10

Scalable Sensor Fusion for Motion Localization in Large RF Sensing Networks

Fetze Pijlman[1,2]

[1]Signify, The Netherlands
[2]Eindhoven University of Technology, The Netherlands

Abstract

RF sensing in wireless communication networks is a novel approach for motion detection, but it faces challenges in accurately localising motion which is crucial for confinement in lighting control use cases. A probabilistic model enables motion localisation through sensor fusion. However, in probabilistic models the posterior estimations do not scale well with large networks as likelihoods of all possible system states need to be computed. It will be demonstrated that variational Bayesian techniques offer attractive approximations to the posteriors where the approximations require computational resources that scale with the number of nodes. This method is of general interest to large networks as it models nonlocal effects through localised updates.

Keywords: RF sensing, sensor fusion, probabilistic models, variational inference, confinement.

10.1 Motivation

In a smart lighting system, one can control the lighting by sensing motion through the wirelessly connected lights (nodes) which is also known as RF sensing. For some background on RF sensing the reader is referred to Liu [1] and Wu [2]. Fluctuations in RSSI values, Received Signal Strength

Indicator, between nodes are strong indicators of nearby motion. For lighting control, the system should be able to detect human arm motions and steps as specified by the NEN norms [3]. As major motions generally lead to large RSSI fluctuations, it becomes a challenge for lighting control to sense minor motions within a room while not being sensitive to major motions just outside the room. An example situation is given in Figure 10.1.

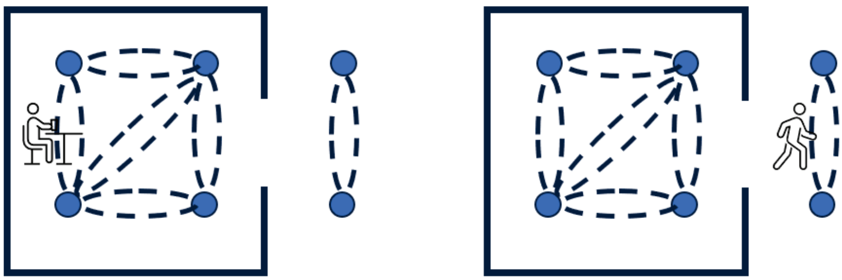

Figure 10.1 Downlights and indicated node pairs that are monitored for RSSI fluctuations. On the left a person working in an office and on the right a person walking on a corridor.

Probabilistic hidden Markov models are popular for motion sensing for lighting control. Typically, parameters such as motion rate during presence, average duration of presence, and probability of entering a room are introduced (see for example Papatsimpa [4]). These parameters can be used for propagating states and updating them with observations.

For modelling the signal variations due to motions in the surroundings, it is natural to extend the probabilistic models by modelling crosstalk. Motion directly underneath a node pair is visible, but a fraction of said motion is also visible to nearby other node pairs. The introduction of crosstalk complicates Bayesian inference as the number of different presence states now scale with 2^N where N is the number of presence areas. This essentially blocks the application of Bayesian inference for large networks. Another hurdle for large networks is the limited bandwidth for communication. In a mesh ZigBee network, the communication budget is 30-50 bytes per second per node, which limits the network size to about 50 nodes.

The goal of this paper is to apply a probabilistic model for motion/presence localisation by monitoring fluctuations in RSSI values between nearby node pairs. Using a variational Bayesian method, one can approximate posteriors by maximising the so-called free energy. As we will see, the maximisation is an iterative process that only involves local interactions. This leads to a scalable approach in which the computational power and memory scale with the number of nodes. Moreover, the calculations can be distributed over the nodes themselves, enabling a truly decentralised approach.

10.2 Spensor Fusion via a Probabilistic Model

For localising motion, a probabilistic model will be constructed (see Bishop [5], Murphy [6] and Morey [7] for an introduction to probabilistic models). The probabilistic model has the fluctuation of RSSI values of the node pairs as observational data *X*. By monitoring the fluctuations in RSSI values of a node pair one effectively obtains a motion sensor. In the remainder of the article node pairs will be referred to as sensors.

