance b between vehicle position $\mathbf{x}$ and $\mathbf{k}_i$ is modelled as zero-mean, normally distributed random variable B with standard deviation σ_B.
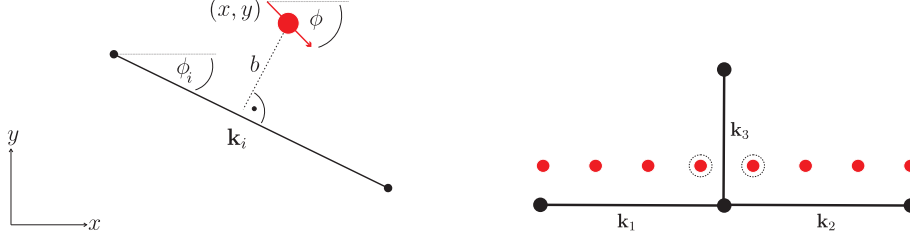
2

Figure 1: Red dots indicate GPS measurements, black line segments denote road elements from the map. *(a)* Classification task: Given a GPS measurement $(x, y, \phi)^{\mathrm{T}}$ the probability of being located on an arbitrary road element $\mathbf{k}_i$ has to be estimated. Orthogonal distance b and orientation difference thereby serve as criteria. *(b)* Example of erroneous assignment to road element $\mathbf{k}_3$ for the highlighted GPS measurements if only positional information is used.

- The angular difference $\delta\phi$ between vehicle orientation ϕ and the orientation of the road element ϕ_i is modelled as zero-mean, normally distributed random variable Φ. The standard deviation is σ_Φ.

The values for the standard deviations σ_B and σ_Φ have to account for the uncertainties in both the map and the GPS receiver data. The GPS orientation information becomes less reliable at lower speeds, therefore σ_Φ is chosen to be inversely proportional to the measured GPS speed.

Equation 1 can now be rewritten as the following Mahalanobis distance:

$$
\begin{aligned}
\hat{i} &= \operatorname{argmax}_i p(\mathbf{x}|\mathbf{k}_i) \\
&= \operatorname{argmax}_i p_B(b(\mathbf{x}, \mathbf{k}_i)) \cdot p_\Phi(\delta\phi(\mathbf{x}, \mathbf{k}_i)) \\
&= \operatorname{argmax}_i \frac{1}{\sqrt{2\pi}\sigma_B} \cdot \exp\left(-\frac{b(\mathbf{x}, \mathbf{k}_i)}{2\sigma_B^2}\right) \cdot \frac{1}{\sqrt{2\pi}\sigma_\Phi} \cdot \exp\left(-\frac{\delta\phi^2(\mathbf{x}, \mathbf{k}_i)}{2\sigma_\Phi^2}\right) \\
&= \operatorname{argmin}_i \frac{b^2(\mathbf{x}, \mathbf{k}_i)}{\sigma_B^2} + \frac{\delta\phi^2(\mathbf{x}, \mathbf{k}_i)}{\sigma_\Phi^2} \quad .
\end{aligned}
\tag{3}
$$

This yields the desired classifier for a single time instant. Figure 2 illustrates the properties of the classifier for one particular road element.

After assigning the current vehicle state to a road element, the vehicle position and orientation estimates are updated accordingly. The updated position is determined by the orthogonal projection of the GPS position on the assigned road element. The updated orientation equals the orientation of the road element.
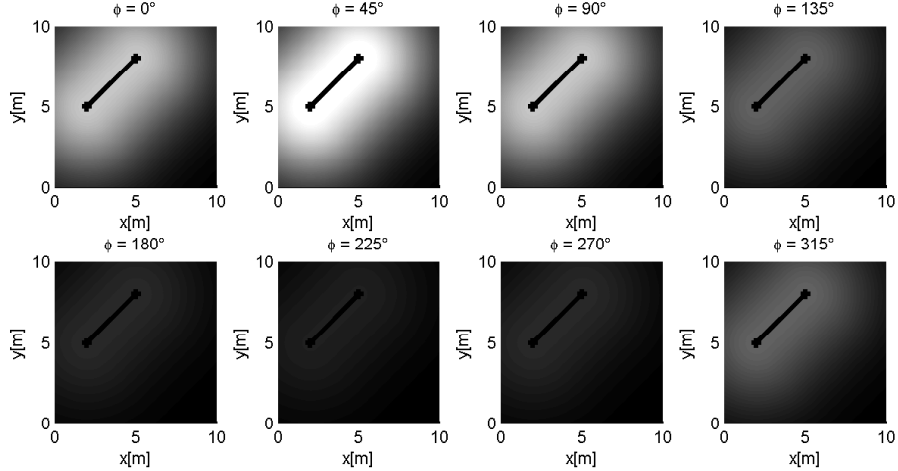
φ = 0° φ = 45° φ = 90° φ = 135°

