



Workshops der
Wissenschaftlichen Konferenz
Kommunikation in verteilten Systemen 2009
in Kassel
(WowKiVS 2009)

Decentralized Probabilistic World Modeling with
Cooperative Sensing

Arjan Peddemors and Eiko Yoneki

8 Pages

Decentralized Probabilistic World Modeling with Cooperative Sensing

Arjan Peddemors

Eiko Yoneki

INCA group, Telematica Instituut
The Netherlands
arjan.peddemors@telin.nl

Computer Laboratory, University of Cambridge
United Kingdom
eiko.yoneki@cl.cam.ac.uk

Abstract: Drawing on the projected increase in computing power, solid-state storage and network communication capacity to be available on personal mobile devices, we propose to build and maintain without prior knowledge a fully distributed decentralized large-scale model of the physical world around us using probabilistic methods. We envisage that, by using the multimodal sensing capabilities of modern personal devices, such a probabilistic world model can be constructed as a collaborative effort of a community of participants, where the model data is redundantly stored on individual devices and updated and refined through short-range wireless peer-to-peer communication. Every device holds model data describing its current surroundings, and obtains model data from others when moving into unknown territory. The model represents common spatio-temporal patterns as observed by *multiple* participants, so that rogue participants can not easily insert false data and only patterns of general applicability dominate. This paper aims to describe – at a conceptual level – an approach for building such a distributed world model. As one possible world modeling approach, it discusses *compositional hierarchies*, to fuse the data from multiple sensors available on mobile devices in a bottom-up way. Furthermore, it focuses on the intertwining between building a decentralized cooperative world model and the opportunistic communication between participants.

Keywords: probabilistic modeling, cooperative sensing, opportunistic communication, compositional hierarchy, mobility modeling, social networks

1 Introduction

More and more people carry with them powerful computing devices on a nearly continuous time basis, mostly in the form of personal devices such as advanced mobile phones. Indeed, on a world-wide scale, the number of mobile phones now exceeds the number of more traditional computing devices such as PCs. These personal computing devices often contain a rich set of different sensors that can be used to almost continuously perceive the surroundings of the owner. Although very distinctive from the physiological senses, they certainly are able to capture relevant features such as owner mobility, network location, and social interaction as well as – with more recently deployed sensors – light, acceleration, device touch, geographical location and limited audio and video. The true revolution of this development lies in the fact that, for the first time in history, we have a computing platform capable of sensing the direct environment of a substantial part of the world population. This platform is not only capable of enhancing the perception of individuals by sensing beyond what is possible with the human



senses, but also facilitates the easy sharing of these perceptions by means of ubiquitous data network technologies.

In this paper we seek to explore at a conceptual level the possibilities to deliver an accurate world model to personal computing devices based on the perception of this world by a large cooperating group of devices. Such a world model is important to many applications, as a primary input or as a source for enhancement of existing functionality. For instance, it may reliably tie stationary objects to geographical locations (a spatial feature) and it may reflect that a main road has been blocked in a certain direction because the associating frequently traveled pattern, represented by a common sequence of observations in time, suddenly is not used anymore (a temporal feature). By sharing the perception of the world between peers in the group, we aim to 1) increase the accuracy of the perception as taking place on a single device – by taking into account what is perceived as common between devices – as well as 2) to provide a model of environments that are new to an owner and have not been observed in the past, drawing solely on what has been sensed by others.

A fundamental question is how to represent the perception as data that can be stored on, used by, and exchanged between computing devices. The data representation must accommodate the different kinds of data coming from the sensors and the fused data from multiple sensors. Also, it must be capable to capture spatio-temporal aspects as well as capture uncertainty. We require a model building process that captures elementary patterns directly derived from simple sensor information such as the concurrent occurrence of sensor values (spatial) or the strong correlation between occurrences of sensor events in time (temporal), and that is capable of generating knowledge at higher levels of abstraction. Another fundamental question is the impact of the distribution of perception data between mobile peers on the timeliness and accuracy of shared data available on individual mobile devices. A primary aspect is the safeguarding of the privacy of individuals in a group of participants and the protection against the insertion of corrupt or false data by an – intentionally or unintentionally – rogue participant. As a consequence, we focus on peer-to-peer (P2P) distribution mechanisms, also applied in various forms in internet file sharing, because a fully distributed P2P mechanism lacks a centralized role which may easily compromise privacy. We aim for a completely decentralized infrastructure, which only relies for its world model building on what is sensed on the devices and shared between devices.

This paper is organized as follows. Section 2 discusses modeling frequent patterns in a probabilistic way, and provides a more in-depth overview of ‘compositional hierarchies’ as a possible way to extract a model structure from observed sensor data on a single device. When capable of structuring frequent patterns on a single device, we focus in section 3 on the exchange of model data using opportunistic communication. Section 4 provides an example application, putting sections 2/3 in context. Section 5 discusses some of the open issues and we wrap up with a summary.