The probabilistic model contains hidden states of the physics in the various areas. These hidden states describe at each time instant *t* for each area: the presence $s(t)$ and the motion $m\,(t)$. In addition, at each time instant the visibility of a motion from an area *i* to a particular sensor *j* is described by $C_{ij}(t)$.

In order to simplify the calculation, one would like to make the Markov assumption which influences the choice for the value of the time-step. A simple model can be obtained by choosing the time-step to be larger than 1 second in which case motion at time *t* only depends on the presence state at *t* but not on a previous motion state. As the RSSI fluctuations *X(t)* only depend on the ***m(t)*** and the probability that a motion is visible *C* to a sensor, one can write the distribution as

$$
\begin{aligned}
&P\left(X\left(t\right),\boldsymbol{m}\left(t\right),\boldsymbol{s}\left(t\right),\boldsymbol{s}\left(t-1\right)|C(t)\right)\\
&= P(X(t)|\boldsymbol{m}(t),C(t))\prod_i P\left(\boldsymbol{m_i}\left(t\right)\mid\boldsymbol{s_i}\left(t\right)\right)P\left(\boldsymbol{s_i}\left(t\right)\mid\boldsymbol{s_i}\left(t-1\right)\right)\\
&\qquad P\left(\boldsymbol{s_i}\left(t-1\right)\right).
\end{aligned}
$$

Note that in the equation above it was assumed for simplicity that the presences in each area are uncorrelated. In Figure 10.2 an example of the interactions is given at a time instant.

With the Markov assumption, the posterior can now be determined by Bayesian inference

$$
\begin{aligned}
&P\left(\boldsymbol{s}\left(t\right),\boldsymbol{m}\left(t\right),C\left(t\right)|X\left(:t\right)\right) \qquad (10.1)\\
&= \frac{P\left(X\left(t\right)|\boldsymbol{s}\left(t\right),\boldsymbol{m}\left(t\right),C\left(t\right)\right)}{P\left(X\left(t\right)\right)}P\left(\mathbf{s}\left(t\right),\boldsymbol{m}\left(t\right),C\left(t\right)|X\left(:t-1\right)\right),
\end{aligned}
$$

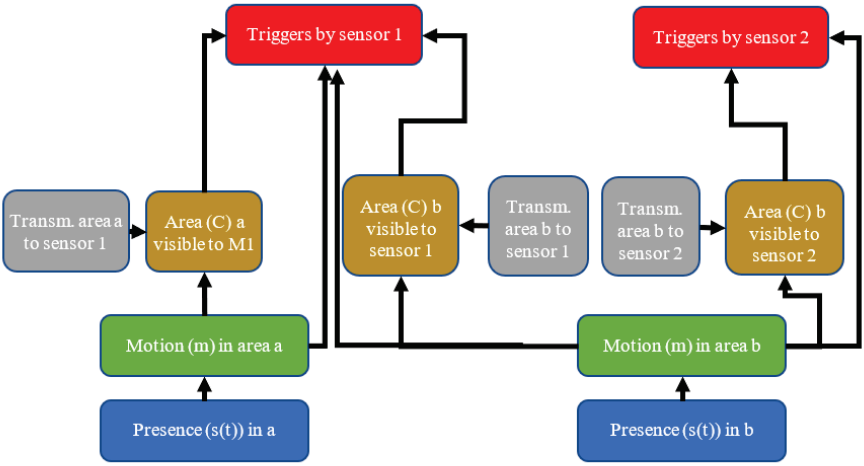

Figure 10.2 Bayesian network at a time instant showing the dependence between the various states and the sensor observations.