y[m] x[m] y[m] x[m] y[m] x[m] y[m] x[m]

φ = 180° φ = 225° φ = 270° φ = 315°

y[m] x[m] y[m] x[m] y[m] x[m] y[m] x[m]

Figure 2: Properties of the classifier. The black line denotes a road element with an angle of 45° with respect to the x axis. The likelihood $p(\mathbf{x}|\mathbf{k}_i)$ is brightness-coded with respect to different vehicle orientations. The highest likelihood is observed for a vehicle orientation of 45°, the lowest for a vehicle driving just in opposite direction, i.e. 225°.

3 Incorporation of position history and network topology

Up to now, the proposed classifier exclusively uses position and orientation information from the current time step. Additionally, all information concerning the topology of the road network is discarded. A considerable increase of classification robustness can be achieved by including the following features in the classifier:

- Position history and orientation history: using all previously measured position and orientation data will lead to a significant reduction of the impact of gross measurement errors.

- Road network topology: considering the relations among different road elements will inhibit impossible consecutive map matchings (i e. a transition from road element $\mathbf{k}_i$ at time t to road element $\mathbf{k}_j$ at time $t+1$, although $\mathbf{k}_i$ and $\mathbf{k}_j$ are not connected).

The described features can be fully incorporated in the map matching process by the Hidden Markov Model described in the next section.

3.1 Hidden Markov Model (HMM)

An important class of Markov Models can be represented by a stochastic finite state machine, with state transitions and outputs being described by probability distributions. A Hidden Markov Model is defined by the five-tuple: state space, set of possible observations, transition probabilities, emission probabilities and initial state distribution. Duda *et al.* (2001), for example, provide an introduction to Hidden Markov modelling. Figure 3 depicts the proposed model. Each road element $\mathbf{k}_i$ constitutes one element of the state space. The emission probabilities $p(\mathbf{x}|\mathbf{k}_i)$ correspond to the classification rule from Eq. 2. The transition probabilities $p(\mathbf{k}_j|\mathbf{k}_i) = a_{ij}$ represent the road network topology: Two elements have a non-zero transition probability only if they share at least one vertex. No state transition is preferred: $a_{ij} = \frac{1}{s_i} \forall j$.
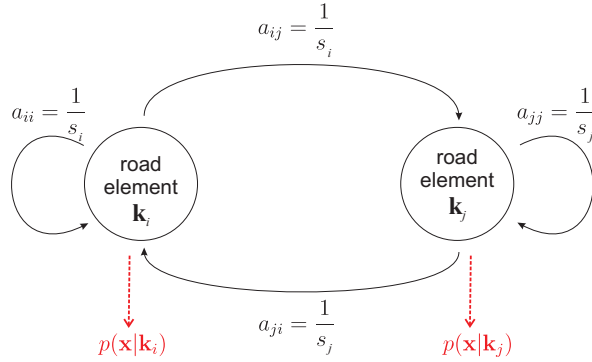

Figure 3: First-order Hidden Markov Model. Circles denote the model states, thin arrows denote state transitions. Road elements $\mathbf{k}_i$ and $\mathbf{k}_j$ are assumed to be connected. The dashed arrows indicate the output probabilities.