2 Bottom-up probabilistic modeling with compositional hierarchies

When modeling the physical world in a generic and probabilistic way, we are interested in finding frequent patterns in the stream of data coming from the sensors that perceive our surroundings. One way to tackle this is by starting from a predefined structure of the world and then learn our surroundings using primitives from this structure – primitives like roads, buildings, people, cars, etc. This approach has a number of disadvantages: it is a very large endeavor to manually construct a sufficiently detailed set of primitives and their

interrelationships and we may not always have knowledge about environments where we want to apply our technology. Furthermore, perhaps most importantly, imposing a predefined structure makes it hard to identify unanticipated but important relationships between occurrences and objects around us. Therefore, we prefer a model that has no such predefined structure.

Assuming no a priori structure in the world around us, it is necessary to learn these patterns from the data only. A specific example of building a dynamic world model is [3], where 3D objects are sensed using ultrasonic ray tracing. More close to the mobile device platform, the work in [7] predicts wireless network visibility events based on the detection of other recent events. Also, the area of probabilistic robotics [10] is closely related to this problem. One possible way to achieve this is by means of *compositional hierarchies* [11][8]. We describe this approach here and use it later to illustrate the interaction between local model building (on a single device) and distributed global model building by means of opportunistic communication between devices.

A compositional hierarchy represents *part-whole relationships*. It can be used to group together sensor values that frequently occur together, in the spatial as well as the time domain. When grouped together as a primitive, a primitive itself can subsequently be used to form new, higher level primitives, building in this way a whole hierarchy of primitives.




Figure 1: primitive building for a spatial and spatio-temporal compositional hierarchy

The figure above depicts an example of the learning over time of a purely spatial (top) and a mixed spatio-temporal compositional hierarchy (bottom). The spatial hierarchy captures frequent patterns that link persons and a group to a specific location as observed from a single device. Initially, the observer identifies simple patterns such as being in close range of person A and B and being in room R (left hand side). Then, it identifies that these simple patterns sometimes occur concurrently, resulting in new primitives at a higher level (middle), followed by the identification of an even higher level primitive that describes group meetings in room R (right). Note that for reasons of clarity we have labeled the simple patterns and primitives with a meaningful description: the construction of the hierarchy, however, in no way relies on manually generated labels and only identifies frequent patterns and assigns probabilities to them. Note also that being in or close to a room and being in company with somebody else may be sensed with network interfaces available on many current mobile phones, such as Bluetooth and 802.11 wireless LAN.

The spatio-temporal hierarchy in the example above is constructed in a similar manner, showing not only sensing data available concurrently, but also spread out in time – for the higher level primitives. Again, it starts with identifying simple frequent patterns such as being at location A and B and having a high speed (left). Speed and location are joined in spatial primitives (middle) and these are again joined in a primitive that indicates being on a highway in between ramps X and Y (right). This top level primitive represents a temporal pattern: it registers the probability of observing high speed at location B *some time after* observing high speed at location A. So, when observing that ‘being at location A with high speed’ is true, it provides the probability of ‘being at location B with high speed’ a little later. And, when ‘being at location A with high speed’ and ‘being at location B with high speed’ have occurred after each other, it gives the probability of ‘being at location C with high speed’ some time later. Obviously, speed and location may be provided by a single sensor (GPS-based) but may also be derived from data from separate sensors, such as cellular network location and the typical sound associated with a car driving fast.

A compositional hierarchy can be used to predict missing or future data. If, for example, the above observer almost exclusively meets person A in room R, sensing the company of A also assigns high probability to being in room R, even without the need to sense the current location in the building. Or, as explained above, when consecutively observing high speed at location A and then at B, it is likely to later observe high speed at location C, or – one step down – being later at location C. This has benefits when it is expensive to sense or when, on a particular device, not all sensors are available, or simply to anticipate upcoming events (see also section 4).

Hierarchical compositional structures play an important role in existing work such as hierarchical hidden Markov models, certain data compression algorithms, and hierarchical reinforcement learning. When building compositional hierarchies, the pruning of infrequent patterns is essential to keep the necessary amount of memory within limits. The approach in [8] focuses on the on-line aspects of building hierarchies – using, amongst others, *hierarchical n-grams* – taking efficient infrequent pattern removal into account. However, even when pruning aggressively, it is likely that world model hierarchies need substantial storage. Also, the amount of hierarchical data depends strongly on the amount of event data coming from the sensors. If, for instance, location and speed are indicated with a high granularity, it may be possible to find hierarchy entities for every location and speed level combination.

