Prof. Dr. Kristof Van Laerhoven | Eingebettete Systeme | kvanlaerhoven@acm.org | Tel. 0173 801 7770

An den Dekan der Fakultät IV
der Universität Siegen,
57068 Siegen

Freiburg, 29.04.2015

Bewerbung auf die Professur für Ubiquitous Computing (W 3)

Sehr geehrter Herr Prof. Dr. Dr. h. c. Ullrich Pietsch,
Sehr geehrte Damen und Herren der Berufungskommission,

hiermit sende ich Ihnen meine Bewerbungsunterlagen zu, insbesondere eine kurze Beschreibung meiner aktuellen und früheren Tätigkeiten in Forschung und Lehre sowie drei rezente Publikationen. Diese bieten Ihnen einen Einblick in meine Forschung, die das Entwerfen und die Anwendung eingebetteter Sensorsysteme für Wearable Computing und Drahtlose Sensornetze beinhaltet.

Ich leite seit Oktober 2014 den Lehrstuhl „Eingebettete Systeme" an der Technischen Fakultät der Albert Ludwig Universität Freiburg. Vorher war ich als Leiter einer Nachwuchsgruppe im Emmy-Noether-Programm der Deutschen Forschungsgemeinschaft (DFG) an der Technischen Universität Darmstadt tätig. Im Wintersemester 2012 war ich Vertretungsprofessor an der Universität Passau. In den letzten fünf Jahren habe ich so erhebliche Erfahrungen in der Lehre und in der Betreuung von Doktoranden gesammelt.

Sollten Ihnen meine Bewerbungsunterlagen zusagen, stehe ich Ihnen gern für ein Vorstellungsgespräch zur Verfügung.

Mit freundlichen Grüßen,

Prof. Dr. Kristof Van Laerhoven
Eingebettete Systeme, Uni Freiburg

http://es.informatik.uni-freiburg.de

Kristof Van Laerhoven

Schwerpunkte meiner Forschungsaktivitäten

Sensoren sind wichtiger Bestandteil vieler eingebetteter Systeme und aus dem heutigen Alltag kaum mehr wegzudenken. Der technologische Fortschritt in Miniaturisierung und Datenübertragung ermöglicht es mittlerweile, Sensoren von der Stange zu kaufen, die immer kleiner und energiesparender werden. Dadurch verschiebt sich der Forschungsschwerpunkt immer mehr auf die Ebenen der Signal- und Datenverarbeitung in eingebetteten Systemen und des systematisch-energieeffizienten Hardwareentwurfs, denn trotz der wachsenden Leistung von miniaturisierten eingebetteten Systemen sind diese immer noch die „Knackpunkte“ für viele Anwendungen.

Aufgrund dessen sind wegen der beschränkten Ressourcen, der komplexen Signale und der interessanten Anwendungen im Alltag tragbare Sensorknoten ein gutes Beispiel. Im Rahmen meiner Arbeit bevorzuge ich vor allem Evaluationen von eingebetteten Systemen in anspruchsvollen und komplexen Anwendungsgebieten wie zum Beispiel in der Schlafmusteranalyse (Somnographie / Circadian Rhythm Analysis), der alltagsnahen Analyse von Verhaltensproblemen (in der Psychologie: Ambulant Assessment) und im Monitoring mit Hilfe von drahtlosen Sensornetzen (z.B. für Gleisüberwachung). In diesen Studien wurden individuelle Hardwaresysteme entworfen um die energieeffiziente Erkennung von anwendungs-relevanten Mustern (z.B., die Bewegungsmustern und der Haltung des Körpers, der das Modul trägt), zu ermöglichen. Dabei sind nicht nur die eingebetteten Algorithmen, die von den rohen Sensordaten abstrahieren, wichtig. Auch an der Schnittstelle zwischen Software und Hardware gibt es eine große Spannbreite, eingebettete Systeme zu gestalten, die sowohl effizient als auch genau funktionieren müssen.

Ich habe langjährige Erfahrung im Einwerben von Drittmittel der DFG als Antragsteller des Emmy-Noether-Projektes “Long-Term Activity Recognition” sowie auch als Mitantragsteller der zwei DFG-Graduiertenkollegs “Topologie der Technik” und “Mixed-Mode Environments”. Weiterhin habe ich erfolgreich kleinere BMBF- und Industrie- Drittmittel eingeworben in Verbundprojekten mit Google Research und Bell Labs (die Forschungseinrichtung von Alcatel-Lucent).

Als Hauptantragsteller:

- DFG Emmy-Noether-Projekt LA 2758/1-1 (2010-2015) — ca. 883.000 Euro
- Industrieprojekte mit Google Research und Bell Labs (2014) — 85.000 Euro
- BMBF Projekt DEDIPAC WP113Da (2014-2017) — 34.827 Euro

Beteiligt als Mitantragsteller (und PI):

- DFG GRK Mixed-mode Environments (2011-2016), 2 Stipendiaten ca. 1.895k Euro
- DFG GRK Topologie der Technik (2011-2016), 2 Stipendiaten ca. 1.750k Euro

Außerdem sind die folgenden Anträge für Forschungsprojekte zur Zeit in der Begutachtungs- bzw. Schreibphase: DFG Sachbeihilfe " DiabetesStories" (mit Prof. Martin Knöll, TU Darmstadt), EU H2020 FET Open "N-CORDER" (mit Prof. Danny Hughes, KU Leuven, Belgien). Zwei "follow-up" Industriekooperationen für 2015 mit Google Research (Mountainview, USA) und Bell Labs (Dublin, Irland) werden ebenfalls geplant.

Durch meine derzeitige Beteiligung im Rahmen des DFG-Graduiertenkollegs „Topologie der Technik" und durch den Einsatz der von mir entwickelten Aktivitätssensoren in Studien zu Bipolarer Störung, Schlafforschung sowie Diabetes, bin ich mit der Zusammenarbeit in interdisziplinären Kontexten bestens vertraut.

In den letzten Jahren war ich mehrfach an der Planung und Durchführung internationaler wissenschaftlicher Tagungen und Kongressen beteiligt, zum Beispiel als Poster Chair für ACM UbiComp 2009 (Orlando, USA), Workshop Co-chair für AMI 2011 (Amsterdam, Niederlande), General Chair für ACM/IEEE ISWC 2013 (Zürich, Schweiz) oder als Program Co-chair für ACM MUM 2013 (Luleå, Schweden). In 2015 bin ich als TPC Chair für die ACM/IEEE ISWC Konferenz (in Osaka, Japan) gewählt worden.

Lehrveranstaltungen / Lehrkonzept

Über meine Forschung hinaus habe ich in den letzten fünf Jahren jedes Semester Lehraufträge übernommen und war dadurch, neben der Betreuung als Co-Referent von Bachelor- und Masterarbeiten, regelmäßig in den "Drahtlose Sensornetze"-Modulen des Studiengangs „Net Centric Systems" an der TU Darmstadt tätig. Als Vertretungsprofessor an der Universität Passau habe ich auch Erfahrung in der Gestaltung von Vorlesungen und Übungen im Bachelor Studiengang „Mobile und Eingebettete Systeme". Die Vorlesungen und Praktika wurden in Lehrevaluationen als sehr gut bewertet.

Insbesondere in Bezug auf Internationalität und Interdisziplinäre Vernetzung könnte ich mir gut vorstellen, meine Erfahrungen und Fähigkeiten auch an der Universität Freiburg einzubringen. Beispiele in diesem Zusammenhang sind englischsprachige Vorlesungen, die Organisation von internationalen Konferenzen, interdisziplinäre Studienprojekte, Netzwerkveranstaltungen und Mentorenprogramme für ausländische Studierende. An der TU Darmstadt habe ich gute Erfahrungen in sogenannten "Interdisziplinären Studieneingangsprojekten" gemacht - im Rahmen des Lehrkonzeptes „Kompetenz-Entwicklung durch interdisziplinäre Vernetzung von Anfang an" (KIVA), ein durch das BMBF gefördertes „Qualitätspakt Lehre"-Programm für bessere Studienbedingungen und interdisziplinäre Vernetzung in der Lehre. Ich habe hierbei Studierende aus der Informatik, Biologie, Philosophie und Politikwissenschaft während einer Projektwoche beraten und betreut (ca. 10 Gruppen von je 4-6 Studierenden).

Da ich fünf Jahre in England gelebt und gearbeitet habe, wurden die Mehrzahl meiner bisherigen Lehrveranstaltungen in englischer Sprache durchgeführt (z.B. Gastvorträge an der Lancaster University, der TU München und der Universität Lugano oder auch die englischsprachige Variante der Grundlagenvorlesungen „Grundlagen der Informatik 1" und „Grundlagen der Informatik 2" an der TU Darmstadt). Dies reflektiert die zunehmende Internationalisierung von sowohl der Forschung als auch der Lehre und Praxis in der Informatik -- gute Englischkenntnisse sind inzwischen unersetzlich geworden. Meine englischsprachigen Vorlesungen "Introduction to Computer Science 1" (2008) und "Introduction to Computer Science 2" (2007) wurden per Video/Audio aufgezeichnet und im Rahmen des E-Learning Pilotprojektes damals an der TU Darmstadt zusammen mit den Folien online zur Verfügung gestellt. Ich habe außerdem in den letzten Jahren praktische Erfahrung mit internetbasierten Lernumgebungen wie Moodle (seit 2010) und STUD.IP (2012) gesammelt. Diese Erfahrungen werde ich in Freiburg selbstverständlich einbringen.

Kristof Van Laerhoven

kvanlaerhoven@acm.org
http://es.informatik.uni-freiburg.de

Office: ES, Georges-Köhler-Allee 010, 79110 Freiburg, DE   +49 761 203 98 585
Home: Geyer-zu-Lauf-strasse 35, 79312 Emmendingen, DE   +49 173 801 77 70

Born on 16/11/1975 in Kapellen, Belgium   Nationality: Belgian
Languages: Dutch, French, English, German   Married, 2 children

Supervision: **8** PhDs (Erstgutachter: 6), **24** completed master/bachelor theses
h-index ≥ **24.** total citations ≥ **2863** (Google Scholar)
89 peer-reviewed publications: 79 conf., 6 journal art., 4 books. **1** patent

Academic Experience Overview

***2014*–now** Professor **University of Freiburg** **Germany**
W2 - Embedded Systems

Projects: (as principal investigator)
- DEDIPAC: Determinants of diet and activity, EU JPI, BMBF-funded (2014-2017)
- Google Glass Academic Award (funded by Google [x]) (2013-2015)
- Wireless Sensor and Micro Display Networks (funded by Bell Labs) (2013-2015)

Teaching:
- Praktikum Wearable Computing Systems (Master Embedded Systems Engineering)
- Systeme II - Rechnernetze (Bachelor Computer Science & Embedded Systems)

2012–2013 Guest Professor **University of Passau** **Germany**
Lehrstuhlvertretung Lukowicz, W3 - Embedded Systems

Teaching:
- Grundlagen der Mensch-Machine Interaktion (Bachelor Mobile & Embedded Systems)
- Intelligente Technische Systeme (Bachelor Intelligente Technische Systeme)

2010–2014 Group Leader **Technische Universität Darmstadt** **Germany**
DFG "Emmy Noether" Leader, with Promotionsrecht - Embedded Sensing Systems

Projects: (as principal investigator)
- Long-Term Activity Recognition: DFG-funded Emmy Noether project (2010-2015)
- Topologie der Technik: graduate school GRK 1343, DFG-funded (2011-2016)
- Mixed-mode Environments: graduate school GRK 1362, DFG-funded (2010-2016)

Teaching:
- Wireless Sensor Networks Seminar and Lab course (2010-2013) co-organiser

2006–2009 Post-Doc **Technische Universität Darmstadt** **Germany**
Prof. Bernt Schiele, *Multimodal Interactive Systems*

Projects:
- MOBVIS: Vision and Intelligent Maps for Mobile Interfaces. 3-year EU-IST STREP
- CESORA: Context Aware Support for Business Applications, 3-year funded by SAP

Teaching:
- Wireless Sensor Networks Seminar and Lab course (2006-2009)
- Introduction to Computer Science 1 (2007-2008) and 2 (2006-2007)

2001–2006 Research Associate **Lancaster University** **United Kingdom**
Prof. Hans–Werner Gellersen, *Embedded Interactive Systems*

Projects:
- UbiCare: 3-yr DTI, with Imperial, Univ. of Southampton, & 5 UK healthcare companies
- EQUATOR: 6-year IRC funded by EPSRC, with 7 top UK universities
- Pin&Play: 1-year EU FET assessment project, with Viktoria Institute, Sweden
- CommonSense: Distributed Context Acquisition, 3-year EPSRC project
- Smart-Its: 3-year EU STREP with ETH Zürich, Univ. of Karlsruhe & VTT Finland

Teaching:
- Introduction in Artificial Intelligence (2006-2007)

1999–2001 Researcher **Starlab Research, Brussels** **Belgium**
Prof. Walter Van de Velde. *Intelligent Systems*

Projects:
- TEA: 3-year EU-funded STREP with Nokia Research, University of Karlsruhe

Teaching:
- Industrial Design projects at Hogeschool Antwerpen

Education	2001–2005 Lancaster University, UK Thesis: "Embedded Perception"	**PhD in Computer Science** Supervisor: Prof. Hans-Werner Gellersen
	1997–1999 University of Brussels, Belgium "met grootste onderscheiding" == "with greatest distinction" Thesis: "Online Adaptive Context Awareness with Low-level Sensors"	**MSc in Computer Science** Supervisor: Prof. Bernard Manderick
	1993–1996 University of Hasselt, Belgium intermediate degree awarded after the first half of university studies	**Kandidaat Informatica**
	1986–1992 College Essen, Belgium	Highschool, Latin-Sciences

Project Reviewing

- external scientific promotor for the IWT-funded SensorTIP project (2001)
- evaluator for the European Commission's IST program (FP7, FET) (since 2005)
- served as expert reviewer for:
  the Hasler Stiftung (Switzerland) (2010), the IeMRC (UK) (2011), K.U. Leuven IDO programme (Belgium) (2012), KAIST (South Korea) (2012), IWT (Belgium) (2014), Stiftung Rheinland-Pfalz für Innovation (2014), etc.