We can now formulate the optimum estimate for the path $\hat{\mathbf{i}} = (\hat{i}_T, \hat{i}_{T-1}, \ldots, \hat{i}_1)$ for an observed input sequence $\mathbf{x}_T, \mathbf{x}_{T-1}, \ldots, \mathbf{x}_1$ using the chain rule as:

$$\hat{\mathbf{i}} = \operatorname{argmax}_{\mathbf{i}} \left(\prod_{t=2}^{T} p(\mathbf{x}_t|\mathbf{k}_{i_t}) \right) \cdot p(\mathbf{k}_{i_t}|\mathbf{k}_{i_{t-1}}) \cdot p(\mathbf{x}_1|\mathbf{k}_{i_1}) . \tag{4}$$

The Viterbi algorithm is used for a minimum cost computation of the best path. It iteratively computes the statistically optimal sequence of state transitions for a given sequence of vehicle states.

3.2 Extended HMM

The Hidden Markov Model is further augmented by considering the roads' driving restrictions (i.e. one way streets) and, moreover, the assumed driving direction of the vehicle. Both are incorporated by the following model enhancements:

- The elements of the state space are enhanced by a flag denoting the driving direction. One road element can thus yield one (for oneways) or two elements in the state space.

- The transition probabilities between two state space elements are set to a very small value for contradictory driving directions, reflecting probability for doing a U-turn.

Figure 4 depicts the proposed model extension. Traversing a road element opposite to its allowed driving direction is no longer permitted. Additionally, paths with contradictory driving directions are assumed very unlikely.
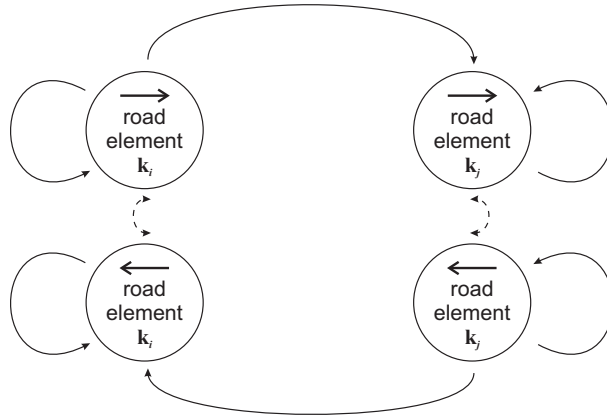

Figure 4: Extended Hidden Markov Model depicting the case where road elements $\mathbf{k}_i$ and $\mathbf{k}_j$ are bidirectional. Emission and transition probabilities have been omitted. Only state transitions with non-zero transition probability are shown. A dashed transition arrow indicates a low transition probability (reflecting the U-turn probability).

3.3 Detection of erroneous map topology

The *maximum a posteriori* probability given by Eq. 2 can directly be used as a measure of goodness of the classification result. A very low value indicates a coarse GPS measurement error, as already indicated by low horizontal/vertical dilution of precision (HDOP/VDOP) values within the GPS receiver protocol, or a modelling error. Modelling errors refer to an erroneous map topology due to missing road elements. Hence, the algorithm inherently provides a means for detecting erroneous map data.

In the case of a detected model error the Hidden Markov Model is reset by discarding all previously aquired position data.

4   Revised position estimate

The proposed classifier assigns GPS data to the most likely road segment of the digital map. This allows for a subsequent update of the vehicle state estimate. Within this contribution the updated position is determined by the orthogonal projection of the GPS position on the assigned road element. The updated orientation equals the orientation of the road element. Another possibility of computing the updated position estimate using the vehicle speed data from the GPS sensor for dead-reckoning is described in Ochieng *et al.* (2003).

5   Examples in complex urban environment

The proposed map matching has been successfully tested on an experimental vehicle in the inner city of Karlsruhe, Germany. All tests are run on standard hardware (Pentium 4, 2 GHz processor). The digital map is commercially available and frequently used in today's navigation systems. A standard low-cost GPS receiver without differential corrections is used with an estimated standard deviation of position and orientation measurements of 10-15 metres and $15°$. Data is acquired at 1 Hz. The processing time of the algorithm is 0.01 seconds per GPS datum. Real time performance can thus be achieved on systems up to hundred times slower, for example on PDA processors. Several test runs were performed with a total amount of more than four hours of on-line testing (equalling about 15000 measured GPS data points) in dense urban area. Examples of the computed path results are shown in Fig. 5.