where X (: t) denotes all RSSI fluctuations up and including t, and X(t) denotes RSSI fluctuations at time t. Assuming the independence of presences in areas one finds for each area *i*

$$P\left(\boldsymbol{s_i}\left(t\right), \boldsymbol{m_i}\left(t\right) | X\left(: t{-}1\right)\right) = P\left(\boldsymbol{m_i}\left(t\right) \mid \boldsymbol{s_i}\left(t\right)\right) P\left(\boldsymbol{s_i}\left(t\right) | X\left(: t{-}1\right)\right)$$
$$P\left(\boldsymbol{s_i}\left(t\right) | X\left(: t{-}1\right)\right) = \sum_{\boldsymbol{s_i}(t-1)} P\left(\boldsymbol{s_i}\left(t\right) \mid \boldsymbol{s_i}\left(t-1\right)\right) P\left(\boldsymbol{s_i}\left(t-1\right) | X\left(: t{-}1\right)\right).$$

Note that the propagation of presence states involves a transition matrix $P\left(\boldsymbol{s_i}\left(t\right) \mid \boldsymbol{s_i}\left(t-1\right)\right)$.

The posterior in Equation 10.1 will be computed at each time-step using variational methods (see also Attias [8] and Jordan [9]). By maximising the free energy one can approximate the posterior $P\left(\boldsymbol{s}\left(t\right), \boldsymbol{m}\left(t\right), C\left(t\right) | X\left(t\right)\right)$ by a probabilitty distribution q in terms of the KL-divergence $D_{KL}(P||q)$. This yields

$$P\left(\boldsymbol{s}\left(t\right), \boldsymbol{m}\left(t\right), C\left(t\right) | X\left(t\right)\right)$$
$$\approx \underset{q}{\operatorname{argmax}} \sum_{\mathbf{s}(t), \mathbf{m}(t), C(t)} q\left(\boldsymbol{s}\left(t\right), \boldsymbol{m}\left(t\right), C\left(t\right)\right) \left[\log \frac{P\left(X\left(t\right), \boldsymbol{s}\left(t\right), \boldsymbol{m}\left(t\right), C\left(t\right)\right)}{q\left(\boldsymbol{s}\left(t\right), \boldsymbol{m}\left(t\right), C\left(t\right)\right)}\right]. \tag{10.2}$$

For notational simplicity the explicit time-dependence will be dropped in the remainder.

In order to solve the equation above one often makes additional assumptions on the probability distribution q. However, the mean-field approximation such as $q(\boldsymbol{s}(t), \boldsymbol{m}(t), C(t)) = q(\boldsymbol{s}(t))\, q(\boldsymbol{m}(t))\, q(C(t))$ is blocked as probability for motion during absence is zero. Zero probabilities lead to problems with the logarithms. Instead, we approximate q by

$$q(\mathbf{s}, \mathbf{m}, C) = \prod_{i,j} q(C_{ij}|m_i)\; q(m_i|s_i)\; q(s_i), \tag{10.3}$$

where $q(m_i = 1|s_i = 0) = 0$ and j refers to the sensor index. For notational simplicity the difference between the q's such as $q(\mathrm{s}_1)$ and $q(s_2)$ has been made implicit. Note that this approximation contains more distributions compared to the mean-field approximation; $q(m_i|s_i = 0)$ has no relation with $q(m_i|s_i = 1)$.

10.3 Update Equations

Equation 10.2 can be solved by iteratively maximizing the lower bound evidence. With each update for one of the $q's$ the lower bound evidence gets increased. Please see Figure 10.3 for a subset of the network that is relevant for determining the states. Please note that $m_{o_{ij}}$ is the visible motion towards sensor j excluding any motion from m_i.