Chairing Duties

a selection

Conferences:

EUSAI 2004, Eindhoven, The Netherlands:	Demo track chair
AmI 2011, Amsterdam, The Netherlands:	Workshops chair
IEEE ISWC 2006, Montreux Switzerland:	Doctoral Colloquium chair
IEEE ISWC 2008, Pittsburgh, Pennsylvania:	Workshops chair
ACM Ubicomp 2009, Orlando, Florida:	Poster track chair
Pervasive 2011, San Francisco, California:	Doctoral Colloquium chair
ACM/IEEE ISWC 2013: Zürich, Switzerland:	**General chair**
MUM 2013: Luleä, Sweden:	**Technical Program Committee chair**
MUM 2015: Linz, Austria:	Poster track chair
ACM/IEEE ISWC 2015, Osaka, Japan:	**Technical Program Committee chair**

Intl. Workshops (selected):

ubiPCMM 2006, Orange County, California	Workshop co-chair
IWSAWC 2007, Toronto, Canada	Workshop co-chair
From Wearable to Mobile, 2015, Copenhagen	Workshop co-chair

PhD Students

✓ Christian Seeger	Event-driven middleware for body&ambient sensor Applications 6.2010-1.2014, **ausgezeichnet**, *co-advised w. Prof.A Buchmann*
✓ Sofia Nikitaki	Location Sensing via Spatial Sparsity and Matrix Completion 10.2011-6.2014, *co-advised with Prof. P. Tsakalides*
✓ Vinay Sachidananda	On maximizing the Efficiency of Multipurpose WSNs 2009 - 10.2014, *Erstgutachter*
✓ Eugen Berlin	Embedded motion characterization with shape matching 1.2010-10.2014, **ausgezeichnet**, *supervised as Erstgutachter*
✓ Marko Borazio	Modeling wearable activities from long-term deployments, 1.2010-11.2014, *supervised as Erstgutachter*
▪ Agha Muhammad	Unsupervised detection methods for group activities, started 8.2011, *supervised as Erstgutachter*
▪ Philipp Scholl	Activity recognition with instrumented artifacts: The wavelet approach started 11.2011, *supervised as Erstgutachter*
▪ Manuel Dietrich	How does Wearable Activity Recognition Influences Users' Actions? started in 9.2013 *supervised as Erstgutachter*

External Reviewer for PhD Students	▪ Nils Hammerla	Activity recognition in naturalistic environments 11.2014, *advisor: Thomas Ploetz, Newcastle University, UK*
	▪ Pablo Guerrero	Workflow support for low-power wireless sensor / actuator nets 1.2015, *advisor: Alejandro Buchmann, TU Darmstadt, DE*
	▪ Tobias Große-Puppendahl	Natural interaction, activity & behaviour from capacitive sensing 5.2015, *advisor: Dieter Fellner, TU Darmstadt, DE*

Masters and Bachelor Thesis Supervision

- Gitte Robeyns: *Cochlear Implant System*. Product Design, Antwerp Academy for Industrial Design, Belgium, 2001
- Alexander Crolla: *EAR: Electronic Auditive Reaction*, Product Design, Antwerp Academy for Industrial Design, Belgium, 2001
- Christian Decker: *Simulation of Routing Algorithms in an Ad-Hoc Network*. Computer Science, University of Karlsruhe, Germany, 2001
- Dikaios Papadogkonas: *Investigation in Activity Recognition Algorithms on Embedded Platforms*, Computer Science, University of Lancaster, UK, 2004
- Martin Berchtold: *Processing Sensor Data with CSTK*, Computer Science, University of Karlsruhe, Germany, 2004
- André Kvist Aronsen: *Long-Term Fine-Grained Actigraphy*, Computer Science, TU Darmstadt, Germany, 2006
- Eugen Berlin: *Optimizing Wearable Sensors under Technical and Usability Constraints*, Computer Science, TU Darmstadt, Germany, 2007
- Urs Glaubitt & Max Lehn: *Analysis and Visualization of Circadian Rhythms from Actigraphy Data*, TU Darmstadt, Germany, 2007
- Marko Borazio: *Sustained Logging in Sleep Studies through Energy-Efficient Actigraphy*, TU Darmstadt, Germany, 2007
- David Kilian: *User-Annotated Activity Data Gathering and Processing*, Computer Science, TU Darmstadt, Germany, 2008
- Jun Liu: *Combining RFID and Acceleration to Detect Living Patterns*, Computer Science, TU Darmstadt, Germany, 2008
- Matthias Altmann: *Monitoring Athletes' Performance during Prolonged Training,* Computer Science, TU Darmstadt, Germany, 2008
- Delphine Christin & François Phillipp: *Collaborative Activity Sensing within a Zigbee Network*, ENSEA, France, 2008
- Holger Becker: *Monitoring of Long-Term, Low-Fidelity Sleep Characteristics*, Computer Science, TU Darmstadt, Germany, 2010
- Alexander Wöhnl: *Visualization and Recognition of Sleep Characteristics*, Computer Science, TU Darmstadt, Germany, 2011
- Johannes Reichard: *Design of a ZigBee Module in an Activity Sensor*, Computer Science, TU Darmstadt, Germany, 2012
- Thomas Pignede: Development of an Activity Recognition System for Fitness Trails, Computer Science, TU Darmstadt, Germany, 2012
- Nagihan Kücükyildiz: *Acquisition and Evaluation of Sensor Data in a Sleep Laboratory*, Computer Science, TU Braunschweig, Germany, 2012
- Patrick Frankenberger: *Evaluation of Visualization Techniques on Wearable Sensor Accelerometry Data,* Computer Science, TU Darmstadt, Germany, 2012
- Jan Hendrik Burdinski: *Evaluation of Long-Term Activity Data from a Mobile Phone* Computer Science, TU Darmstadt, Germany, 2013
- Jakob Weigert: Automatisierte Testläufe des Audi Infotainment-Systems, Audi AG & Uni Passau, Germany, 2013.
- Martin Jaensch: Aktivitätserkennung mit Google Glass, Smartphone und Handgelenkssensor, Informatik TU Darmstadt, Germany, 2014.
- Georg Schneider: Long-term Auditory Habitat Monitoring with Android, Embedded Systems Engineering, Uni Freiburg, Germany, 2015
- Marian Theiss: Detecting objects through grasp and grasp intention, Embedded Systems Engineering, Uni Freiburg, Germany, 2015

Courses Taught

- Systeme II: Rechnernetze (Deutsch, 2. Semester Bachelor)…… University of Freiburg
  2015, co-organizer with Prof. Schindelhauer (±150 students) 6 ECTS / 4 SWS
- Praktikum Wearable Computing (English, MSc) ………………..… University of Freiburg
  2014 - now, Initiator (±30 students) 6 ECTS / 4 SWS
- Seminar Wearable and Networked Systems (English, MSc) …… University of Freiburg
  2014 - now, Initiator (±20 students)) 6 ECTS / 4 SWS
- Intelligente Technische Systeme (ITS) (Deutsch, MSc) …………. University of Passau
  2012, Lehrstuhlvertretung Eingebettete Systeme (12 students) 7 ECTS / 6 SWS
- Grundlagen der Mensch-Machine Interaktion (Deutsch) .... University of Passau
  2012, Lehrstuhlvertretung Eingebettete Systeme (±40 students) 5 ECTS / 3 SWS
- Drahtlose Sensornetze / WSN Seminar (English)................................... TU Darmstadt
  2007 - now, Initiator and co-organizer (±20 students) 6 ECTS/4 SWS
- Drahtlose Sensornetze / WSN Praktikum (English) .... TU Darmstadt
  2007 - now, Initiator and co-organizer (±20 students) 6 ECTS / 4 SWS
- Grundlagen der Informatik 1 / Introduction to Computer Science 1 .... TU Darmstadt
  2008, English lectures (±80 students) 10 ECTS
- Grundlagen der Informatik 2 / Introduction to Computer Science 2 .... TU Darmstadt
  2007, Englisch lectures (±80 students) 10 ECTS

selected guest lectures:

- Ubiquitous Computing (Prof. Matthias Kranz) .... TU München
- Ubiquitous and Pervasive Computing (Prof. Marc Langheinrich) .... Univ. Lugano
- Artificial Intelligence 1 (Prof. Hans-Werner Gellersen) ....Univ. of Lancaster

Invited Talks

a selection

- “The Pervasive Sensor” - Opening keynote at UCS 2004, Tokyo, Japan
- “Multi-Sensor Context Awareness” - BSN 2004, Imperial College, London, UK
- “Balancing Dimensionality and Quality of Sensing”, ETH Zürich 2005, Switzerland
- “Embedded Sensing and Learning” - FITLab, University of Swansea, 2006, UK
- “Wearable Activity Recognition” - Fraunhofer IAIS, 2007, Bonn, Germany
- “Intro in Wearable Activity Systems” – Microsoft Research, 2008, Cambridge, UK
- “Detecting Activity in Time Series” - NEC Research, 2009, Heidelberg, Germany
- “Fine-grained Activity Monitoring” - EPA Conference, 2009, Thessaloniki, Greece
- “Walking with Porcupines” - invited lecture KAIST, 2010, Daejeon, Republic of Korea
- “Activity Recognition: What is Taking so Long?” - BMI, 2011, Berlin, Germany
- “Long-Term Challenges for Activity Recognition” - Opportunity WS, Anchorage, USA
- “Long-Term Activity Recognition” -iCareNET autumn school, 2011, Passau, Germany
- “Activity Assessment” - Karlsruhe Institute of Technology, 2012, Karlsruhe, Germany
- “Beyond Actigraphy: Discovery of Routines and Activities” -- SAA 2013, Amsterdam

Publications http://es.informatik.uni-freiburg.de/index.php/publications **Google Citations h-index: 24**

Total amount of peer-reviewed publications: **89** total citations ≥ **2863**

conference articles: **79** journal articles: **6** books: 4 issued patents: **1**

Agha Muhammad and Kristof Van Laerhoven, "DCS: A Divide and Conquer Strategy for Detecting Overlapping Communities in Social Graphs [Extended Abstract]", International Conference on Computational Social Science (ICCSS), Helsinki, Finland, 06/2015.

Agha Muhammad and Kristof Van Laerhoven, "DUKE: A Solution for Discovering Neighborhood Patterns in Ego Networks", The 9th International Conference on Web and Social Media (ICWSM), Oxford, England, AAAI, 05/2015. **[19% accept.rate]**

Eugen Berlin, Martin Zittel, Michael Bräunlein and Kristof Van Laerhoven, "Low-power Lessons from Designing a Wearable Logger for Long-term Deployments", 2015 Sensors Applications Symposium (SAS 2015), Zadar, Croatia, IEEE, 2015.

Marko Borazio, Eugen Berlin, Nagihan Kücükyildiz, Philipp M Scholl and Kristof Van Laerhoven, "Towards Benchmarked Sleep Detection with Inertial Wrist-worn Sensing Units", ICHI 2014, Verona, Italy, IEEE Press, 09/2014.