3 World model creation using opportunistic communication

Once a device has collected sufficient amounts of data from its sensors to build an initial compositional hierarchy of frequent patterns, it may start to engage with others in the community in a P2P fashion to share model data, using means of short range communication. The incentive of sharing data with others comes from the potential improvements of the model data, for refining patterns that are self-observed or for obtaining patterns in new surrounding [6]. Data exchange may happen among community members or spontaneous neighborhood nodes. In this section we address potential communication mechanisms and *open issues* we see as important for such an exchange mechanism to work.

Every device has its own local view of the model, representing frequent patterns inferred from its direct surroundings. When comparing primitives with similar ones from others, the primitive probabilities must be updated to incorporate the perception of both devices (and increase accuracy). Exactly how this process takes place depends on how the compositional hierarchy is represented, where the choice of representation must take into account the possibility for cross hierarchy updates. It is likely to involve exchanging recent raw sensory data. A fundamental problem is that the raw sensor data may have substantial influence on how the hierarchy is build over time: the same frequent patterns may be captured by equally performing but very different hierarchy primitives, as shown in [11]. This means that individual participants may form different hierarchies simply by the order in which they have observed their surroundings, even when these surroundings are exactly the same.

The spread of model data depends on the possibilities of opportunistic communication between peers. Obviously, if the number of participants is small and dispersed over a larger geographical area, the possibilities for data exchange are minimal. An open question is how fast observed changes in the model can be communicated to other nearby participants. In the Huggle project [4][1], we introduced a new communication paradigm using dynamic interconnectedness via wireless enabled devices, called a Pocket Switched Network (PSN). People carry devices in their pockets, which communicate directly with other devices within their range or with the infrastructure. As people move around they can exchange messages with nearby devices, carrying a message for someone until it is close to another device. We explore epidemic/gossip mechanisms to achieve fully decentralized communication, combined with infrastructure networks where possible. Messages are spread via epidemic mechanisms, which are robust against disconnection, mobility and node failures, and are simple and decentralized. The experiments show a great potential of efficient data spread through proximity communication. Because the communication is human-to-human, we integrate social networks for improving communication efficiency. Our approach uses the centrality and community structures of social networks to deploy various types of communication mechanisms [4]. We have also shown how to uncover social structures in a distributed manner [5], where we discover cliques or tightly connected clusters, i.e. communities. To bootstrap a hierarchy quickly for someone who starts participating, it may be possible to copy a full tree from another trusted person (when both agree).

When exchanging information between nodes, a participant may disclose information that is privacy sensitive, even though no semantics are coupled with these patterns because the semantics can be added later. For instance, a frequently observed pattern is seeing a Bluetooth node in combination with a geographical location, which may link the observing participant to another person at that location because the Bluetooth node corresponds with that person's



mobile device. This not only exposes personal information of the observing participant, but also the personal information of others, even when not participating in the community. Privacy breaches are somewhat limited because model data is only dispersed locally and every participant only keeps track of the locally relevant patterns. Nevertheless, a privacy protecting mechanism must be part of the overall mechanism of cooperative model building. One possible approach to protect a participant's privacy is to not disclose, when exchanging model data with others, whether this data comes from personal observations or from the observations of somebody else. This requires that participants relay a fraction of the model data that is not observed by them.

A known problem of ad hoc data distribution in large communities is the possible insertion of false data into the network. A participant may benefit from giving wrong information to others when this leads to different behavior of these others. By checking the integrity of a part of the data received from others, it is possible to prevent flooding the whole community with false data, as shown in [2]. This solution, however, depends on signing the original content with a digital signature, which conflicts with our approach because we do not want to disclose the identity of the participant inserting the pattern into the network and, furthermore, we want to update patterns depending on the perception of others (which changes the original content). Another problem is free loader behavior of a part of the participants, resulting in a reduced community effect. This, however, can most likely be efficiently tackled with a 'tit-for-tat' mechanism when exchanging model data between peers.

4 Application examples

Having a probabilistic model of a person's surroundings (spatial) and the changes in these surroundings (temporal) can be useful for a wide range of applications. It provides information about, for example, locations where resources are likely to be available or where and when this person is likely to meet with other people. The example we describe here employs such a generic model to determine whether a person is in an *unusual situation*, or is likely to be at a location with an unusual situation in the future: it is a case of metering the *level of disorder* we observe in a person's surroundings. We take two cases to illustrate this 'uncommon state' example: a case where we detect unusual inter-person meeting and a case where we detect uncommon traffic jams.

Suppose we have a community consisting of individual nodes building up a hierarchy of locally perceived frequent patterns. These patterns are exchanged in a P2P fashion within the community, resulting in a model per node that captures common community knowledge. Some of these primitives in the hierarchy correspond with very stable observations: the nodes have high confidence that when observing one part of the primitive, the other part(s) are also observed. When a person meets another person always in the same location, the primitive joining the other person and the meeting location assigns high probability to being in the meeting room when the other person is detected. Or, on a stretch of high way with almost no traffic jams, being at location A with high speed assigns high probability to being at location B with high speed a little later (as in the spatio-temporal part of Figure 1) – higher than at stretches with frequent traffic jams.