Few intermediate misclassifications occurred for 0.4% of the data points due to severe deviations of the GPS measurement (up to 80 metres with respect to the correct road element) or due to occasionally coarse map digitisation (up to 40 metres deviation from ground truth data). All of those cases were based on the following configuration: The vehicle was standing still close to an intersection and the GPS points were slowly drifting away from the true position. It is believed that those few cases will elegantly be circumvented by preferring self-transitions over transitions to any follow-up road elements for low vehicle speeds within the emission probabilities of the Hidden Markov model.

All intermediate misclassifications are completely compensated by the algorithm leading to an completely error-free posterior vehicle path estimate! Figure 6 illustrates how an intermediate misclassification in a very complex road configuration was automatically corrected by the HMM towards the correct path as soon as enough measurements corroborated the belief in the correct path. One exceptional case leading to one erroneous path estimate was observed which is analyzed in Fig. 7.

The travelled route contained three map topology errors, referring to missing road elements. All three cases have been successfully classified as *modelling error* by the algorithm (cf. Sec. 3.3).

6 Conclusion and Future Developments

A map matching method has been presented which exclusively relies upon information from a standard GPS receiver. The method exploits the vehicle position and orientation history, information about road network topology, driving restrictions and the assumed driving direction for each road element. A Hidden Markov Model has been formulated leading to a statistically optimal, iterative Bayesian estimation procedure. In contrast to conventional map matching, the whole vehicle path is estimated at each time step. An initialization procedure as required by other methods is not needed. The method has proven to be robust even on challenging inner city scenarios. Real-time performance has been shown. An improvement of navigation quality for platforms without access to vehicle sensors (odometer or gyro, respectively) is anticipated by the method. This is especially the case for handheld navigation systems. Furthermore, the studies revealed that additional vehicle sensors are not necessary in order to yield robust map matching results.

The current contribution focused on *global* map matching, i.e determining the most likely road segment within the network. However, *local* map matching, i.e. determining the lateral vehicle pose within the road segment still remains an open challenge (an approach using Differential GPS is described in Du *et al.* (2004)). Current research focuses on video sensor based estimation of the number of lanes and the subsequent estimation of the vehicle ego pose, namely lateral offset and orientation with respect to the road segment. Further work concerns methods for an automated extension of the digital map by the estimated attributes.

References

Bernstein, D. and Kornhauser, A. (1998). 'An introduction to Map matching for personal navigation assistants'. In *The Transportation Research Board - 77th Annual Meeting, Washington, D.C.*

Du, J., Masters, J., and Barth, M. (2004). 'Lane-Level Positioning for In-Vehicle Navigation and Automated Vehicle Location (AVL) Systems'. In *Proceedings of the Intelligent Transportations Systems Conference*.

Duda, R. O., Hart, P., and Stork, D. G. (2001). *'Pattern Classification'*. John Wiley & Sons, New York.

French, R. L. (1989). 'Map Matching Origins, Approaches and Applications'. In *Proceedings of the Second International Symposium on Land Vehicle Navigation*, pages 91–116.

Lakakis, K., Savvaidis, P., Ifadis, M., and Doukas, D. I. (2004). 'Quality of Map-Matching Procedures Based on DGPS and Stand-Alone GPS Positioning in an Urban Area'. In *Proceedings of the FIG Working Week*.

NCHRP (2002). *'NCHRP Synthesis 301: Collecting, Processing, and Integrating GPS Data into GIS'*. National Academy Press, Washington, D.C.

Ochieng, W. Y., Quddus, M. A., and Noland, R. B. (2003). 'Map-Matching in Complex Urban Road Networks'. *Brazilian Journal of Cartography (Revista Brasileira de Cartografia)*, **55**(2), 1–18.

Quddus, M. A., Ochieng, W. Y., Zhao, L., and Noland, R. B. (2003). 'A general map matching algorithm for transport telematics applications'. *GPS Solutions Journal*, **7**(3), 157–167.

Scott, C. A. and Drane, C. R. (1994). 'Increased Accuracy of Motor Vehicle Position Estimation by Utilising Map Data, Vehicle Dynamics and Other Information Sources'. In *Proceedings of the Vehicle Navigation and Information Systems Conference*.