Manuel Dietrich and Kristof Van Laerhoven, "How does Wearable Activity Recognition Influence Users' Actions? A Computer Science and Philosophy Interdisciplinary Investigation", IACAP 2014, Thessaloniki, Greece, Springer, 08/2014.

Matthias Wille, Sascha Wischniewski, Philipp M Scholl and Kristof Van Laerhoven, "Comparing Google Glass with Tablet-PC as Guidance System for Assembling Tasks", Glass & Eyewear Computers (GEC), Zurich, Switzerland, IEEE Press, 06/2014.

Philipp M. Scholl, Marko Borazio, Martin Jänsch and Kristof Van Laerhoven, "Diary-Like Long-Term Activity Recognition: Touch or Voice Interaction?", Glass & Eyewear Computers (GEC), Zurich, Switzerland, IEEE Press, 06/2014.

Iliya Gurov, Pablo Guerrero, Martina Brachmann, Silvia Santini, Kristof Van Laerhoven and Alejandro Buchmann, "A Site Properties Assessment Framework for Wireless Sensor Networks", The 11th ACM Conference on Embedded Networked Sensor Systems (SenSys 2013), Rome, Italy, ACM Press, 12/2013.

Marko Borazio and Kristof Van Laerhoven, "Using Time Use with Mobile Sensor Data: A Road to Practical Mobile Activity Recognition?", 12th International Conference on Mobile and Ubiquitous Multimedia, Lulea, Sweden, ACM Press, 12/2013

Christian Seeger, Kristof Van Laerhoven, Jens Sauer and Alejandro Buchmann, "A Publish/Subscribe Middleware for Body and Ambient Sensor Networks that Mediates between Sensors and Applications", Int'l Conf on Healthcare Informatics (ICHI 2013), Philadelphia, PA, IEEE Press, 10/2013. **[best paper award]**.

Philipp M Scholl, Nagihan Kücükyildiz and Kristof Van Laerhoven, "When Do You Light a Fire? Capturing Tobacco Use with Situated, Wearable Sensors", First Workshop on Human Factors and Activity Recognition in Healthcare, Wellness and Assisted Living, Zurich, Switzerland, ACM Press, 09/2013.

Philipp M Scholl, Brahim El Mahjoub, Silvia Santini and Kristof Van Laerhoven, "Connecting Wireless Sensor Networks to the Robot Operating System", second International Workshop on Cooperative Robots and Sensor Networks (RoboSense 2013), Nova Scotia, Canada, IEEE Press, 06/2013.

Sofia Nikitaki, Philipp M Scholl, Kristof Van Laerhoven and Panagiotis Tsakalides, "Localization in Wireless Networks via Laser Scanning and Bayesian Compressed Sensing", IEEE 14th Workshop on Signal Processing Advances in Wireless Communications (SPAWC 2013): IEEE Press, 06/2013.

Agha Muhammad and Kristof Van Laerhoven, "Quantitative Analysis of Community Detection Methods for Longitudinal Mobile Data", International Conference on Social Intelligence and Technology (SOCIETY) 2013, State College, Pennsylvania USA, 05/2013.

Eugen Berlin and Kristof Van Laerhoven, "Sensor Networks for Railway Monitoring: Detecting Trains from their Distributed Vibration Footprints", 9th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS 2013), Cambridge, MA, USA, IEEE Press, 05/2013.

Marko Borazio and Kristof Van Laerhoven, "Improving Activity Recognition without Sensor Data: A Comparison Study of Time Use Surveys", 4th International Augmented Human Conference , Stuttgart, Germany, ACM Press, 03/2013.

Matthias Kreil, Kristof Van Laerhoven and Paul Lukowicz, "Allowing Early Inspection of Activity Data from a Highly Distributed Bodynet with a Hierarchical-Clustering-of-Segments Approach", IEEE BSN, Boston, MA, IEEE Press, 2013

Daniel Roggen, Daniel Gatica-Perez, Masaaki Fukumoto and Kristof Van Laerhoven, "ISWC 2013-Wearables Are Here to Stay.", IEEE Pervasive Computing, vol. 13, no. 1, pp. 14-18, 02/2014.

Pablo Guerrero, Iliya Gurov, Alejandro Buchmann and Kristof Van Laerhoven, "Diagnosing the Weakest Link in WSN Testbeds: A Reliability and Cost Analysis of the USB Backchannel", SenseApp 2012, Clearwater Beach, FL, USA, IEEE Press, **[20% accept.rate]**.

Agha Muhammad and Kristof Van Laerhoven, "Discovery of User Groups within Mobile Data", Nokia Mobile Data Challenge (MDC 2012), Newcastle, UK.

Eugen Berlin and Kristof Van Laerhoven, "Detecting Leisure Activities with Dense Motif Discovery", 14th ACM International Conference on Ubiquitous Computing (UbiComp 2012), Pittsburgh, PA, USA, ACM, 09/2012, **[19.3% accept.rate]**.

Philipp M Scholl, Stefan Kohlbrecher, Vinay Sachidananda and Kristof Van Laerhoven, "Fast Indoor Radio-Map Building for RSSI-based Localization Systems", The Ninth International Conference on Networked Sensing Systems (INSS 2012), Antwerp, Belgium, IEEE, 06/2012.

Albrecht Schmidt, Alireza Sahami Shirazi and Kristof Van Laerhoven, "Are You in Bed with Technology?", IEEE Pervasive Computing, vol. 11, no. 4, pp. 4-7, 10/2012.

Philipp M Scholl and Kristof Van Laerhoven, "A Feasibility Study of Wrist-Worn Accelerometer Based Detection of Smoking Habits", esIoT 2012, Palermo, Italy, IEEE Press, 05/2012.

Eugen Berlin and Kristof Van Laerhoven, "Trainspotting: Combining Fast Features to Enable Detection on Resource-constrained Sensing Devices", The Ninth International Conference on Networked Sensing Systems (INSS 2012), Antwerp, Belgium, IEEE Press, 06/2012.

Philipp M Scholl, Matthias Berning, Kristof Van Laerhoven, Markus Scholz and Dawud Gordon, "jNode: a Sensor Network Platform that Supports Distributed Inertial Kinematic Monitoring", The Ninth International Conference on Networked Sensing Systems (INSS 2012), Antwerp, Belgium, IEEE Press, 06/2012.

Marko Borazio and Kristof Van Laerhoven, "Combining Wearable and Environmental Sensing into an Unobtrusive Tool for Long-Term Sleep Studies", 2nd ACM SIGHIT International Health Informatics Symposium (IHI 2012), Miami, Florida, USA, ACM Press, 01/2012, **[18% accept.rate]**.

Christian Seeger, Alejandro Buchmann and Kristof Van Laerhoven, "An Event-based BSN Middleware that supports Seamless Switching between Sensor Configurations", 2nd ACM SIGHIT International Health Informatics Symposium (IHI 2012), Miami, FL, ACM Press, 01/2012, **[18% accept.rate]**.

Pablo Guerrero, Alejandro Buchmann, Kristof Van Laerhoven, Immanuel Schweizer, Max Mühlhäuser, Thorsten Strufe, Stefan Schneckenburger, Manfred Hegger, and Birgitt Kretzschmar "A Metropolitan-Scale Testbed for Heterogeneous Sensor Networks to Support CO2 Reduction", GreeNETS 2012, Gandia, Spain.

Christian Seeger, Alejandro Buchmann and Kristof Van Laerhoven, "Wireless Sensor Networks in the Wild: Three Practical Issues after a Middleware Deployment", the Sixth International Workshop on Middleware Tools, Services and Run-time Support for Networked Embedded Systems (MidSens 2011), Lisbon, Portugal, ACM Press, 12/2011.

Christian Seeger, Alejandro Buchmann and Kristof Van Laerhoven, "myHealthAssistant: A Phone-based Body Sensor Network that Captures the Wearer's Exercises throughout the Day", The 6th International Conference on Body Area Networks, Beijing, China , ACM Press, 11/2011, **[best paper award]**.

Agha Muhammad, Niklas Klein, Kristof Van Laerhoven and Klaus David, "A Feature Set Evaluation for Activity Recognition with Body-Worn Inertial Sensors", Workshop on Interactive Human Behavior Analysis in Open or Public Spaces (InterHub) 2011, Amsterdam, Springer Verlag, 11/2011.

Marko Borazio and Kristof Van Laerhoven, "Predicting Sleeping Behaviors in Long-Term Studies with Wrist-Worn Sensor Data", International Joint Conference on Ambient Intelligence (AmI-11), vol. LNCS 7040, Amsterdam, Springer Verlag, pp. 151–156, 11/2011.

Christian Seeger, Alejandro Buchmann and Kristof Van Laerhoven, "Poster Abstract: Adaptive Gym Exercise Counting for myHealthAssistant", Body Area Networks (BodyNets), Beijing, China, ACM Press, 11/2011.

Holger Becker, Marko Borazio and Kristof Van Laerhoven, "How to Log Sleeping Trends? A Case Study on the Long-Term Capturing of User Data", The 5th European Conference on Smart Sensing and Context 2010 (EuroSSC 2010), vol. 6446, Passau, Germany, Springer Verlag, pp. 15-27, 2010.

Kyu-ho Park, Hoi-jun Yoo and Kristof Van Laerhoven, "Proceedings of the 14th International Symposium on Wearable Computers (ISWC 2010)", ISWC 2010, Seoul, South Korea, IEEE Press, 12/2010.

Daniel Jacobi, Pablo Ezequiel Guerrero, Khalid Nawaz, Christian Seeger, Arthur Herzog, Kristof Van Laerhoven and Ilia Petrov, "Towards Declarative Query Scoping in Sensor Networks", From Active Data Management to Event-Based Systems and More, vol. 6462: Springer Verlag, pp. 281-292, 11/2010.

Eugen Berlin and Kristof Van Laerhoven, "An On-Line Piecewise Linear Approximation Technique for Wireless Sensor Networks", 5th IEEE International Workshop on Practical Issues in Building Sensor Network Applications (SenseApp 2010), Denver, Colorado, USA, IEEE Computer Society, pp. 921-928, 10/2010.

K. Van Laerhoven, "ISWC 2010: The Latest in Wearable Computing Research", IEEE Pervasive Computing, vol. 10, no. 1, New York, IEEE Press, 01/2011.

Marko Borazio, Ulf Blanke and Kristof Van Laerhoven, "Characterizing Sleeping Trends from Postures", Proceedings of the 14th IEEE International Symposium on Wearable Computers (ISWC 2010), Seoul, South Korea, IEEE Press, pp. 167-168, 10/2010.

Ulf Blanke, Diane Larlus, Kristof Van Laerhoven and Bernt Schiele, "Standing on the Shoulders of Other Researchers - A Position Statement", Proc. of the Workshop "How to do good activity recognition research? Experimental methodologies, evaluation metrics, and reproducibility issues" (Pervasive 2010), Helsinki, Finland, 05/2010.

Kazuya Murao, Kristof Van Laerhoven, Tsutomu Terada and Shojiro Nishio, "A Method for Context Recognition Using Peak Values of Sensors", Transactions of Information Processing Society of Japan, vol. 51, no. 3: Information Processing Society of Japan (IPSJ), pp. 1068-1077, 03/2010.

Eugen Berlin, Jun Liu, Kristof Van Laerhoven and Bernt Schiele, "Coming to Grips with the Objects We Grasp: Detecting Interactions with Efficient Wrist-Worn Sensors", International Conference on Tangible and Embedded Interaction (TEI 2010), Cambridge MA, USA, ACM Press, pp. 57-64, 01/2010, [34% accept.rate].

Eugen Berlin, Pablo Guerrero, Arthur Herzog, Daniel Jacobi, Kristof Van Laerhoven, Bernt Schiele and Alejandro Buchmann, "Demo Abstract: Whac-A-Bee -- A Sensor Network Game", Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys 2009), Berkeley, California, USA, ACM Press, pp. 333-334, 11/2009.

Kristof Van Laerhoven and Eugen Berlin, "When Else Did This Happen? Efficient Subsequence Representation and Matching for Wearable Activity Data", Proceedings of the 13th International Symposium on Wearable Computers (ISWC 2009): IEEE Press, pp. 69-77, 2009. [28% accept.rate]

Kristof Van Laerhoven, Eugen Berlin and Bernt Schiele, "Enabling Efficient Time Series Analysis for Wearable Activity Data", Proceedings of the 8th International Conference on Machine Learning and Applications (ICMLA 2009), Miami Beach, FL, USA, IEEE Press, pp. 392-397, 2009.