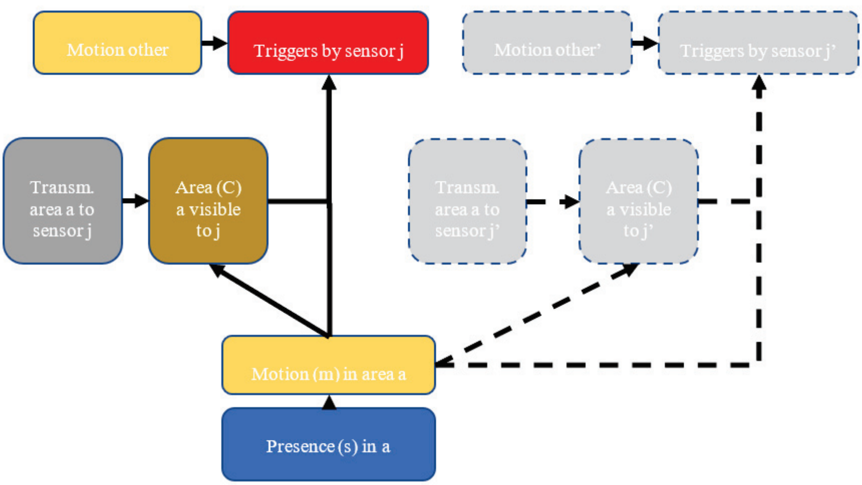

Figure 10.3 Isolated part of the network that is relevant for determining the states.

In the following subsections the update equations for the various factors will be determined by plugging Equation 10.3 in Equation 10.2. In the derivation the following relation will be often employed

$$\mathrm{argmax}_{q(a)} \sum_a q(a)\left(\log P(X,a) - \log q(a)\right) = \frac{P(X|a)P(a)}{\sum_{a'} P\left(X|a'\right)P\left(a'\right)}. \tag{10.4}$$

The relation can be proven by using the method of Lagrange multipliers and treating $\sum_a q(a) = 1$ as a constraint (see also Arfken [10]).

It is instructive to work out the update equation in case a mean-field approximation would apply. Consider the example in Figure 10.4 and write $q(s,t,u)$ as $q(s)\,q(t)\,q(u)$. In this example the update equation for $q(t)$ would be

$$\begin{aligned} q(t) &= \mathrm{argmax}_{q(t)} \sum_{s,t,u} q(s)\,q(t)\,q(u) \log \frac{P(X,s,t,u)}{q(s)q(t)q(u)} \\ &= \mathrm{argmax}_{q(t)} \sum_{s,t,u} q(s)\,q(t)\,q(u) \log \frac{P(s|t)P(t|u))}{q(s)q(t)q(u)} \\ &= \mathrm{argmax}_{q(t)} \sum_t q(t) \\ &\qquad \left[\log\left(\prod_s P(s|t)^{q(s)} \prod_u P(t|u)^{q(u)}\right) - \log q(t)\right] \\ &= \frac{\prod_s P(s|t)^{q(s)} \prod_u P(t|u)^{q(u)}}{\sum_{t'} \prod_s P\left(s|t'\right)^{q(s)} \prod_u P(t'|u)^{q(u)}} \end{aligned} \tag{10.5}$$

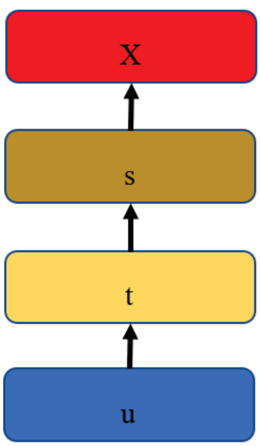

Figure 10.4 Example network with a mean-field assumption.

In case t only takes two values (0,1) and t is the opposite value of t then

$$\mathrm{q}\,(\mathrm{t}) = \mathrm{S}\left[\log\frac{q(t)}{1-q(t)}\right] = \mathrm{S}\left[\sum_{\mathrm{s}}\mathrm{q}\,(\mathrm{s})\log\frac{P(s|t)}{P(s|t)} + \sum_{\mathrm{u}}\mathrm{q}\,(\mathrm{u})\log\frac{P(t|u)}{P(t|u)}\right] \tag{10.6}$$

where S is the Sigmoid function.