Now, at a certain point in time, observations are made that contradict the high probability of a pattern seen in the past. These observations mark an uncommon situation. When you meet another person in a location not the same as where you always meet this person, this marks an uncommon situation. A contradiction can also be made by others participating in the

community: when they detect that a roadblock breaks the high confidence pattern described above (having high speed at location A does no longer imply having high speed at location B a little later), this information is spread in a P2P fashion amongst nearby nodes, also reaching those nodes that are interested in this information but have not observed this contradiction themselves. This information then can be used to predict that a person will enter an uncommon state in the future, possibly far ahead in time, if it is clear that this person will be at location B in the future.

Although we have taken unusual inter-person meetings and unexpected traffic jams as cases, measuring the level of disorder can also be useful for detecting unusual crowds, alarm levels of noise, and natural disasters.

5 Discussion

We realize that there are many open ends in our approach to build a distributed probabilistic world model as a community effort. Some of these are discussed below.

One thing we need to consider is the difference in sensor quality and calibration, which may interfere with building a common model. Because of the slight differences in perception, the hierarchies build by individuals will be slightly different and therefore not shared easily. In particular the lower levels in the hierarchy may be affected by this.

We have to deal with the fact that different temporal patterns evolve at widely varying rates. This is nicely illustrated in [9]. The hierarchical modeling is likely to provide this feature, as low level features happen on a short time scale and high level features on a longer time scale, but this needs to be further investigated.

We propose to build a bottom-up model, without predefined structure. This means that semantics are not automatically assigned to hierarchy primitives. The label 'driving on highway between ramp X and Y' for a certain primitive is just that: a label we put on it to illustrate an example. If this entity is active, it does not automatically have meaning to an application. This can be considered an impedance mismatch between the probabilistic world and the world of application logic. One mode of usage could be that we associate application actions – again using probabilistic techniques – with hierarchy primitives, so that, when primitive activation is imminent, we can already initiate the application action.

It is likely that building the model hierarchy is costly in terms of computing resources, which are limited on mobile devices. A possible solution is to offload the calculation to a back end machine, but even then computational complexity remains an issue. Another restriction on mobile devices is battery power, which may limit the time that sensors can be active.

6 Summary

We have provided an outline for building a fully decentralized probabilistic world model using the sensors available on the personal mobile devices of a community of participants. We discussed open issues in building, maintaining and the opportunistic distribution of such a model, using the principle of compositional hierarchies as an example.

7 Acknowledgement

This research is funded in part by the Haggel project under the EU grant IST-4-027918 and the SOCIALNETS project under the EU grant 217141.



8 References

- [1] C. Diot et al., The Huggle Project, <http://www.huggleproject.org/>, Retrieved on 2008-08-08.
- [2] D.Gavidia and M. van Steen, Enforcing Data Integrity in Very Large Ad Hoc Networks, In Proceedings of the 8th International Conference on Mobile Data Management (MDM'07), May 2007
- [3] R. Harle and A. Hopper, Dynamic World Models from Ray-tracing, In Proceedings of the 2nd International Conference on Pervasive Computing and Communications (PerCom'04), March 2004
- [4] P. Hui, J. Crowcroft, and E. Yoneki. BUBBLE Rap: Social Based Forwarding in Delay Tolerant Networks. In Proceedings of the 9th Intl. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'08), May 2008
- [5] P. Hui, E. Yoneki, et al., Distributed community detection in delay tolerant networks, In Proceedings of the 2nd International Workshop on Mobility in the Evolving Internet Architecture (MobiArch'07), August 2007
- [6] V. Kostakos, T. Nicolai, E. Yoneki, et al. Understanding and Measuring the Urban Pervasive Infrastructure". Journal of Personal and Ubiquitous Computing, pp.1617-4909, Springer, 2008.
- [7] A. Peddemors, H. Eertink, and I. Niemegeers, Density Estimation for Out-of-Range Events on Personal Mobile Devices, In Proceedings of the 1st International Workshop on Mobility Models for Networking Research (MobilityModels'08), May 2008
- [8] K. Pflieger, On-Line Learning of Predictive Compositional Hierarchies, Ph.D. Dissertation, Computer Department, Stanford University, June 2002
- [9] S. Saria, U. Nodelman, and D. Koller, Reasoning at the Right Time Granularity, In Proceedings of the 21st Conference on Uncertainty in AI (UAI'05), July 2005
- [10] S. Thrun, W. Burgard, and D. Fox, Probabilistic Robotics, The MIT Press, Sept. 2005
- [11] J. Utans, Learning in Compositional Hierarchies: Inducing the Structure of Objects from Data, Advances in Neural Information Processing Systems 6, 1994