Albert Hein*, Florian Grützmacher, Christian Haubelt and Thomas Kirste

Fast care – real-time sensor data analysis framework for intelligent assistance systems

Distributed inference algorithms on multi-sensor platforms for medical assistance

Abstract: Main target of *fast care* is the development of a real-time capable sensor data analysis framework for intelligent assistive systems in the field of Ambient Assisted Living, eHealth, Tele Rehabilitation, and Tele Care. The aim is to provide a medically valid integrated situation model based on a distributed, ad-hoc connected, energy-efficient sensor infrastructure suitable for daily use. The integrated situation model combining physiological, cognitive, and kinematic information about the patient is grounded on the intelligent fusion of heterogeneous sensor data on different levels. The model can serve as a tool for quickly identifying risk and hazards as well as enable medical assistance systems to autonomously intervene in real-time and actively give telemedical feedback.

Keywords: multi-sensor platform, situation model, inference, assistance system, home care, AAL.

1 Motivation

One important research objective during the last years is how to empower people in need of care (e.g. in rehabilitation, the elderly, people with cognitive impairments) to stay at their own home as long as possible. Unfortunately most medical measuring technology for diagnosis, monitoring, and risk stratification cannot be simply transferred to the domestic area. Assistance systems in the field of Ambient Assisted Living and medical care have to detect situations which require assistive intervention in real-time. In addition many of these situations can only be detected with a combination of multiple sensing modalities (e.g. heart rate, kinematics, location). The major challenge within the *fast care* project is to provide an integrated picture of the current situation, which requires knowledge about the environmental state, the user's current and possible future activities, and his/her goals. In addition, for estimating this kind of information mostly only heterogeneous, error-prone, and noisy sensor data is available. Because of the complex technical requirements an interdisciplinary cooperation is needed and *fast care* is designed as a compound project of multiple partners. In the project consortium the academic partners are responsible for developing, evaluating, and optimizing new technologies while the industrial partners work on novel technical concepts for practical product solutions.

2 Project targets

The central task at the University of Rostock is to create an integrated probabilistic situation model as well as the development and prototypical realization of real-time capable distributed inference methods. The integrated situation model here not only contains kinematic parameters, physiological measurements, or activity labels, but – depending on the requirements of the application – has to mirror a preferably holistic picture of a person's state within his/her environment (Fig 1). It will be represented by a probabilistic generative state space model (state of the environment, location, motion state, physiological parameters, cognitive and emotional state, activity goals) and can be realized for example in form of a Dynamic Bayesian Network.

_____
*Corresponding author: Albert Hein: Institute of Computer Science, University of Rostock, 18059 Rostock, Germany, e-mail: albert.hein@uni-rostock.de
Florian Grützmacher, Christian Haubelt: Institute of Applied Microelectronics and Computer Engineering, University of Rostock, 18051 Rostock, Germany, e-mail: florian.gruetzmacher2@uni-rostock.de, christian.haubelt@uni-rostock.de
Thomas Kirste: Institute of Computer Science, University of Rostock, 18059 Rostock, Germany, e-mail: thomas.kirste@uni-rostock.de
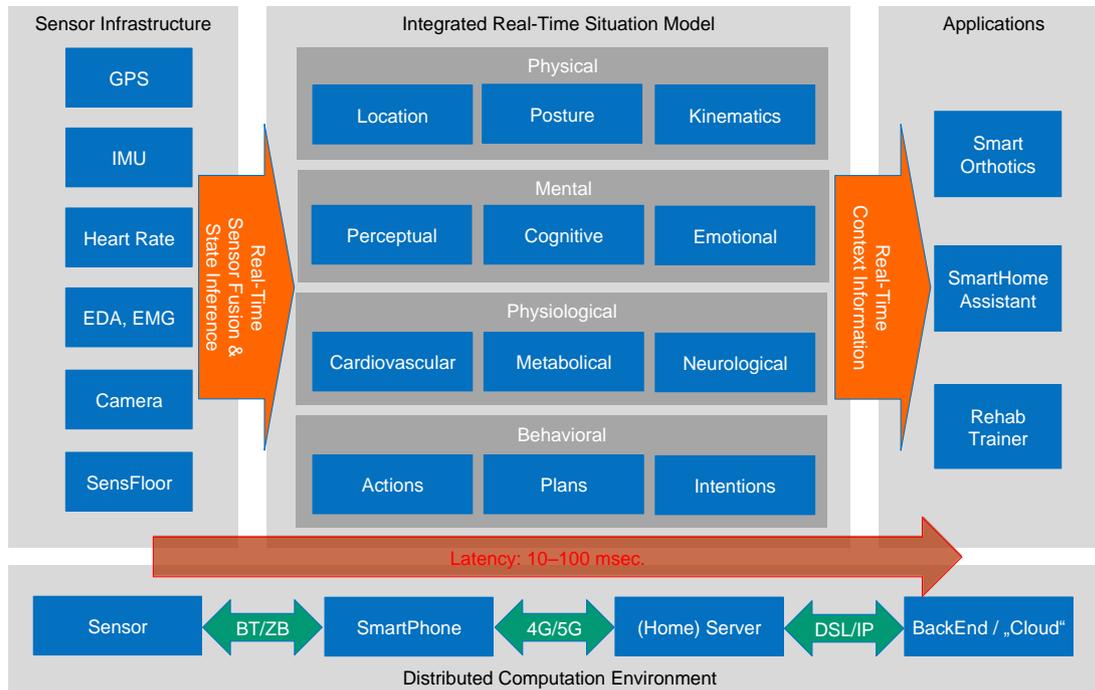
Figure 1: Overview of the fast care system consisting of a heterogeneous sensor infrastructure and an integrated situation model allowing assistive real-time intervention.