Therefore, in order to update $q\,(t)$ one only need to know the approximated posteriors of its direct neighbours. Moreover, Equation 10.5 is just a Bayesian inference in which the prior and likelihood are products of powers by how much the surroundings appear which is according to expectation. In the following subsection the update equations will be derived for the various $q's$. In Figure 10.3 a part of the network is visualized which is useful when deriving the update equations.

10.3.1 Update equation for $q\,(C_{ij}|m_i\ =\ 0)$

Plugging Equation 10.3 in Equation 10.2 one finds

$$\begin{aligned} q\,(C_{ij}|m_i\ =\ 0)\ &= \underset{q(C_{ij}|m_i=0)}{\operatorname{argmax}} \sum_{s_i,m_i,m_{oij},C_{ij}} q\left(m_{oij}\right) q\left(s_i\right) q\left(m_i|s_i\right) \\ &\qquad q(C_{ij}|m_i)\log\frac{P\left(X_j,m_{oij},s_i,m_i,C_{ij}\right)}{q\left(m_{oij}\right)q(s_i)q(m_i|s_i)q(C_{ij}|m_i)} \\ &= \underset{q(C_{ij}|m_i=0)}{\operatorname{argmax}} \sum_{s_i,m_{oij},C_{ij}} q(m_{oij})q\left(s_i\right) q\left(m_i=0|s_i\right) \\ &\qquad q\left(C_{ij}|m_i{=}0\right)\log\frac{P\left(X_j,m_{oij},s_i,m_i{=}0|C_{ij}\right)P(C_{ij})}{q(m_{oij})q(s_i)q(m_i{=}0|s_i)q(C_{ij}|m_i{=}0)} \\ &= \underset{q(C_{ij}|m_i=0)}{\operatorname{argmax}} \sum_{s_i,C_{ij}} q\left(s_i\right),q\left(m_i{=}0|s_i\right) q\left(C_{ij}|m_i{=}0\right) \\ &\qquad \log\frac{P(C_{ij})}{q(C_{ij}|m_i{=}0)} \\ &= P\left(C_{ij}\right) \end{aligned} \tag{10.7}$$

The result above is according to expectation as one cannot estimate transparency/visibility in absence of a source m.

10.3.2 Update equation for $q\,(C_{ij}|m_i\ =\ 1)$

Plugging Equation 10.3 in Equation 10.2 one finds

$$\begin{aligned} q\,(C_{ij}|m_i\ =\ 1)\ &= \underset{q(C_{ij}|m_i=1)}{\operatorname{argmax}} \\ &\qquad \sum_{s_i,m_i,m_{oij},C_{ij}} q(m_{oij})q\left(s_i\right),q\left(m_i|s_i\right) q\left(C_{ij}|m_i\right) \end{aligned}$$

$$\begin{aligned}
&\log \frac{P\left(X_j, m_{oij}, s_i, m_i, C_{ij}\right)}{q(m_{oij}) q(s_i) q(m_i|s_i) q(C_{ij}|m_i)} \\
&= \underset{q(C_{ij}|m_i=1)}{\operatorname{argmax}} \sum_{m_{oij}, C_{ij}} q(m_{oij}) q\left(C_{ij}|m_i=1\right) \\
&\quad \log \frac{P\left(X_j, m_{oij}, m_i=1, s_i=1, C_{ij}\right)}{q(C_{ij}|m_i=1)} \\
&= \underset{q(C_{ij}|m_i=1)}{\operatorname{argmax}} \sum_{m_{oij}, C_{ij}} q(m_{oij}) q\left(C_{ij}|m_i=1\right) \\
&\quad \left[\log P\left(X_j|m_{oij}, m_i=1, C_{ij}\right) - \log \frac{q(C_{ij}|m_i=1)}{P(C_{ij}|m_i=1))}\right] \\
&= \frac{P(C_{ij}|m_i=1)\left(P\left(X_j|m_{oij}=0, m_i=1, C_{ij}\right)\right)^{q\left(m_{oij}=0\right)}}{\sum_{C'_{ij}} P(C'_{ij}|m_i=1)\left(P\left(X_j|m_{oij}=0, m_i=1, C'_{ij}\right)\right)^{q\left(m_{oij}=0\right)}}
\end{aligned} \tag{10.8}$$