Maja Stikic, Kristof Van Laerhoven and Bernt Schiele, "Exploring Semi-Supervised and Active Learning for Activity Recognition", Proceedings of the 12th International Symposium on Wearable Computers (ISWC 2008), Pittsburgh, USA, IEEE Press, pp. 81-90, September, 2008, [26% accept.rate].

Maja Stikic, Tâm Huynh, Kristof Van Laerhoven and Bernt Schiele, "ADL Recognition Based on the Combination of RFID and Accelerometer Sensing", Proceedings of the 2nd International Conference on Pervasive Computing Technologies for Healthcare (Pervasive Health 2008), Tampere, Finland, IEEE Xplore, pp. 258–263, January, 2008.

Kristof Van Laerhoven, Marko Borazio, David Kilian and Bernt Schiele, "Sustained Logging and Discrimination of Sleep Postures with Low-Level, Wrist-Worn Sensors", Proceedings of the 12th International Symposium on Wearable Computers (ISWC 2008): IEEE Press, pp. 69-77, 2008, [26% accept.rate].

Kristof Van Laerhoven, David Kilian and Bernt Schiele, "Using Rhythm Awareness in Long-Term Activity Recognition", Proceedings of the 12th International Symposium on Wearable Computers (ISWC 2008), Pittsburgh, PA, USA, IEEE Press, pp. 63-68, 2008.

Martin Berchtold, Till Riedel, Christian Decker and Kristof Van Laerhoven, "Gath-Geva Specification and Genetic Generalization of Takagi-Sugeno-Kang Fuzzy Models", Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC 2008): IEEE Press, 2008.

Maja Stikic and Kristof Van Laerhoven, "Recording Housekeeping Activities with Situated Tags and Wrist-Worn Sensors: Experiment Setup and Issues Encountered", Proceedings of the 1st International Workshop on Wireless Sensor Networks for Health Care (WSNHC 2007), Braunschweig, Germany, June, 2007.

Paula Alexandra Silva and Kristof Van Laerhoven, "BadIdeas for Usability and Design of Medicine and Healthcare Sensors", Proceedings of the 3rd Human-computer interaction and usability engineering (USAB'07), LNCS vol. 4799, Graz, Austria, Springer, pp. 105–112, 11/2007.

Andreas Zinnen, Kristof Van Laerhoven and Bernt Schiele, "Toward Recognition of Short and Non-repetitive Activities from Wearable Sensors", European Conference on Ambient Intelligence (AmI 2007), vol. 4794, Darmstadt, Germany, Springer, pp. 142-158, 11/2007.

Kristof Van Laerhoven and Andre Aronsen, "Memorizing What You Did Last Week: Towards Detailed Actigraphy With A Wearable Sensor", 27th International Conference on Distributed Computing Systems Workshops (ICDCSW 2007), Toronto, Ontario, Canada, IEEE Computer Society, pp. 47-53, 2007.

Kristof Van Laerhoven and Hans-Werner Gellersen, "Fair Dice: A Tilt and Motion-Aware Cube with a Conscience.", 26th International Conference on Distributed Computing Systems Workshops (ICDCSW 2006): IEEE Computer Society, pp. 66-72, 07/2006.

Kristof Van Laerhoven, "Embedded Perception - Concept Recognition by Learning and Combining Sensory Data", Computer Science, vol. Ph.D., Lancaster, Lancaster University, 03/2006.

Kristof Van Laerhoven, Hans-Werner Gellersen and Yanni G Malliaris, "Long-Term Activity Monitoring with a Wearable Sensor Node", International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2006), Cambridge, Massachusetts, USA, IEEE Computer Society, pp. 171–174, 04/2006.

Kristof Van Laerhoven and Martin Berchtold, "Real-Time Analysis of Correlations Between On-Body Sensor Nodes", Proceedings of the 2nd International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2005): IEE Press, pp. 27-32, 2005.

Nicolas Villar, Gerd Kortuem, Kristof Van Laerhoven and Albrecht Schmidt, "The Pendle: A Personal Mediator for Mixed Initiative Environments", IEE International Workshop on Intelligent Environments, Essex, UK, IEE Press, 2005.

Kristof Van Laerhoven, Benny PL Lo, Jason WP Ng, Surapa Thiemjarus, Rachel King, Simon Kwan, Hans-Werner Gellersen, Morris Sloman, Oliver Wells, Phil Needham, et al., "Medical Healthcare Monitoring with Wearable and Implantable Sensors", International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications (UbiHealth), September, 2004.

Kristof Van Laerhoven and Hans-Werner Gellersen, "Spine versus Porcupine: A Study in Distributed Wearable Activity Recognition", 8th Intl. Symposium on Wearable Computers (ISWC 2004): IEEE Computer Society, pp. 142-149, 2004. [23% accept.rate]

Kristof Van Laerhoven, "The Pervasive Sensor - Invited Talk", Ubiquitous Computing Systems, Second International Symposium, UCS 2004: Springer Verlag, pp. 1-9, 2004. [opening keynote, peer-reviewed for post-proceedings]

Tara Matthews, Hans-Werner Gellersen, Kristof Van Laerhoven and Anind K Dey, "Augmenting Collections of Everyday Objects: A Case Study of Clothes Hangers As an Information Display", Pervasive Computing, Second International Conference, vol. 3001: Springer, pp. 340-344, 2004. [12% accept.rate]

Nicolas Villar, Kristof Van Laerhoven and Hans-Werner Gellersen, "A Physical Notice Board with Digital Logic and Display", Adjunct Proceedings of the European Symposium on Ambient Intelligence 2004 (AmI 2004): ACM Press, pp. 207-217, 2004.

Nicky Kern, Kristof Van Laerhoven, Hans-Werner Gellersen and Bernt Schiele, "Towards an Inertial Sensor Network", IEE EuroWearable, Birmingham, UK, September, 2003.

Jennifer Sheridan, Ben Short, Gerd Kortuem, Kristof Van Laerhoven and Nicolas Villar, "Exploring Cube Affordance: Towards A Classification Of Non-Verbal Dynamics Of Physical Interfaces For Wearable Computing", IEE EuroWearable 2003, pp. 113-118, September, 2003.

Kristof Van Laerhoven, Nicolas Villar and Hans-Werner Gellersen, "A Layered Approach to Wearable Textile Networks", IEE EuroWearable 2003, pp. 61-67, September, 2003.

Kristof Van Laerhoven, Nicolas Villar, Albrecht Schmidt, Gerd Kortuem and Hans-Werner Gellersen, "Using an autonomous cube for basic navigation and input", Proceedings of the 5th International Conference on Multimodal Interfaces, ICMI 2003: ACM, pp. 203-210, 2003.

Kristof Van Laerhoven, Albrecht Schmidt and Hans-Werner Gellersen, "Multi-Sensor Context Aware Clothing", 6th International Symposium on Wearable Computers (ISWC 2002): IEEE Press, pp. 49-56, 2002, [19% accept.rate].

Kristof Van Laerhoven, Albrecht Schmidt and Hans-Werner Gellersen, "Pin&Play: Networking Objects through Pins", UbiComp 2002: Ubiquitous Computing, 4th International Conference, vol. 2498: Springer, pp. 219-228, 2002, [18% accept.rate].

Albrecht Schmidt, Martin Strohbach, Kristof Van Laerhoven, Adrian Friday and Hans-Werner Gellersen, "Context Acquisition Based on Load Sensing", UbiComp 2002: Ubiquitous Computing, 4th International Conference, vol. 2498: Springer, pp. 333-350, 2002, [18% accept.rate].

Albrecht Schmidt, Martin Strohbach, Kristof Van Laerhoven and Hans-Werner Gellersen, "Ubiquitous Interaction - Using Surfaces in Everyday Environments as Pointing Devices", Universal Access: Theoretical Perspectives, Practice, and Experience, 7th ERCIM International Workshop on User Interfaces for All, vol. 2615: Springer, pp. 263-279, 2002.

Kristof Van Laerhoven, Nicolas Villar, Maria Hakansson and Hans-Werner Gellersen, "Pin&Play: Bringing Power and Networking to Wall-Mounted Appliances", Proceedings of the 5th IEE International Workshop on Networked Appliances, Liverpool, UK: IEEE Press, pp. 131-137, 2002.

Ozan Cakmakci, Joelle Coutaz, Kristof Van Laerhoven and Hans-Werner Gellersen, "Context Awareness in Systems with Limited Resources", Proc. of the third workshop on Artificial Intelligence in Mobile Systems, pp. 21-29, 2002.

Kristof Van Laerhoven, "Combining the Self-Organizing Map and K-Means Clustering for On-Line Classification of Sensor Data", Artificial Neural Networks - ICANN 2001, International Conference, vol. 2130: Springer, pp. 464-469, 2001.

Kristof Van Laerhoven, Kofi A Aidoo and Steven Lowette, "Real-time Analysis of Data from Many Sensors with Neural Networks", 5th International Symposium on Wearable Computers (ISWC 2001): IEEE Computer Society, pp. 115-122, 2001.

Kristof Van Laerhoven and Ozan Cakmakci, "What Shall We Teach Our Pants?", ISWC '00: Proceedings of the 4th IEEE International Symposium on Wearable Computers, Washington, DC, USA, IEEE Computer Society, pp. 77, 2000.

Albrecht Schmidt, Kofi A Aidoo, Antti Takaluoma, Urpo Tuomela, Kristof Van Laerhoven and Walter Van de Velde, "Advanced Interaction in Context", Proc. of the First International Sympoosium on Handheld and Ubiquitous Computing (HUC'99), September, Springer Verlag,, Karlsruhe, Germany, pp. 89-101, 1999.

Refereed Journal Publications:

- Christian Seeger, Kristof Van Laerhoven and Alejandro Buchmann, "MyHealthAssistant: An Event-driven Middleware for Multiple Medical Applications on a Smartphone-mediated Body Sensor Network", IEEE Journal of Biomedical and Health Informatics (J-BHI), vol. 19, no. 2, 03/2015.
- Alireza Sahami Shirazi, James Clawson, Yashar Hassanpour, Mohammad J Tourian, Ed Chi, Marko Borazio, Albrecht Schmidt and Kristof Van Laerhoven, "Already Up? Using Mobile Phones to Track & Share Sleep Behavior", International Journal of Human-Computer Studies, vol. 71, no. 9, Elsevier Press, 2013.
- K. Van Laerhoven and B. Schiele, “Energieeffiziente Datenverarbeitung auf modularen Sensorknoten”, in Thema Forschung 1/2007 “Ambient Intelligence”, pp. 34-39. 2007.
- Kristof Van Laerhoven, Nicolas Villar, Albrecht Schmidt, Hans-Werner Gellersen and Lars Erik Holmquist, "Pin&Play: The Surface as Network Medium" In IEEE Communications Magazine, April 2003, Vol.41 No.4. IEEE Press, pp. 90-96. 2003.
- Albrecht Schmidt and Kristof Van Laerhoven, "How to Build Smart Appliances" In IEEE Personal Communications, Special Issue on Pervasive Computing, August 2001, Vol. 8, No. 4. IEEE Press, pp. 66-71. 2001.
- Kristof Van Laerhoven and Kofi A Aidoo, "Teaching Context to Applications". In Personal and Ubiquitous Computing: Situated Interaction and Context-Aware Computing, A. Schmidt, G. Kortuem, D. Morse and A. Dey (Eds). Vol. 5-1, Springer Verlag, Berlin, Germany. 2001. pp.46-49. 2001.

Books / Proceedings:

- Kristof Van Laerhoven, Daniel Roggen, Daniel Gatica-Perez, Masaaki Fukumoto:Proceedings of the 17th Annual International Symposium on Wearable Computers. ISWC 2013, Zurich, Switzerland, September 8-12, 2013. ACM 2013, ISBN 978-1-4503-2127-3
- Matthias Kranz, Kåre Synnes, Sebastian Boring, Kristof Van Laerhoven: 12th International Conference on Mobile and Ubiquitous Multimedia, MUM '13, Luleå, Sweden - December 02 - 05, 2013. ACM 2013, ISBN 978-1-4503-2648-3
- Reiner Wichert, Kristof Van Laerhoven and Jean Gelissen, "Constructing Ambient Intelligence". AmI 2011 Workshops Proceedings. Volume 277, Springer Verlag, Berlin, Germany. 2012.
- Hoi-Jun Yoo and Kristof Van Laerhoven, Proceedings of the 14th IEEE Intl. Symposium on Wearable Computers (ISWC 2010). IEEE Press. 2010.