Because of the real-time requirement of the application it is crucial to efficiently make use of the underlying computing platform as much as possible, which will be further described in Section 3.1.

In summary the following targets can be identified:
– Requirement analysis of hard- and software for distributed real-time capable inference methods
– Probabilistic generative system model
– Observation model for heterogeneous noisy sensor data
– Energy reduction and synchronization approaches for heterogeneous on-body sensor networks
– Real-time capable distribution of inference algorithms
– Example application with real-time feedback for people with cognitive impairments

Under the term *activity recognition* multiple approaches can be found in the literature, which derive the current activity from motion sensor data (e.g. [1], for an overview see [2,3,4]). Therefor successfully machine learning techniques like Support Vector Machines and Decision Trees are used. While these approaches show high accuracies in detecting of activity labels (like sit, walk, lie), they are inherently unable to produce a comprehensive picture of the overall situation. Methods like Dynamic Bayesian Networks beyond that may contain a model of the environment and represent temporal

dependencies and allow the prediction of potential future events [5].

In parallel to the research on recognizing activities from noisy sensor data, approaches have been investigated which contain a complete description of the environment, possible activities, and goals and therefore can infer a comprehensive picture of the current situation [6,7]. These methods are capable of determining both the current situation as well as possible courses of action and goals, and are not dependent on direct and error-free observations, which cannot be provided by real-world sensor hardware. By the combination of complex descriptions of the environment and permitted courses of activities with probabilistic inference methods this gap could be overcome in some works [8,9,10] using position data. In [11] it was shown furthermore, that, using Computational Causal Behavior Models (CCBM), this is also possible with wearable sensors.

3 Hardware architecture

In typical activity recognition scenarios, data from multiple sensors, e.g., accelerometers, gyroscopes, and magnetometer, located at different parts on the human body are processed to infer the current activity of the subject. However, such a setup leads to a high amount of sensor data which needs to be gathered, communicated, preprocessed and finally used as

input for inference algorithms. In order to achieve reasonable processing times, the amount of data needs to be reduced as early as possible, in order to reduce communication overhead and to accelerate the preprocessing of gathered data from multiple sensors.

We propose a system architecture consisting of several sensor nodes attached to different parts of the human body. Those sensor nodes are equipped with a sensor, a micro controller acquiring the sensors data and a wireless communication interfaces to send the data to a central computing platform like a smartphone for further processing. Since sensors-hubs already entered the market, which combine different inertial sensors on a chip together with processing capabilities for sensor data fusion [12], such sensor systems are promising for the proposed architecture. Such a setup leads to a heterogeneous multi-processor architecture with different levels of processing power distributed over the human body. An example of such an architecture is sketched in Figure 2.

Algorithm distribution:
The sensor fusion capabilities offer a data reduction and preprocessing at a very first processing stage. The microcontroller on the sensor node would be a second stage, at which preprocessing can be done. Similar approaches calculating a selection of features from accelerometer signals on a sensor node have already been introduced [13].
We will extend those approaches by generalizing it to an evaluation of how to efficiently distribute inference algorithms on heterogeneous sensor architectures in a generic way. This evaluation will be based on earlier work, which focused on how to efficiently distribute gesture recognition algorithms on homogeneous multi-processor architectures in a generic way [14].