In the last line terms involving $q\left(m_{oij} = 1\right)$ can be omitted as they do not lead to differentiation for the different values of C_{ij}. The result can be understood as a Bayesian inference where the amount of observation is determined by $q\left(m_{oij} = 0\right)$.

10.3.3 Update equation for $q\left(m|s = 1\right)$

Plugging Equation 10.3 in Equation 10.2 one finds

$$\begin{aligned}
&q\left(m_i|s_i = 1\right) \\
&= \max_{q(m_i|s_i=1)} \sum_{j, s_i, m_i, m_{oij}, C_{ij}} q\left(m_{oij}\right) q\left(s_i\right) q\left(m_i|s_i\right) q\left(C_{ij}|m_i\right) \\
&\quad \log \frac{P\left(X_j, m_{oij}, s_i, m_i, C_{ij}\right)}{q\left(m_{oij}\right) q(s_i) q(m_i|s_i) q(C_{ij}|m_i)} \\
&= \max_{q(m_i|s_i=1)} \sum_{j, m_i, m_{oij}, C_{ij}} q\left(m_{oij}\right) q\left(m_i|s_i=1\right) q\left(C_{ij}|m_i\right) \\
&\quad \log \frac{P\left(X_j, m_i, C_{ij}|m_{oij}, s_i=1\right)}{q(m_i|s_i=1) q(C_{ij}|m_i)} \\
&= \frac{\prod_{j, m_{oij}, C_{ij}} P(m_i|s_i=1)\left(\frac{P\left(C_{ij}|m_i\right)}{q\left(C_{ij}|m_i\right)} P\left(X_j|m_i, C_{ij}, m_{oij}\right)\right)^{q\left(C_{ij}|m_i\right) q\left(m_{oij}\right)}}{\sum_{m'_i} \prod_{j, m_{oij}, C_{ij}} P(m'_i|s_i=1)\left(\frac{P\left(C_{ij}|m'_i\right)}{q\left(C_{ij}|m'_i\right)} P\left(X_j|m'_i, C_{ij}, m_{oij}\right)\right)^{q\left(C_{ij}|m'_i\right) q\left(m_{oij}\right)}} \\
&= \frac{P(m_i|s_i=1) \prod_{j, C_{ij}} \left(\frac{P\left(C_{ij}|m_i\right)}{q\left(C_{ij}|m_i\right)}\right)^{q\left(C_{ij}|m_i\right)} \left(P\left(X_j|m_i, C_{ij}, m_{oij}\right)\right)^{q\left(m_{oij}=0\right)}}{\sum_{m'_i} P(m'_i|s_i=1) \prod_{j, C_{ij}} \left(\frac{P\left(C_{ij}|m'_i\right)}{q\left(C_{ij}|m'_i\right)}\right)^{q\left(C_{ij}|m'_i\right)} \left(P\left(X_j|m'_i, C_{ij}, m_{oij}\right)\right)^{q\left(m_{oij}=0\right)}}
\end{aligned} \tag{10.9}$$

The result above contains terms that originate from complexity $P\left(C_{ij}|m_i\right) / q\left(C_{ij}|m_i\right)$.