Patents
issued

“Context acquisition based on load sensing”, US Pat. 7434459 -
Filed 26 Sep 2003 - Issued 14 Oct 2008 – SAP

Referees

Prof. Dr. Bernt Schiele
Max Planck Institute – Informatics, D2
Building E1 4, Stuhlsatzenhausweg 85,
66123 Saarbrücken, Germany
schiele@mpi-inf.mpg.de, +49 681 9325 2000

Prof. Dr. Hans-Werner Gellersen
Lancaster University – School of Computing and Communications
Infolab21, South drive
Lancaster LA1 4WA, United Kingdom
hwg@comp.lancs.ac.uk, +44 1524 510 323

Prof. Dr. Walter Van de Velde
European Commission – Future and Emerging Technologies
BU25 05 140, Av. De Beaulieue 25
B1160 Brussels, Belgium
Walter.VAN-DE-VELDE@ec.europa.eu, +32 229 880 85

Prof. Dr. Michael Beigl
Karlsruhe Institute of Technology – Pervasive Computing Systems
Vincenz-Priessnitz-Str.1,
76131 Karlsruhe, Germany
micael@teco.edu, +49 (721) 608 417 00

Prof. Dr. Gerhard Tröster
ETH Zürich – Institut für Elektronik
ETZ H 89, Gloriastr. 35,
8092, Zürich, Switzerland
troester@ife.ee.ethz.ch, +41 44 632 39 64

Prof. Dr. Alejandro Buchmann
TU Darmstadt – Computing Department
Hochschulstr. 10,
64289 Darmstadt, Germany
buchmann@informatik.tu-darmstadt.de, +49 (0) 6151 16 6230

5550V Intelligente Technische Systeme

Liebe Dozentin, lieber Dozent,
anbei erhalten Sie die Ergebnisse der Evaluation Ihrer Lehrveranstaltung.
Zu dieser Veranstaltung wurden 5 Bewertungen abgegeben.
Erläuterungen zu den Diagrammen befinden sich am Ende dieses Dokuments.
Mit freundlichen Grüßen,
Das Evaluationsteam

Persönliche Angaben

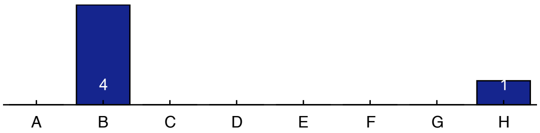

In welchem Studiengang sind Sie eingeschrieben?
- **A** Diplom Informatik (0)
- **B** Bachelor Informatik (4)
- **C** Bachelor Internet Computing (0)
- **D** Master Informatik (0)
- **E** Master IT-Sicherheit (0)
- **F** Lehramt Gymnasium Mathematik (0)
- **G** Lehramt Gymnasium Informatik (0)
- **H** andere (1)

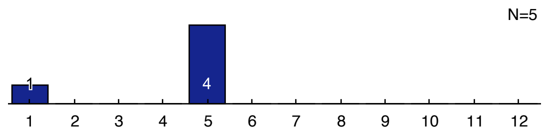

In welchem Fachsemester studieren Sie?

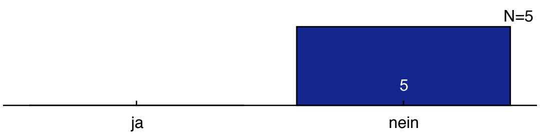

Haben Sie die Vorlesung schon einmal gehört?

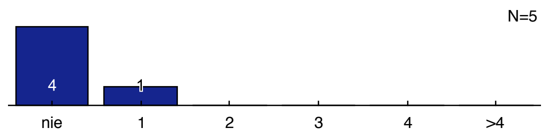

Wie oft haben Sie die Vorlesung versäumt?

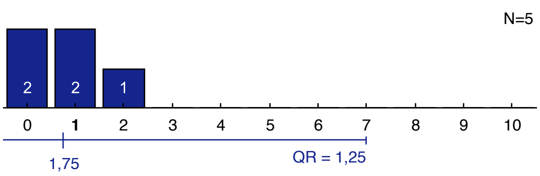

Wie viele Stunden pro Woche wenden Sie in etwa zusätzlich zur Veranstaltungsdauer auf? (Vorbereitung, Nachbereitung, selbstständiges Lösen der Übungsaufgaben)

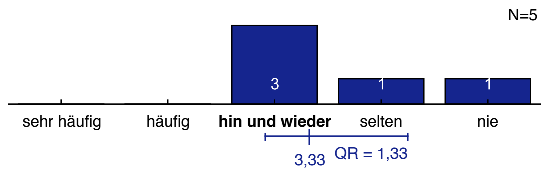

Wie oft haben Sie in der Vorlesung Fragen an den Dozenten / die Dozentin gestellt?

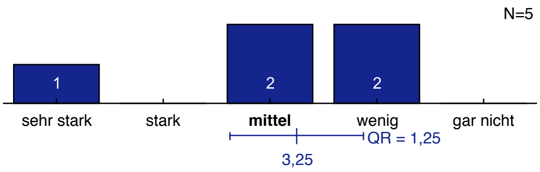

Wie sehr haben Sie sich in den Übungen engagiert? (Fragen stellen, Diskutieren, Lösungen vorstellen)

Fragen zur Vorlesung

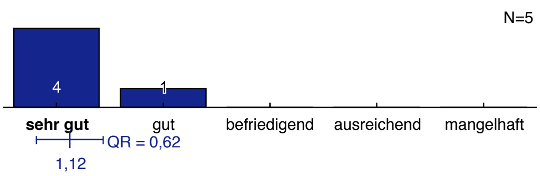

Wie bewerten Sie das Engagement des Dozenten / der Dozentin?

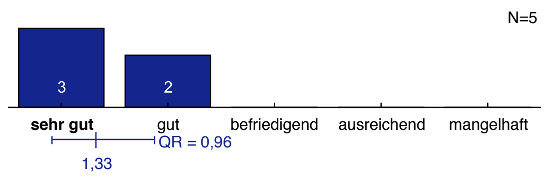

Wie erklärt der Dozent / die Dozentin den Vorlesungsstoff?

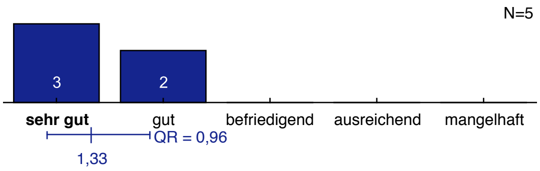

Wie geht der Dozent / die Dozentin auf studentische Fragen ein?

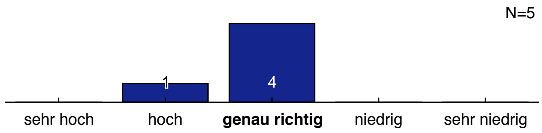

Wie beurteilen Sie das Tempo der Vorlesung?

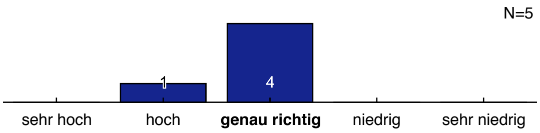

Wie schätzen Sie den Schwierigkeitsgrad der Vorlesung ein?

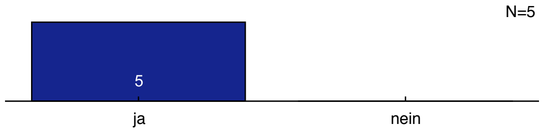

Gibt es schriftliche Unterlagen zur Vorlesung (Bücher, Folien, Skripte, Tafelfotos etc.)

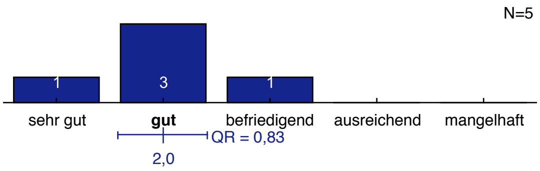

Wie beurteilen Sie ggf. die Nützlichkeit der schriftlichen Vorlesungsunterlagen?

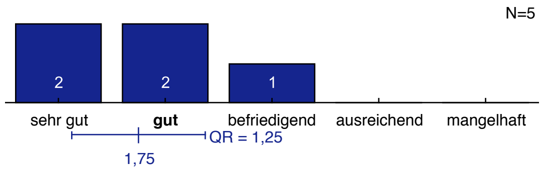

Wie beurteilen Sie die Präsentation in der Vorlesung (Tafelbild, Folien, Filme etc.)

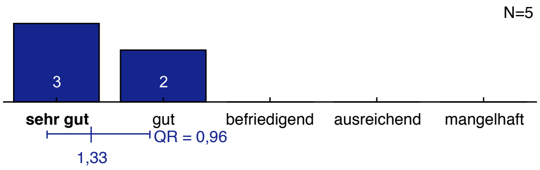

Wie beurteilen Sie die Vorlesung insgesamt?

Freitextkommentare

Persönliche Angaben

Fragen zur Vorlesung

Haben Sie Kommentare und Anregungen zur Vorlesung?

War genau richtig

Gute und lockere Athmosphäre

Erläuterungen zur Visualisierung

- Im oberen Teil des Bildes befindet sich ein Histogramm der absoluten Häufigkeiten. Hierbei ist die Höhe des Balkens proportional zur Anzahl der Nennungen.
- Darunter sind die möglichen Antworten abgetragen. Die Median-Antwort ist durch eine erhöhte Schriftgröße gekennzeichnet.
- Im unteren Bildteil befinden sich zwei gleichartige Visualisierungen von Median und Quartilen. Die obere, blaue Grafik kennzeichnet die Werte dieser Veranstaltung, die untere, graue diejenigen der Vergleichsgruppe.
- Als Vergleich dienen alle Veranstaltungen dieses Semesters, bei denen diese Frage gestellt wurde.
- N ist die Gesamtzahl der Nennungen

5806V Grundlagen der Mensch-Maschine-Interaktion

Liebe Dozentin, lieber Dozent,
anbei erhalten Sie die Ergebnisse der Evaluation Ihrer Lehrveranstaltung.
Zu dieser Veranstaltung wurden 17 Bewertungen abgegeben.
Erläuterungen zu den Diagrammen befinden sich am Ende dieses Dokuments.
Mit freundlichen Grüßen,
Das Evaluationsteam

Persönliche Angaben

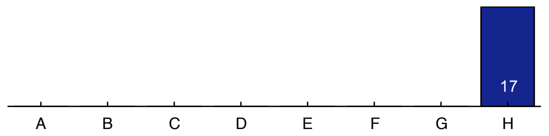

In welchem Studiengang sind Sie eingeschrieben?
- **A** Diplom Informatik (0)
- **B** Bachelor Informatik (0)
- **C** Bachelor Internet Computing (0)
- **D** Master Informatik (0)
- **E** Master IT-Sicherheit (0)
- **F** Lehramt Gymnasium Mathematik (0)
- **G** Lehramt Gymnasium Informatik (0)
- **H** andere (17)

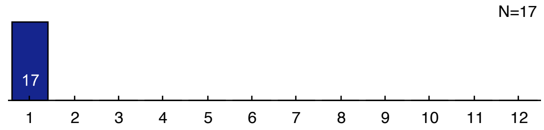

In welchem Fachsemester studieren Sie?

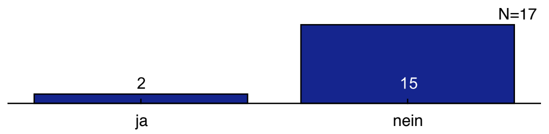

Haben Sie die Vorlesung schon einmal gehört?

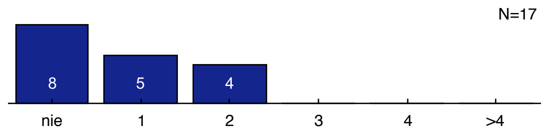

Wie oft haben Sie die Vorlesung versäumt?

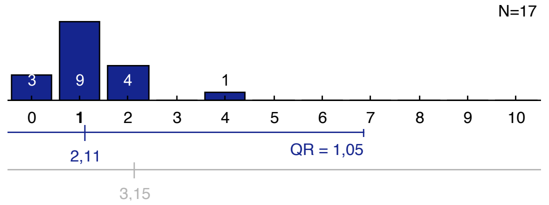

Wie viele Stunden pro Woche wenden Sie in etwa zusätzlich zur Veranstaltungsdauer auf? (Vorbereitung, Nachbereitung, selbstständiges Lösen der Übungsaufgaben)

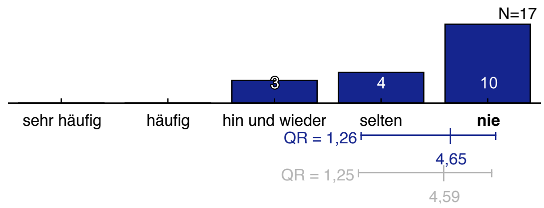

Wie oft haben Sie in der Vorlesung Fragen an den Dozenten / die Dozentin gestellt?