These approaches, which mainly focused on multi-processor architectures with the same processing power of all cores residing on a single chip, will be extended to efficiently use heterogeneous architectures with different levels of processing power. By doing so, also much higher communication delays have to be considered due to distribution of multi-processor architectures over multiple sensor nodes, which communicate over wireless connections.

Power management:
Using wireless sensor nodes, energy consumption plays a crucial role, as it directly affects battery life and therefore the usability of the entire system. By reducing the amounts of data as explained in 3.1, we can already reduce the power consumption of each sensor node as well as the power consumption of the central computing platform.

However, in addition we want to make use of the power modes of the individual processing cores. By detecting activities, which require less information to be processed in order to provide the full service of the activity recognition system, the calculation of some features or even the gathering of sensor data at all can be turned off. As a consequence, the processing time is reduced significantly. A reduction of processing time even leads to longer periods a micro controller can stay in low power modes, which reduces energy consumption further. This idea can be extended to completely turn off certain sensors which we don't need in selected scenarios detected by the activity recognition.

Synchronization:
Due to small differences in the clock rate between sensor, also known as clock drifts, it is possible that sensor data requested at a certain rate actually arrive in slightly different rates on the gathering host platform. In order to provide proper features for activity recognition the data gathered from multiple sensors needs to be synchronized before features are extracted. We will evaluate, which sensor synchronisation techniques are applicable in our scenario, regarding accuracy of the synchronisation and the delays caused by such synchronisation techniques.
A possible approach would be the timestamp reconstruction described in [15], performed centrally on the host platform.

4 Situation model

The situation model itself consists of three dependent parts: The system model describing the probabilistic structure, dependencies and state transitions; the observation model for updating the internal representation with measurements; and the inference algorithm.

System model:
For creating the probabilistic generative system model for the overall situation, at first suitable probabilistic modeling concepts have to be analyzed and developed. After that, relevant parameters have to be represented as random variables and their correlation structure identified, which can then be transferred into a Dynamic Bayesian Network. For this conditional probability distributions and parameters have to be estimated using training data. The result is the algorithmic description and prototypical implementation of the system model.
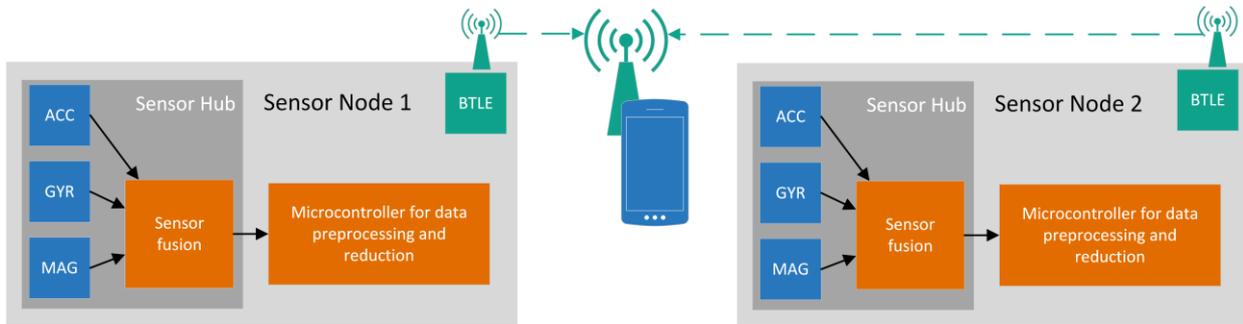
Figure 2: Example architecture of an on-body sensor network for activity recognition resembling a heterogeneous multi-processing architecture with data fusion and processing capabilities at different levels.

Observation model:

To develop the observation model, at first prerecorded data from the sensor nodes has to be analyzed with statistical signal processing methods. The sensor data stream will then be segmented using change point detection methods and meaningful features can be calculated. Now suitable parametric generative (e.g. Gaussian Mixture) as well as discriminative methods (e.g. Support Vector Machines) have to be evaluated for disambiguating the anticipated target parameters of the system model. The observation model will then be linked probabilistically via distributions to the system model.

