Demo Abstract: Lightweight Continuous Indoor Tracking System

Zhuoling Xiao, Hongkai Wen, Andrew Markham, and Niki Trigoni
Department of Computer Science, University of Oxford. Email: firstname.lastname@cs.ox.ac.uk.

***Abstract*—Indoor tracking and navigation is a fundamental need for pervasive and context-aware smartphone applications. Although indoor maps are becoming increasingly available, there is no practical and reliable indoor map matching solution available at present. In this demo abstract, we describe a working prototype of a novel, robust and responsive tracking technique that is extremely computationally efficient (running in under 10 ms on an Android smartphone), does not require training in different sites, and tracks well even when presented with very noisy sensor data. The tracking system requires zero user effort (wardriving, fingerprinting, etc.) – only the floor plan is required. We also demonstrate how it is able to accurately track the position of a user from accelerometer and magnetometer measurements only (i.e. gyro- and WiFi-free). We believe that such an energy-efficient approach will enable always-on continuous background localisation, enabling a new era of location-aware applications to be developed.**

I. Introduction

Whereas GPS is the *de facto* solution for outdoor positioning, no clear solution has as yet emerged for indoor positioning despite intensive research and the commercial significance. Applications of indoor positioning include smart retail, navigation through large public spaces like transport hubs, and assisted living. The ultimate objective of an indoor positioning system is to provide continuous, reliable and accurate positioning on smartphone class devices [1]. We identify maps as the key to providing accurate indoor location. Based on a time-series of observations, such as inertial trajectories or RF scans, the goal is to reconcile the observations with the constraints provided by the maps in order to estimate the most feasible trajectory of the user, i.e. the sequence that violates the fewest constraints.

In this demo abstract, we describe a working prototype of continuous indoor tracking system that is lightweight and computationally efficient, but also robust to noisy data, allowing it to provide always-on and real-time location information to mobile device users. The proposed system uses an undirected graphical model, known as linear chain conditional random fields (CRFs) [2] which is particularly flexible and expressive, which allows us to capture correlations among observations over time, and to express the extent to which observations support not only states, but also state transitions.

The working prototype is very computationally efficient, running in <10 msec on an Android phone, enabling real-time location computation online. It also offers high location accuracy even when it uses ultra low power sensors (e.g. accelerometer and magnetometer). With the trend of low power digital motion processors (DMPs), e.g. InvenSense MPU-6000/MPU-6050 which are able to task and process inertial data in bursts, while the system processor remains in a low-power sleep mode, our system could offer an always-on tracking service for a large range of low power devices.

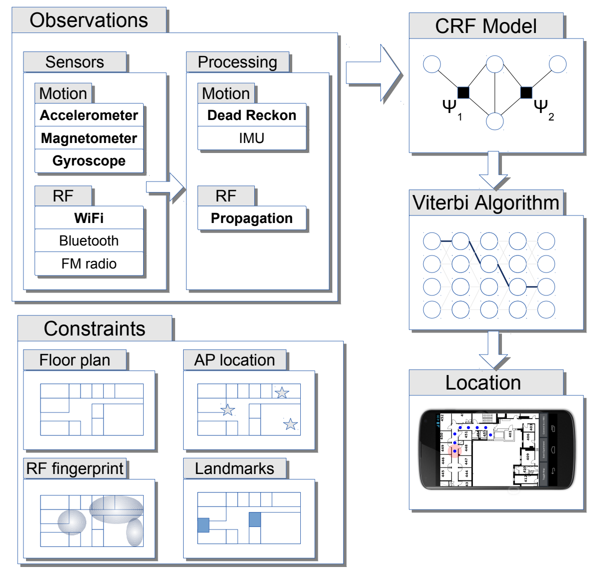

Fig. 1. Overview of the tracking system.

II. Tracking System Model

The system architecture is shown graphically in Fig. 1, and is described through the use of an example.

When a user enters a building and launches the tracking application, the application requests a floor plan (along with other meta-data as generated by other systems, which could include fingerprint maps) from the server, if not already within the cache. Note that this is the only time that a user needs to reveal any data about their coarse position to a third party. The floor plan provides constraints over the set of possible positions a user can take, as well as allowed transitions between locations (i.e. a user cannot directly travel from one end of the building to the other without visiting intermediate locations). As such, the floor plan forms a sparse graph and can thus be efficiently stored in memory. Sensors on the user's phone collect data about the motion and (radio) environment. Motion sensors can include accelerometers, magnetometers and gyroscopes. Radio sensors can include WiFi, Bluetooth (low energy), FM radio and so forth. Raw sensor data is typically not immediately usable and needs to be processed. In the case of motion data, this could include dead reckoning trajectories based on counting steps and estimating heading, or