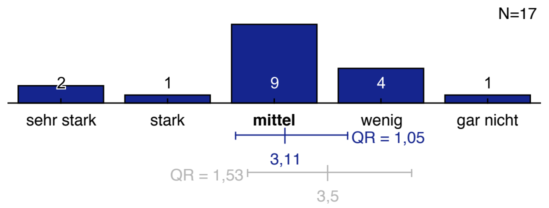

Wie sehr haben Sie sich in den Übungen engagiert? (Fragen stellen, Diskutieren, Lösungen vorstellen)

Fragen zur Vorlesung

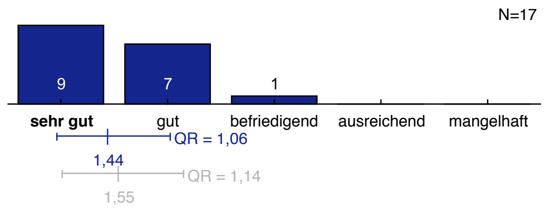

Wie bewerten Sie das Engagement des Dozenten / der Dozentin?

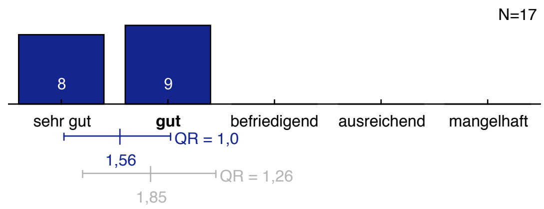

Wie erklärt der Dozent / die Dozentin den Vorlesungsstoff?

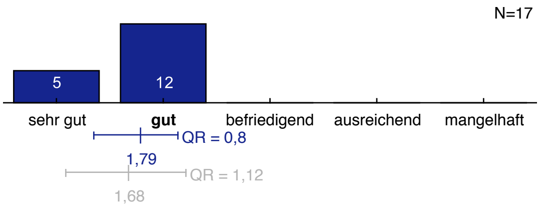

Wie geht der Dozent / die Dozentin auf studentische Fragen ein?

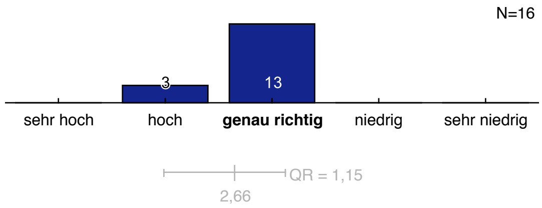

Wie beurteilen Sie das Tempo der Vorlesung?

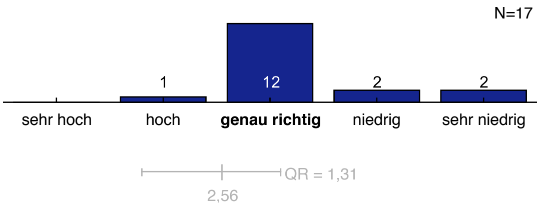

Wie schätzen Sie den Schwierigkeitsgrad der Vorlesung ein?

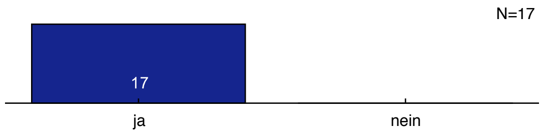

Gibt es schriftliche Unterlagen zur Vorlesung (Bücher, Folien, Skripte, Tafelfotos etc.)

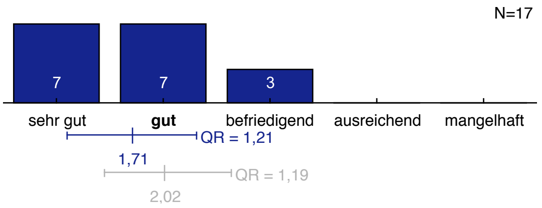

Wie beurteilen Sie ggf. die Nützlichkeit der schriftlichen Vorlesungsunterlagen?

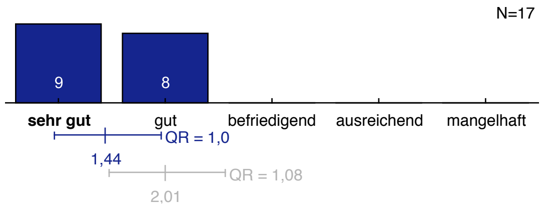

Wie beurteilen Sie die Präsentation in der Vorlesung (Tafelbild, Folien, Filme etc.)

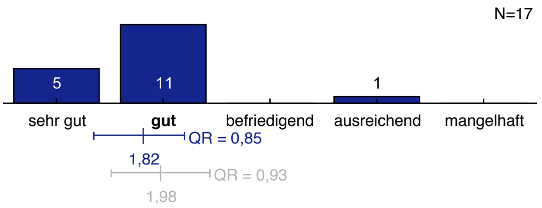

Wie beurteilen Sie die Vorlesung insgesamt?

Freitextkommentare

Persönliche Angaben

Fragen zur Vorlesung

Haben Sie Kommentare und Anregungen zur Vorlesung?

Die Unterlagen / Folien sind häufig auf Englisch, was etwas gewöhnungsbedürftig ist.

Folien auf englisch machen mir einwenig zu schaffen, ansonsten habe ich kein Problem :-)

Nein

Erläuterungen zur Visualisierung

- Im oberen Teil des Bildes befindet sich ein Histogramm der absoluten Häufigkeiten. Hierbei ist die Höhe des Balkens proportional zur Anzahl der Nennungen.
- Darunter sind die möglichen Antworten abgetragen. Die Median-Antwort ist durch eine erhöhte Schriftgröße gekennzeichnet.
- Im unteren Bildteil befinden sich zwei gleichartige Visualisierungen von Median und Quartilen. Die obere, blaue Grafik kennzeichnet die Werte dieser Veranstaltung, die untere, graue diejenigen der Vergleichsgruppe.
- Als Vergleich dienen alle Veranstaltungen dieses Semesters, bei denen diese Frage gestellt wurde.
- N ist die Gesamtzahl der Nennungen

Auswertung zur Übungsgruppe von Prof. Ph. D. Van Laerhoven zur Veranstaltung "Grundlagen der Mensch-Maschine-Interaktion"

Liebe Dozentin, lieber Dozent,
Diese Auswertung der Übungsgruppe ist nur Ihnen und der Übungsgruppenleiterin / dem Übungsgruppenleiter zugegangen. Bitte besprechen Sie gegebenenfalls die Ergebnisse persönlich mit ihr / ihm und behandeln Sie sie ansonsten vertraulich.

Fragen zur Übung

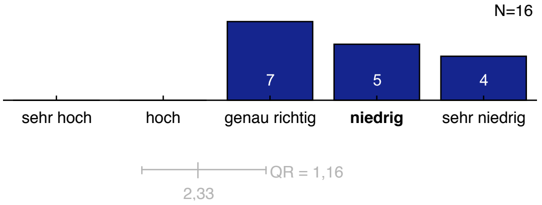

Wie schätzen Sie den Schwierigkeitsgrad der Übungsaufgaben ein?

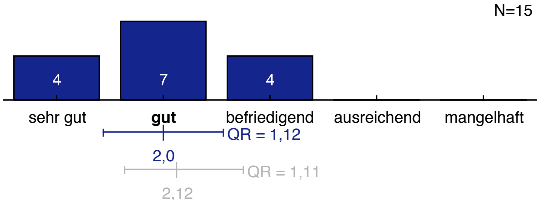

Wie sind Vorlesung und Übungsaufgaben aufeinander abgestimmt?

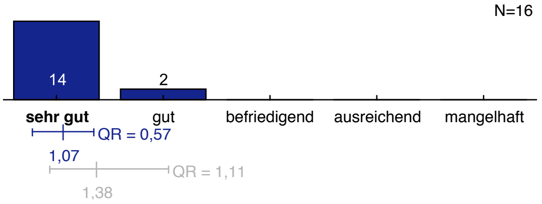

Wie beurteilen Sie die fachliche Kompetenz Ihres Übungsleiters / Ihrer Übungsleiterin?

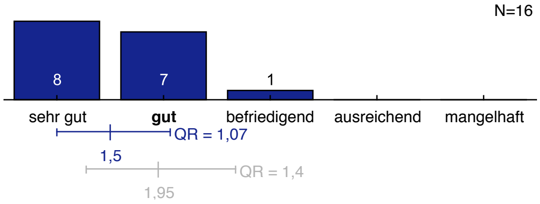

Wie beurteilen Sie die didaktische Kompetenz Ihres Übungsleiters / Ihrer Übungsleiterin?

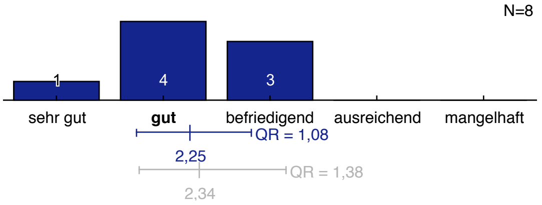

Wie beurteilen Sie die schriftliche Korrektur Ihrer Übungsaufgaben?

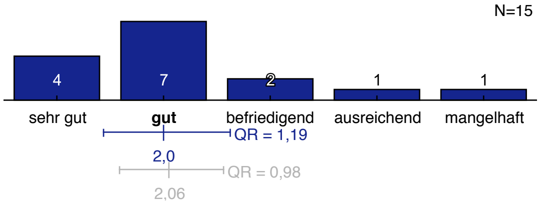

Wie beurteilen Sie die Übung insgesamt?

Erläuterungen zur Visualisierung

- Im oberen Teil des Bildes befindet sich ein Histogramm der absoluten Häufigkeiten. Hierbei ist die Höhe des Balkens proportional zur Anzahl der Nennungen.
- Darunter sind die möglichen Antworten abgetragen. Die Median-Antwort ist durch eine erhöhte Schriftgröße gekennzeichnet.
- Im unteren Bildteil befinden sich zwei gleichartige Visualisierungen von Median und Quartilen. Die obere, blaue Grafik kennzeichnet die Werte dieser Veranstaltung, die untere, graue diejenigen der Vergleichsgruppe.
- Als Vergleich dienen alle Veranstaltungen dieses Semesters, bei denen diese Frage gestellt wurde.
- N ist die Gesamtzahl der Nennungen

Semester: WiSe 2013/2014
Nr. Lehrveranstaltung: 20-00-0354-pr
Name: Drahtlose Sensor Netzwerke / Wireless Sensor Networks Lab
Anmeldungen: 17, Bestätigt: 17