Distributed inference:

Due to the complex nature of the anticipated model, its potentially infinite state space, and the real-time requirements of the application mainly approximate inference methods will be investigated. At first a monolithic particle filter based inference algorithm for non-linear, non-gaussian, and continuous state spaces will be implemented. Single partial models can now be identified which are suitable for more efficient specialized inference methods, which can then be integrated into a Rao-Blackwellized particle filter. Finally components of the inference algorithm and the observation model can be outsourced and distributed for pre-calculation on the sensor nodes.

5 Results and outlook

Based on a requirements analysis we deployed a first setup consisting of 5 sensor nodes each equipped with a Bosch BHI160 sensor hub [12], a microcontroller and a Bluetooth Low Energy interface. The Sensor hub performs a first sensor data fusion. As a central host platform we use a smartphone, which gathers the data from all sensor nodes. Our current research focuses on the systematic distribution of inference algorithms used in [11] on the deployed hardware.

References

[1] L. Bao, S. Intille, Activity recognition from user-annotated acceleration data. In Pervasive, volume 3001 of Lecture Notes in Computer Science, pages 1–17. Springer, 2004.

[2] J. K. Aggarwal, M. S. Ryoo. Human activity analysis: A review. ACM Computing Surveys, 43(3):16:1–16:43, April 2011.

[3] Liming Chen, J. Hoey, C. D. Nugent, D. J. Cook, and Zhiwen Yu. Sensor-based activity recognition. IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews, 42(6):790–808, 2012.

[4] The Anh Han and Luís Moniz Pereira. State-of-the-art of intention recognition and its use in decision making. AI Communications, 26(2):237–246, 2013.

[5] A. Hein, C. Burghardt, M. Giersich, T. Kirste. Model-based Inference Techniques for detecting High-Level Team Intentions. Björn Gottfried, Hamid Aghajan (eds.). Behaviour Monitoring and Interpretation - BMI: Smart Environments. In Behaviour Monitoring and Interpretation - BMI: Smart Environments, IOS Press, Amsterdam, 2009

[6] C. L. Baker, R. Saxe, and J. B. Tenenbaum, Action understanding as inverse planning, Cognition, vol. 113, no. 3, pp. 329–349, 2009.

[7] M. Ramírez and H. Geffner, Goal recognition over POMDPs: inferring the intention of a POMDP agent, in Proceedings of the 22nd IJCAI, 2011, pp. 2009–2014.

[8] H. H. Bui, S. Venkatesh, and G. A. W. West, Policy Recognition in the Abstract Hidden Markov Model, Journal of Artificial Intelligence Research (JAIR), vol. 17, pp. 451–499, 2002.

[9] L. Liao, D. J. Patterson, D. Fox, and H. Kautz, Learning and Inferring Transportation Routines, AI, vol. 171, no. 5–6, pp. 311–331, 2007.

[10] A. Sadilek and H. Kautz, Location-based Reasoning about Complex Multi-Agent Behavior, Journal of Artificial Intelligence Research, vol. 43, no. 1, pp. 87–133, Jan. 2012.

[11] F. Krüger, M. Nyolt, K. Yordanova, A. Hein, and T. Kirste. Computational State Space Models for Activity and Intention Recognition. A Feasibility Study. PLOS ONE, 9(11):e109381, Nov. 2014.

[12] Bosch Sensortec (2017m March), "BHI160 / BHI160B," BST-BHI160(B)-DS000-01, https://ae-bst.resource.bosch.com/media/_tech/media/datasheets/BST-BHI160B-DS000-01-Datasheet.pdf [Revised May. 2017]

[13] Van Laerhoven, K., & Aronsen, A. K. (2007, June). Memorizing what you did last week: Towards detailed actigraphy with a wearable sensor. In Distributed Computing Systems Workshops, 2007. ICDCSW'07. 27th International Conference on (pp. 47-47). IEEE.

[14] Grützmacher, F., Beichler, B., Haubelt, C., & Theelen, B. (2016, April). Dataflow-based modeling and performance analysis for online gesture recognition. In *Modelling, Analysis, and Control of Complex CPS (CPS Data), 2016 2nd International Workshop on* (pp. 1-8). IEEE.

[15] Stieber, S., Dorsch, R., & Haubelt, C. (2016, April). Accurate Sample Time Reconstruction for Sensor Data Synchronization. In *International Conference on Architecture of Computing Systems* (pp. 185-196). Springer International Publishing