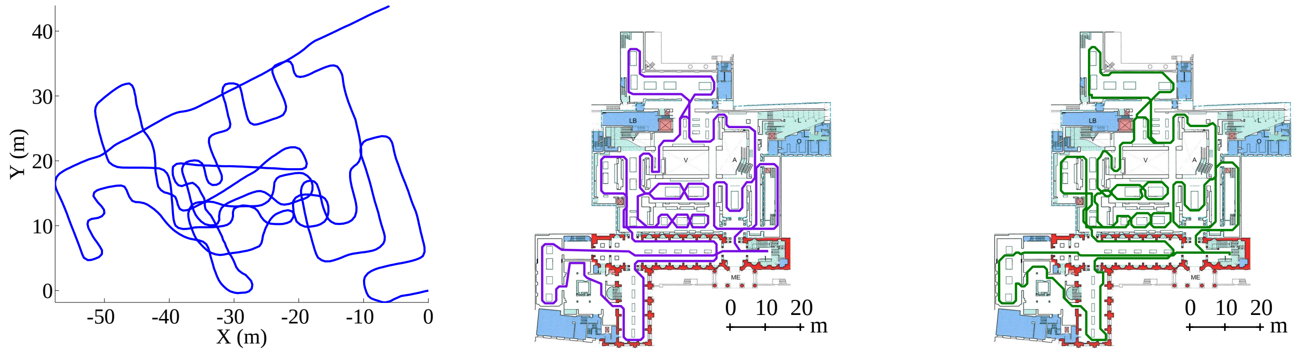

(a) Raw trajectory (b) Ground truth trajectory (c) Matched trajectory

Fig. 2. Experiments in the museum (109m × 89m), showing raw, ground-truth and matched trajectories.

using full IMU tracking in the case of foot mounted sensors. For RF data, a channel/propagation model can be used to relate received signal strengths to physical distances. Alternatively, raw signal strengths may be directly forwarded to the CRF model, to be later combined with RF fingerprint map data if available.

Maps and observations are combined using conditional random fields, an undirected graphical model which is particularly well suited to this sequential problem because it allows us to flexibly define *feature functions* that capture the extent to which observations support states and state transitions, given map constraints. As a user moves through the building, certain paths become unlikely, as they violate map constraints.

More specifically, this process involves four distinct steps:

Map pre-processing

This step takes a floor plan as input, and produces a graph that a) encodes a set of discrete states (locations), and b) represents physical constraints between discrete states imposed by the map. This information will then be fed to the second step, to help us define the CRF's states and feature functions.

Definition of states and feature functions

Our system uses feature functions to elegantly model different sensor data including the inertial measurements, RF measurements, visual measurements, user inputs, etc. and uses potential function to fuse them together;

Training to determine feature weights (optional)

With one or more true trajectories paired with respective sequences of sensor observations, training the CRFs to estimate weights is then performed by maximising the conditional probability of states given observations. Our system works well without training but training can help the tracking system to capture the special per-site features, which could help improve the tracking accuracy.

Inference to estimate location over time

The Viterbi algorithm is used to efficiently find the most likely sequence of states through the transition graph, culminating in an estimate of the user's location and quality thereof.

The first three steps are performed once for each building. The fourth step is performed online on the user's smartphone to track themselves.

Fig. 2 shows an example of our experimental results in the museum, including the raw trajectory, ground truth, and matched trajectory, demonstrating an RMS error of 1.14m.

III. Demo

In the demo session, we will show a working prototype of the continuous tracking system that offers accurate pedestrian location information without delay. The tracking system requires only the floor plan of the test site. In addition, to demonstrate the robustness of our system in different environments, we will also show the videos of our experiments in three other experimental sites: an office building, a museum, and a market.

IV. Conclusion

We demonstrated the merit of a novel continuous indoor tracking system, based on the application of conditional random fields. We have shown how it is robust, being able to operate with very noisy sensor data; lightweight, running in under 10 ms on a smartphone; and accurate, achieving the lowest RMS errors compared with other state-of-the-art approaches. Our system is able to establish a user's position using only dead-reckoned trajectories and a floorplan, without any external information such as a starting location or knowledge of WiFi access point locations. We believe that our tracking system has widespread application to a number of domains, as this single approach can be used with a wide variety of sensors and map information. One particularly relevant area is estimating location online and in real-time in resource-constrained body-worn sensors. In summary, we have presented a system that addresses the very pressing problem of providing accurate, low power, indoor tracking, that is responsive, robust and scalable.

Acknowledgements The authors would like to acknowledge the support of the EPSRC through grants EP/J012017/1 and the anonymous reviewers for their helpful comments.

References

[1] H. Wen, Z. Xiao, N. Trigoni, and P. Blunsom, "On assessing the accuracy of positioning systems in indoor environments," in *EWSN*, 2013.

[2] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *ICML*, 2001.