Die Gesamtzahl der abgegebenen Fragebogen: N = 13

1. What is your course of Study?

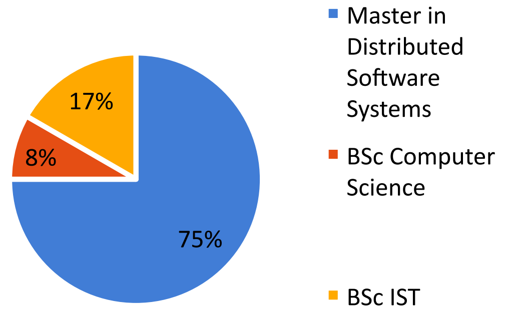

2. Which semester are you in?

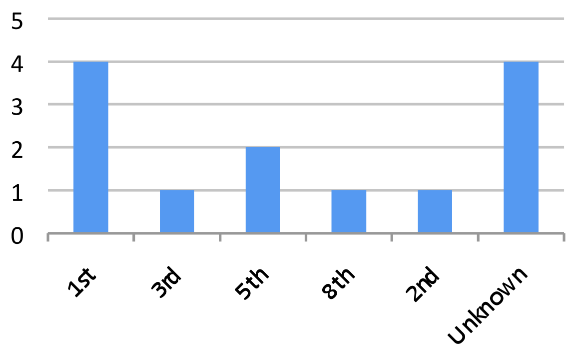

3. Have you attended the lab previously?

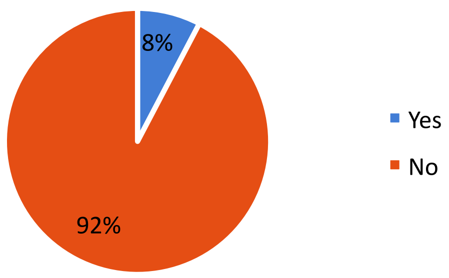

4. How many hours per week have you invested in addition to the lab sessions?

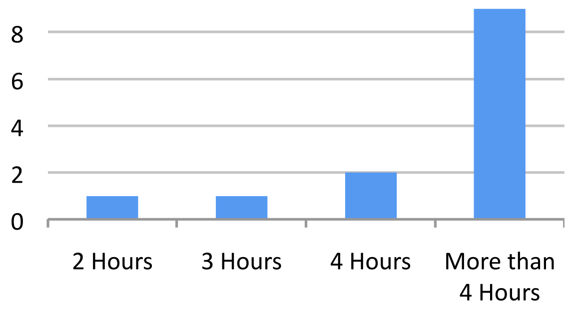

5. How would you rate the commitment of the lecturers?

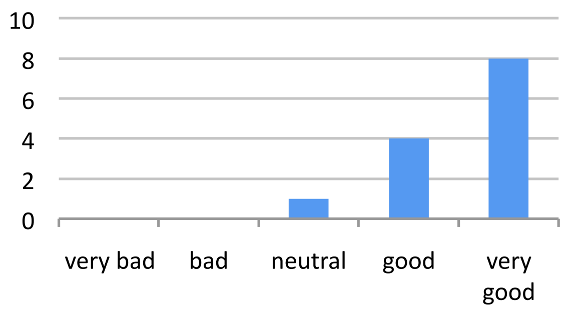

6. How well did the lecturers explain lab material?

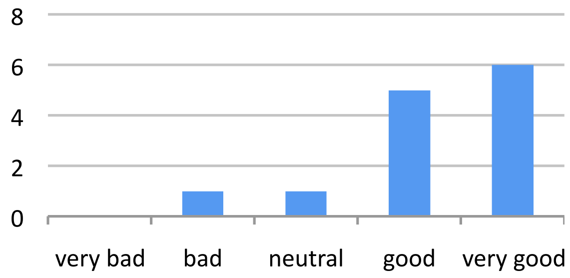

7. How well do the lecturers answer questions by the students?

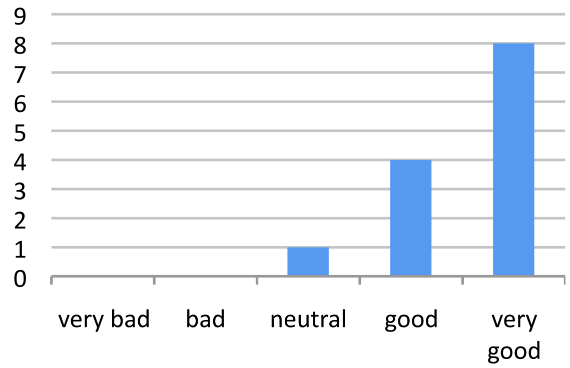

8. How would you rate the speed of the lab?

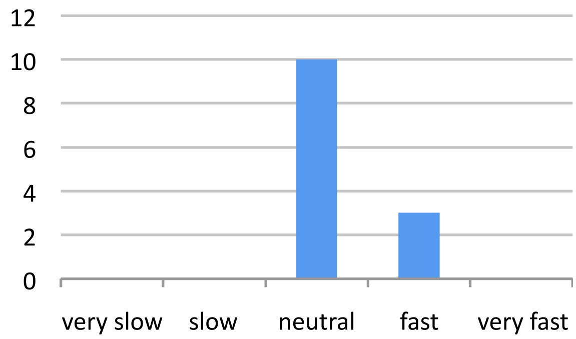

9. Were written material available for the lab?

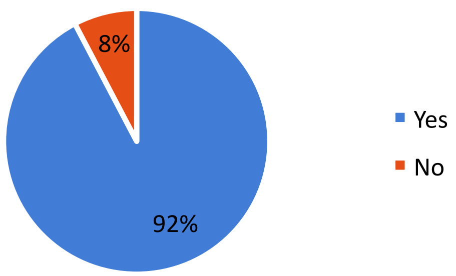

10. How would you rate the lectures in the lab?

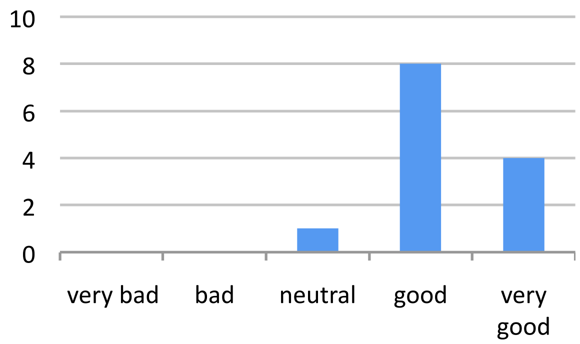

11. How would you rate the final project?

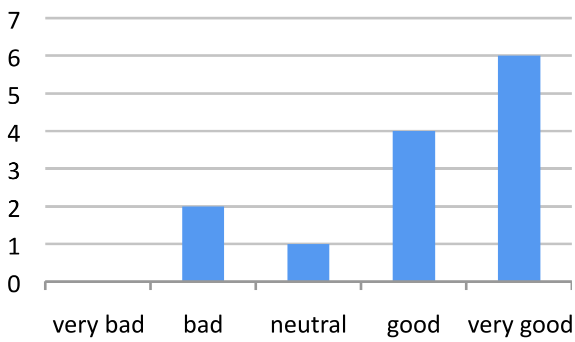

12. How would you rate the lab overall?

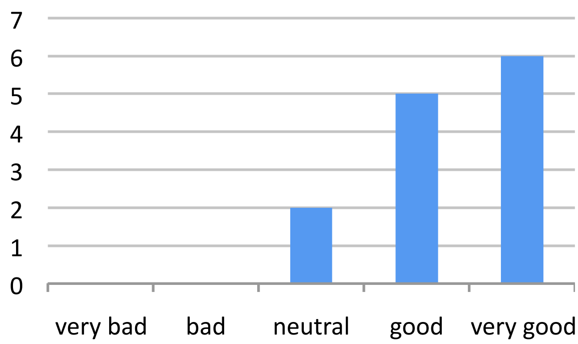

13. Other feedback or comments:

Please try to have more coding from the part of students. We were given many important parts of the code and we were just asked to fill in the rest of the parts.

I would like to recommend to reduce the size of the project group.

It would be nice if the final project gets over before other written exams as it really hampers other subjects. Other than that I had a pretty good time in the WSN Lab. It was very intresting and nice.

+ Very interesting topics. + I liked the exercises
- Project not that interesting - Not well planned - API and details announced too late.
- Work in group didn't work -> individuals didn't participate

It is nice that groups are created Randomly!
Some times I felt there was lack of Clarity upon what we are supposed to do.

This lab have been a great interest for me. The supervisors were really helpful everytime. The final project was also a great learning for me.
Overall I enjoyed attending this lab

Mir würde lieber dass die Versuche in Gruppen gelöst werden, dafür aber sollen die schon umfangreich auffallen. Grund dafür ist, dass man ohne hin der Projekt in eine Gruppen lösen muss. Dadurch könnte das Einspielen der Gruppen im Vorfeld getestet werden. Ansonsten es hat mich sehr gefreut das praktikum ablegen zu können.

Lab was interesting. Project was more interesting. There was full support by all team members and supervisors. Supervisor cleared our doubt on time. Sometimes confusion occurred because same information was not synchronized between supervisors. However, my doubts were cleared by supervisor very well on time. I enjoyed doing Project.

Detecting Leisure Activities with Dense Motif Discovery

Eugen Berlin and Kristof Van Laerhoven
Department of Computer Science
Technische Universität Darmstadt
{berlin,laerhoven}@ess.tu-darmstadt.de

ABSTRACT

This paper proposes an activity inference system that has been designed for deployment in mood disorder research, which aims at accurately and efficiently recognizing selected leisure activities in week-long continuous data. The approach to achieve this relies on an unobtrusive and wrist-worn data logger, in combination with a custom data mining tool that performs early data abstraction and dense motif discovery to collect evidence for activities. After presenting the system design, a feasibility study on weeks of continuous inertial data from 6 participants investigates both accuracy and execution speed of each of the abstraction and detection steps. Results show that our method is able to detect target activities in a large data set with a comparable precision and recall to more conventional approaches, in approximately the time it takes to download and visualize the logs from the sensor.

Author Keywords

activity detection, motif discovery, psychiatric monitoring

ACM Classification Keywords

H.5.m Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Algorithms, Design, Experimentation, Measurement.

INTRODUCTION

The automated recognition of the user's activities has been suggested for over a decade as an attractive system feature in pervasive computing literature. By relying on observations from sensors that are deployed in the environment of the user, or worn on his or her body, knowledge of recognition activities can be extracted. This technology is motivated by establishing a more effective dialogue between user and computer, reducing cognitive load in pervasive computing scenarios, or delivering an improved service by proactively responding to given situations. Numerous applications have been suggested to benefit from activity recognition: The authors of [20] for instance demonstrate how Activities of Daily Living (ADLs) can be detected, which are used in estimating the quality of self-care for elderly users. Further application scenarios for activity recognition include detecting office activities [17], maintenance tasks performed by engineers [21] and specific sports activities [8], finding appropriate advertising based on the user's physical activity [18] and eating and drinking activities [1]. Depending on the application, algorithms can go beyond recognition of activities and detect certain characteristics, such as the number of counts for selected gym workouts [3].

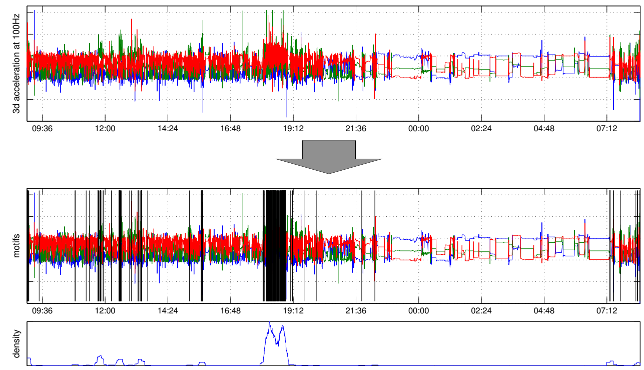

Figure 1. Day-and-night recordings from an unobtrusive sensor (top) are in this paper's approach analyzed for certain leisure activities. Salience detection upon motif discovery (bottom) is used to find typical activity patterns (black marks) as supporting evidence for the activity.

This paper's activity recognition approach is motivated by an application scenario that is relatively new: Psychiatric patient monitoring aims at characterizing both mood and behavioral trends by recording activity data over a period of typically *several months*. Current commercial solutions[1] are able to detect sleep and wake cycles for such long deployments, and come with tools for facilitating the recording of certain physical activities. In this scenario, a few of the general problems in activity recognition become trivial to solve: Patients already detail their activities in diaries so supervised learning methods can be employed, and only a few key leisure activities are of interest among the logged data. Other requirements, however, form new challenges: Sensors need to record for long stretches of time, the large amount of logged data needs to be analyzed fast enough, and detection needs to be robust against a deluge of background data.

[1] e.g., the Actiwatch: `http://www.camntech.com/` [3/2012]

UbiComp '12, Sep 5-Sep 8, 2012, Pittsburgh, USA.

Figure 1 illustrates the dense motif discovery method that forms the basis of our detection approach. Occurrences of so-called *motifs* are searched in the data, which are then used to substantiate the presence of an activity within the data. Motifs are discovered at training time by abstracting the raw acceleration samples in function of sequences of peak patterns, and efficiently searching for such similar patterns by means of a data structure called suffix tree. This search can be implemented in linear time, and performs a strong abstraction in comparison to traditional acceleration features (such as mean, variance, or FFT coefficients), at the cost of a processing overhead. Classification is then implemented with a straightforward bag-of-words classifier. An extra advantage of this method is that the illustration of motif occurrences in the time series allows for visual inspection of the activity recognition before the classification step. The main contributions of this work are threefold:

- We use a specific mood monitoring scenario as an activity recognition application with interesting constraints that impact both sensor and the software system design.
- A deployable system has been built consisting of a minimally invasive, wrist-worn sensor that is able to last for two weeks on a single battery charge while recording acceleration samples at 100Hz, and a data analysis tool that can efficiently process the recorded data.
- A novel detection approach is suggested that is suitable for classifying large amounts of long-term activity data, relying on local shape features within the acceleration signal and dense motif discovery.

The remainder of this paper is structured as follows: First, a psychiatry monitoring scenario motivates the need for the proposed fast and accurate detection of when a user performs physical leisure activities. Then a section is dedicated to related work in activity recognition, with particular focus on methods that aim for long-term deployments, and other motif discovery research in particular. The next section will then go into details on the different design choices and steps that constitute our method, such as the linear abstraction of inertial data and the use of suffix trees for finding motifs. An experiment then presents results on several long-term datasets how fast and robustly the chosen activities can be recognized among a large amount of daily activity data. The paper is wrapped up with the conclusions section enumerating the key findings of this paper, as well as the future research potential.