10.3.4 Update equation for $q(s)$

Plugging Equation 10.3 in Equation 10.2 one finds

$$
\begin{aligned}
q(s_i) =& arg \max_{q(s_i)} \sum_{j,s_i,m_i,m_{oij},C_{ij}} q\left(m_{oij}\right) q(s_i)\, q(m_i|s_i)\, q\left(C_{ij}|m_i\right) \\
& \log \frac{P\left(X_j, m_{oij}, s_i, m_i, C_{ij}\right)}{q\left(m_{oij}\right) q(s_i)\, q(m_i|s_i)\, q\left(C_{ij}|m_i\right)} \\
\propto & P(s_i) \prod_{j,m_i,m_{oij},C_{ij}} \\
& \left[\frac{P\left(X_j|m_{oij}, m_i, C_{ij}\right) P\left(m_{oij}\right) P(m_i|s_i)\, P\left(C_{ij}|m_i\right)}{q\left(m_{oij}\right) q(m_i|s_i)\, q\left(C_{ij}|m_i\right)}\right]^{q\left(m_{oij}\right) q(m_i|s_i) q(C_{ij}|m_i)} \\
\propto & P(s_i) \prod_{m_i,j} \left(P\left(X_j|m_{oij}{=}0, m_i, C_{ij}{=}1\right)\right)^{q\left(m_{oij}=0\right) q(m_i|s_i) q(C_{ij}=1|m_i)} \\
& \left(\frac{P(m_i|s_i)}{q(m_i|s_i)}\right)^{q(m_i|s_i)} \times \prod_{C_{ij}} \left(\frac{P\left(C_{ij}|m_i\right)}{q\left(C_{ij}|m_i\right)}\right)^{q(C_{ij}|m_i)}.
\end{aligned} \tag{10.10}
$$

10.4 Conclusions and Discussion

In this article, a probabilistic model is developed for localising motion in a wireless IoT network. A hidden Markov model is employed to propagate presence states over time using a transition matrix, forming a new prior for presence at the next time instant. At each time instant, a variational Bayesian approach was utilised to translate sensor data (RSSI values) into posterior presence and motion states, accounting for potential crosstalk. A first preliminary result is given in Figure 10.5.

The derived equations do not need to be applied to states in a particular order as each order will improve the lower bound evidence. This enables a distributed approach in which nodes apply asynchronous updates by taking information from their neighboring nodes. However, the iterations do increase network traffic, and it remains to be seen whether the network can handle the increase in communication.

The number of iterations needed for convergence will generally depend on how sensitive RSSI fluctuations of node pairs are to motion in the vicinity as the probabilistic model handles nonlocal effects via iterations involving

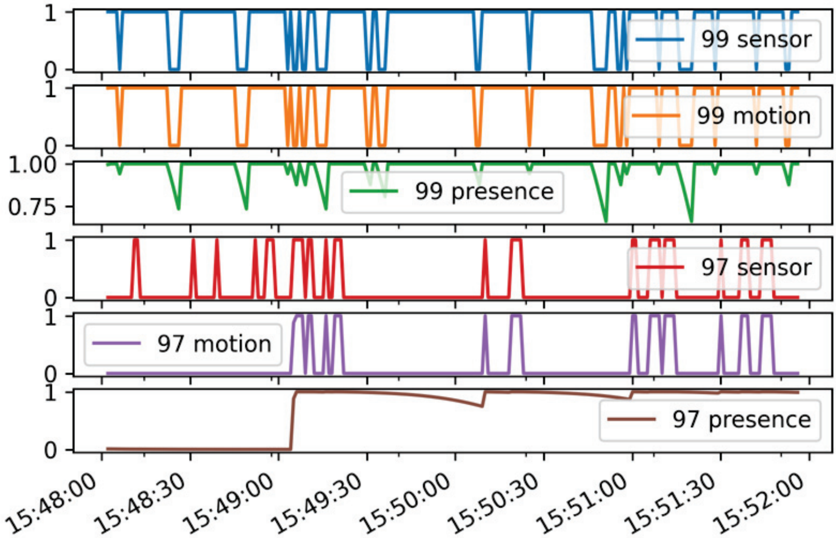

Figure 10.5 First results for two areas being named 99 and 97. The model parameters were chosen such that area 99 is a corridor while area 97 is a meeting room. In the figure one can see that isolated sensor events of 97 before 15:49 do not lead to presence in area 97 as it is more likely they originated from area 99. After probablity for presence is high in area 97 then isolated sensor events of 97 at 15:50:15 do lead to an increase in presence in area 97.

local updates. Future experimental studies are necessary to test the model for convergence and to quantify the improvements in motion localisation.