White, C. E., Bernstein, D., and Kornhauser, A. L. (2000). 'Some map matching algorithms for personal navigation assistants'. *Transportation Research Part C*, **8**(1-6), 91–108.
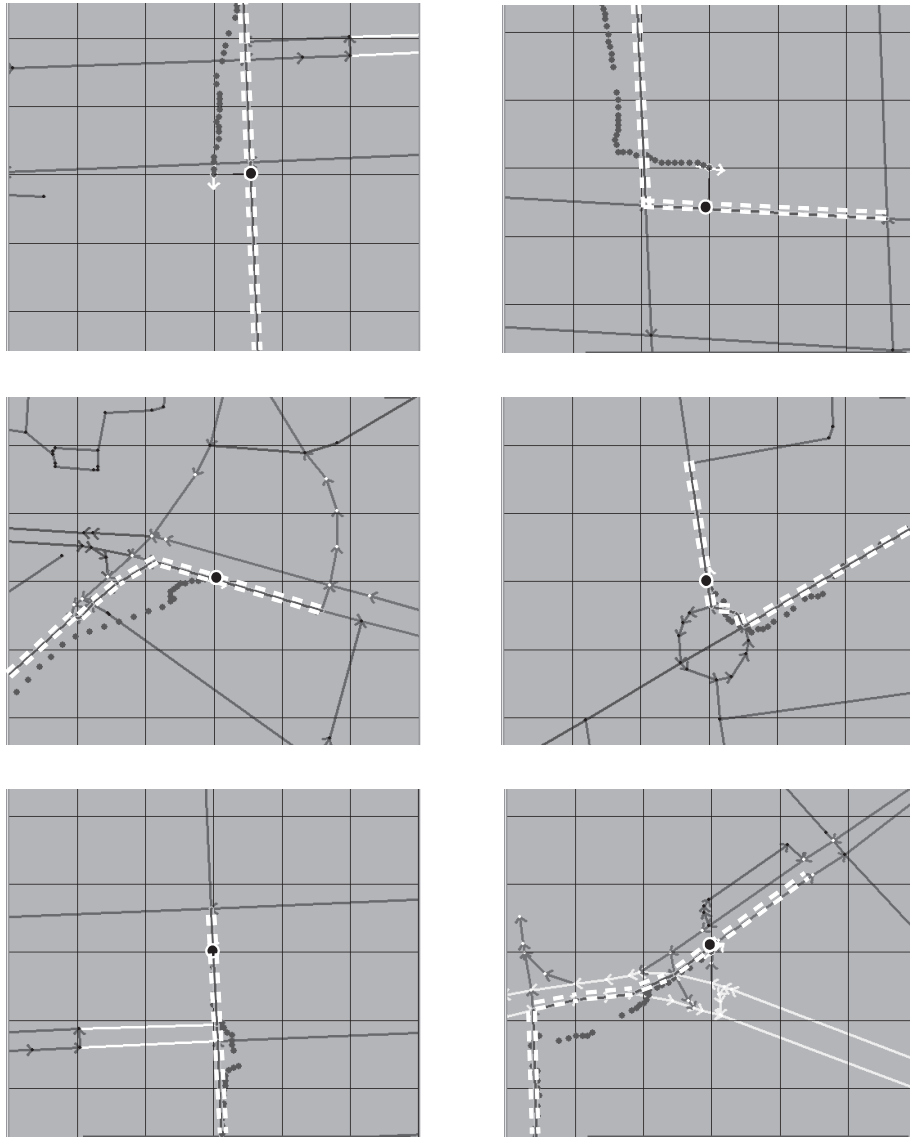
Figure 5: Map Matching results in complex situations. One element of the background grid covers a 50 m area. The small dots correspond to measured raw GPS positions. The computed path is marked by white dashed double lines. The black dot with white surrounding corresponds to the map matched vehicle position for the current time step. Despite severe deviations between GPS measurements and the road elements from the map, the correct path (according to the classification by a human observer) has been successfully found in all situations.
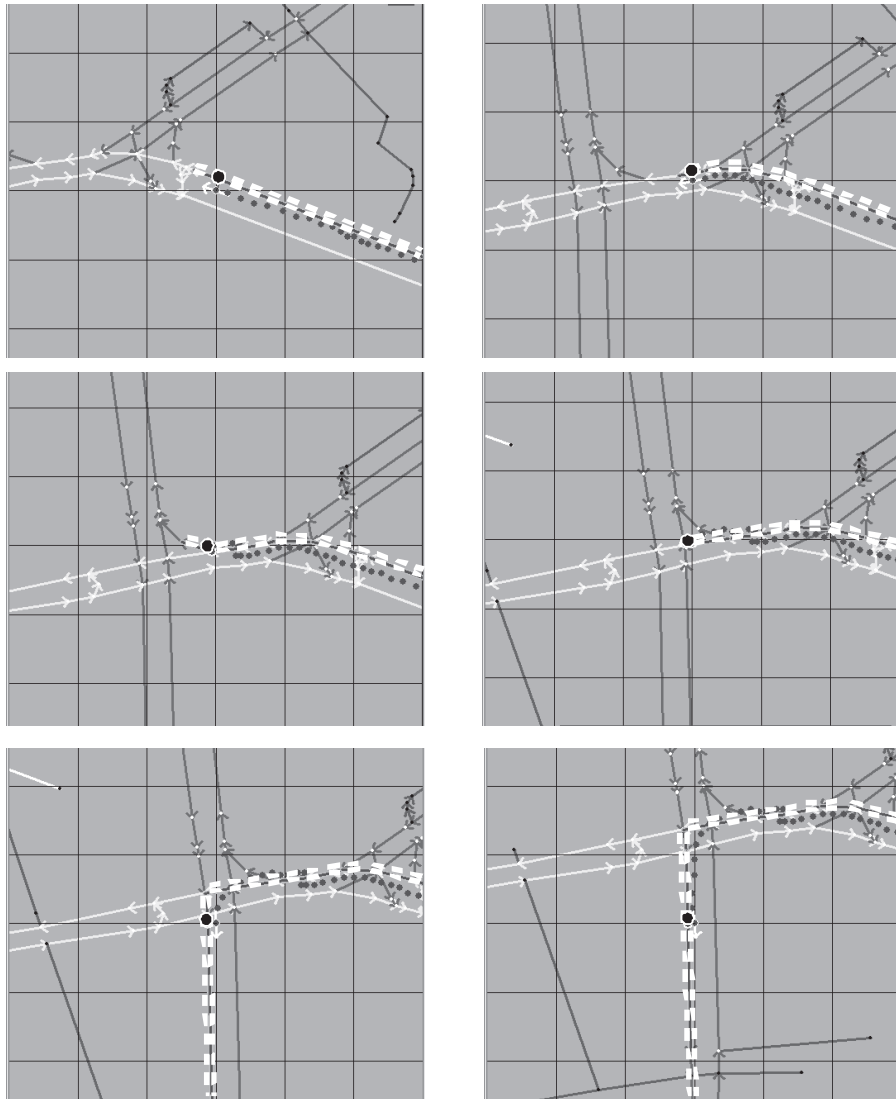
Figure 6: Sequence of Map Matching results in a complex situation. An initially correctly assigned path (top left and top right) was intermediately misclassified due to erroneous GPS and map data (middle left), but was corrected towards the correct path as soon as enough measurements had corroborated the belief in the correct path (middle right and bottom).
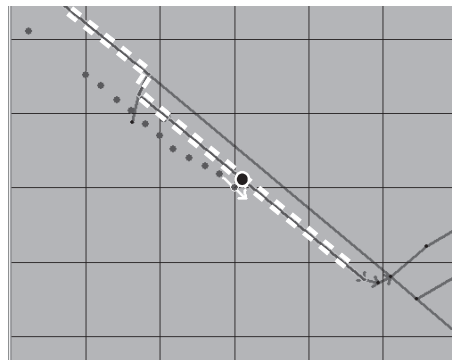
11

Figure 7: Erroneous Map Matching result. The vehicle was erroneously assigned to the lower, parallel running road element although the measured vehicle orientation didn't indicate any right turn. This is due to the fact that the positions of approximately twenty follow-up position measurements yielded a significantly larger error for the upper, correct element, leading to a larger overall error. That case could only be corrected if the standard deviation of the orientation would be set to a very small value compared to the standard deviation of the position. This decision is not justified because of the often large orientation deviations of the road elements in the map compared to ground truth.