Centralized or distributed, the number of iterations needed for convergence may be reduced when considering techniques such as gradient descent. Using gradient descent one updates a group of states by a small amount. Iteratively updating (see also Kuusela [11]) may lead to faster convergence compared to iteratively updating each state.

Although the use of variational Bayesian techniques led to simple update equations, a problem with the approach is that the obtained posteriors are often narrower than what would be justified. This is an effect of minimizing the reversed KL-divergence. Obtaining over-confident posteriors is especially problematic for crosstalk parameters C_{ij} as it highly influences future updates. Therefore, it is recommended to use the whole time-series when updating model parameters with variational Bayes. As an alternative one may consider in future studies expectation propagation as proposed by Minka [12] in which sequential updates should be feasible. Similarly, it would be of interest to benchmark the obtained iterative algorithm with message passing algorithms such as RxInfer [13].

Acknowledgements

Jean-Paul Linnartz, Michel Wijker, Bert de Vries, Marco Haverlag, and Peet van Tooren are acknowledged for fruitful discussions.

EdgeAI "Edge AI Technologies for Optimised Performance Embedded Processing" project has received funding from Key Digital Technologies Joint Undertaking (KDT JU) under grant agreement No 101097300. The KDT JU receives support from the European Union's Horizon Europe research and innovation program and Austria, Belgium, France, Greece, Italy, Latvia, Luxembourg, Netherlands, Norway.

References

[1] H. Liu, H. Darabi, P. Banerjee, J. Liu, "Survey of Wireless Indoor Positioning Techniques and Systems", IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS, 2007.

[2] Z-H. Wu, Y. Han, Y. Chen, K. J. R. Liu, "A Time-Reversal Paradigm for Indoor Positioning System", IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, 2015.

[3] "Methods of measurement and declaration of the detection range of detectors - Passive infrared detectors for major and minor motion detection", NEN-EN-IEC 63180, 2020.

[4] C. Papatsimpa, J.P. Linnartz, "Distributed Fusion of Sensor Data in a Constrained Wireless Network", MDPI Sensors, (2019).

[5] C. Bishop, "Pattern Recognition and Machine Learning", Springer, 2006.

[6] K.P. Murphy, "Machine learning - a probabilistic perspective". *The MIT Press*, 2012.

[7] R.D. Morey, J.W. Romeijn, J.N. Rouder, "The philosophy of bayes factors and the quantification of statistical evidence". *Journal of Mathematical Psychology*, 2016.

[8] H. Attias. "A variational bayesian framework for graphical models". *neurIPS* 1999.

[9] M.I. Jordan, Z. Ghahramani, T.S. Jaakola, L.K. Saul, "An introduction to variational methods for graphical models". *Springer* 1999.

[10] G.B. Arfken, H.J. Weber, "Mathematical Methods for Physicists." *Academic Press*, 1995.
[11] M. Kuusela, T. Raiko, A. Honkela, J. Karhunen, "A Gradient-Based Algorithm Competitive with Variational Bayesian EM for Mixture of Gaussians". Proceedings of International Joint Conference on Neural Networks 2009.
[12] T. Minka, "Expectation Propagation for approximate Bayesian inference", Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, 2001.
[13] D.V. Bagaev, A. Podusenko, B. de Vries, "RxInfer: A Julia package for reactive real-time Bayesian inference". The Journal of Open Source Software, 2023.