MOTIVATION: PSYCHIATRIC MONITORING SCENARIOS

Research in mood disorders relies frequently on the patients' self-reports, as well as semi-structured interviews with a psychiatrist, both during diagnosis and therapy. Emerging work with actigraphy tools and activity measurement in psychiatry [28, 24] has started to deploy wrist-worn sensors in conjunction with these tools that are recording the activity intensities observed for the patient over intervals of several seconds to minutes at a time. Such studies have found to be valuable in a range of mood disorder studies, such as attention deficit hyperactivity disorder (ADHD) and bipolar disorder [6].

Characterized by severe mood swings between manic or hypomanic, mixed, as well as depressive episodes, it is important in the diagnosis of a bipolar disorder to observe the patient's activity behavior over multiple weeks to months at a time. For mania for instance, energy levels tend to be high and activities tend to be performed in an interleaved fashion or especially vigorously (e.g., performing sports exercises longer without breaks). Similarly, depression tends to show in lower activity levels or the way patients behave during key activities, from not performing them at all or sparsely, to not fully completing them. Apart from daily activities such as sleep and food intake, especially physical and leisure activities are very likely to be impacted: Patients might for instance stop playing tennis when depressed, or vigorously practice for several hours in a manic episode.

In collaboration with psychiatrists specializing on actigraphy and ambulatory assessment of bipolar disorders, several interviews were held to list the basic requirements and expectations that an activity recognition method should adhere to. These were grouped in three categories that are important factors to consider when designing an activity recognition system for this field:

- **Supervised learning.** Patients are normally interviewed at regular intervals of several weeks, and provide in many cases log entries where they report on performed tasks and their mood assessment. Recent actigraphy systems already combine this information with the sensor data, so that in a similarly-developed activity recognition approach this can be used as an approximate annotation to train a patient-specific classifier.
- **Week-long, 24/7 data.** It is crucial that data is continuously captured at all hours of the day, as patients that go through depression or manic episodes are known to perform activities at irregular times, including night time. As a result, the sensor units need to be robust and power-efficient enough to keep recording for such long timespans without breaks, and the amount of data that need to be recorded will be substantial to process.
- **Leisure activities.** The number of activity classes that need to be recognized is relatively small (often 1) and can be determined by the psychiatrist during the first phases of diagnosis. This makes it easier for patients to keep track of what activities were performed, and this also impacts activity recognition, since only few activities need to be detected amongst a large amount of background data that might produce false positives.

This paper focuses first and foremost on a practical capturing and detection method that is able to recognize particular activities, and this within large datasets that tend to include a massive amount of background data, generally holding weeks of activity data at a time. The next section will review some of the literature on activity recognition methods that tackle similar problems, as well as technically related approaches in wearable sensing, data mining, and classification, and situate the proposed approach among peer research.

RELATED WORK

Activity recognition has been suggested as a promising tool before for bipolar studies. Both [23] and [27] have pointed out that the use of automatically monitoring activities would be a useful tool to support the diagnosis of bipolar disorder and detect onsets of depression and mania. In particular the so called Hamilton Depression Scale (HAMD) and Bech-Rafaelsen Mania scale (BRMS) [2] tools contain elements where physical activities are of considerable interest. However, to our knowledge no research has thus far focused on implementing an activity recording method that can practically be deployed on a patient's wrist for a week and allows almost-instant analysis at the psychiatrist's office.

A significant amount of work in the context of activity recognition has focused on automatic feature selection for inertial data and using strong classifiers upon these features to detect activities. Common candidates that have proven worthwhile in previous studies (e.g., [9], [13]) have found basic statistics, in particular mean and variance, and frequency-based features (FFT and Cepstral Coefficients, spectral entropy and energy) over a sliding window to be distinctive features to characterize. Lester et al. [13] use in a combined discriminative-generative classification approach the AdaBoost algorithm to automatically select the best of these features and to learn an ensemble of static classifiers to recognize different activities. Strong classifiers that have proved valuable in activity recognition include Naïve Bayes, Bayesian Networks, Hidden Markov Models (HMMs) or Support Vector Machines (SVMs) [1, 3, 8, 13, 15, 17, 18, 19, 20, 21].

The use of motif discovery has been suggested as an alternative approach in activity recognition that is especially useful when a fully supervised method is not feasible, or when short characteristic gestures need to be spotted that are hard to annotate individually by the system's users. Minnen et al. [16] use motifs to automatically discover gym work-out gestures in inertial data recorded form body-worn sensors, by mapping the sensor data to symbols and using a suffix tree to search efficiently through the resulting large symbolic strings. Similarly, Hamid et al. [7] analyze activities in an instrumented kitchen, and [26] uses motif discovery to detect activities such as walking and falling without supervision. We use motif discovery primarily because (1) the annotations that describe which activity was done when are provided by the system's wearer using self-recall and are thus only approximate, (2) we assume that physical leisure activities can typically be characterized by occurrences of certain short gestures, and (3) because it is an especially fast method that allows parsing of large data sets at once.

Motif discovery techniques generally rely on symbolic abstraction of the original raw sensor data to obtain an especially fast detection method using the suffix tree representation. In [22], the symbolic representation of inertial data is used to facilitate efficient matching of motion patterns. The inertial trajectory in space is approximated after which it is mapped to a character based on the minimum angular distance to the 3 axes that are represented by a small alphabet of 6 symbols, thus resulting in a motion string. The

Figure 2. Our custom-made sensor platform, designed to be worn and used as a wrist watch, is light-weight enough to be worn for recordings of up to 2 weeks. Left: board with controller, accelerometer, micro-SD storage, and USB connector for data access and charging the battery.

approach of [16] uses discrete mapping based on a Gaussian distribution fit on the data, whereas others use probabilistic approaches such as [4], or approximate the sensor signal's time-series first by piecewise constant segments of fixed length [14], which are then mapped to a set of discrete symbols. Our approach uses similarly an approximation of the inertial time-series, but uses for the mapping the subsequent segments' slopes to capture the essence of these short gestures' patterns in accelerometer data from the wrist.

The importance of long-term recording of inertial data, in an unobtrusive manner, has been stressed in several key publications on activity recognition (most notably [5, 10]). Although data sets have been recorded over similar time frames as in this paper, none so far have recorded day and night for several days consecutively. Actigraphy on the other hand does log for extended periods of time, but abstracts the inertial data on-board the sensor and does not retain the original time series at the resolution of this paper's (100Hz).

DATA LOGGING PROTOTYPE

Deploying sensors that are lightweight and wearable is one of the hard challenges in the creation of robust recognition systems, as has been identified in previous research such as the Mobile Sensing Platform [5]. Since there are no current off-the-shelf platforms that allow continuous logging of high-resolution accelerometer data, the experiments in this paper have been recorded from custom-built sensor units that measure and store regularly-timestamped 3D acceleration on 2 Gigabytes of local flash storage (Figure 2). This persistent memory is required since the recording is done by sampling 3D acceleration data at 100Hz. In addition to the accelerometer, light and temperature sensors are also on board.

While recording, the sensor regularly acquires 33 equidistant (10 ms) samples at a time from an ADXL345 accelerometer's first-in first-out (FIFO) buffer and logs these using run-length encoding (RLE) on the on-board micro-SD card. Regular time stamps are produced by a precise real-time clock (RTC) unit embedded in the PIC18F46j50 microcontroller, allowing detailed verification of separate annotations taken by the subjects and individual sensor readings. Between FIFO reads, only the accelerometer is active on the board, with the micro controller in sleep mode preserving the charge of its miniature $180mA$ Li-Polymer battery.

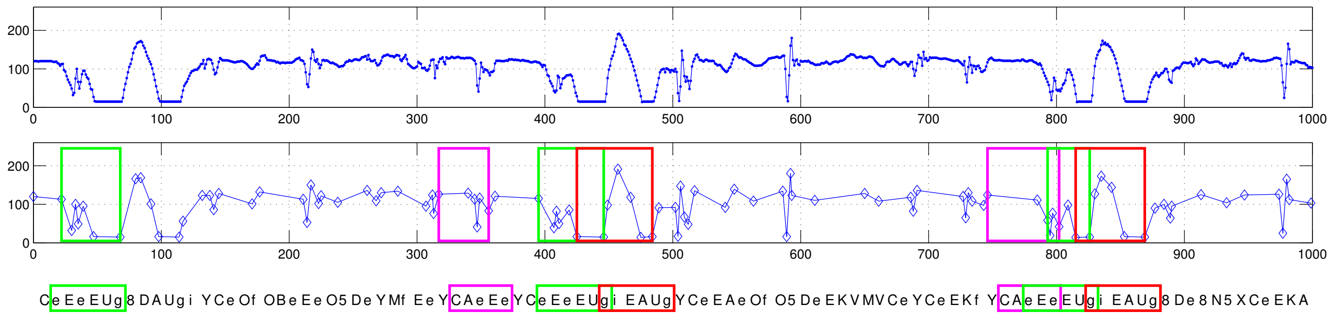

Figure 3. The raw 3D 100 Hz inertial data (top plot) are transformed by a piecewise linear approximation algorithm into segments (bottom plot) that preserve the shape of the signal to facilitate storage and analysis. The segments are subsequently abstracted in discrete symbols (bottom row) to allow fast discovery and matching of motifs. Occurrences for three motifs are highlighted by colored boxes; Note that they can overlap and vary in length.

The entire sensor draws $480uA$ on average while logging. Non-stop logging of 100Hz inertial data should thus last for 15.6 days. From our own experiments, we have observed that the battery is usually drained after 14.5 days, which still allows 2-weeks of uninterrupted logging. To meet the requirements of long-term 24/7 deployment, the unit is packed in a custom shock-proof case and provided with an anti-allergic textile wrist strap. For more inconspicuous deployments, as advised by our collaborators from psychiatry, an OLED display has been added to display the current time from the real-time clock. The display is by default turned off and activated by double-tapping the watch, driven by an on-accelerometer function interrupt. With the OLED attached, deployments generally last one to two days less.

DENSE MOTIF DISCOVERY

This section gives an overview of the search and selection procedure for motifs from raw inertial data, and motivates the use of *dense* motifs. A set of early abstraction steps of the accelerometer data, together with a search-optimized data structure called suffix tree, guarantee that searching through weeks of data becomes feasible on standard computing hardware, and that classification can be done almost simultaneously with the downloading of the prototype's data.

Method Overview

Motif discovery refers to the search for recurring sequences or patterns within a data stream. For this to be applicable in real-world scenarios, previous research has identified several techniques to represent the original data, which often tend to be noisy and hard to match exactly, in a discrete symbolic string. This paper's approach implements a discrete mapping that applies a two-step abstraction process while aiming to characterize patterns in the data (i.e., potential activity-specific gestures) by the shape of the time series.

Figure 3 illustrates how the proposed method transforms inertial data to a string that facilitates the finding of recurring motifs: The original data consists of 3D accelerometer samples taken in equidistant 10 ms steps. Sets of linear segments are created from these, using an online approximation algorithm that minimizes the residual error between original data and segments. The segments are discretized into symbols using the slopes of connected linear segments. Using a suffix tree representation of the target activity's training data, a set of motifs is found using adaptive length thresholds. Motifs that also have occurrences in the training's background data, i.e., the vast set of data that does not belong to the activity, are withheld. As a *dense* set of motifs is trained for, classification is done by searching new data for time windows in which motifs from one particular activity are frequently occurring. This is implemented with a bag-of-words classifier which uses the detected motif occurrences as evidence. To perform the described motif discovery efficiently, the original sensor data needs to be abstracted: The next section will discuss a linear segmentation step.

Approximation from Raw Data

The first abstraction step is crucial from a practical and efficiency point of view: Raw sensor data is sampled at a relatively high frequency (100 Hz) to capture the essence of the gestures and typical motions performed by the sensor's wearer, but this also means that analysis of larger data sets quickly becomes challenging. Even with fast and lossless compression techniques such as run-length encoding, an entire day worth of data typically contains millions of 3D acceleration samples.

We argue that for motif discovery in inertial data, primarily the shape of the acceleration time series is important to retain. The applied technique to reduce the amount of data on the one hand and to preserve the shape of the signal on the other, belongs to the Piecewise Linear Approximation (PLA) family of abstraction algorithms. We used in our algorithm a modification of the original Sliding Window and Bottom-Up (SWAB) algorithm [11] that has been verified to perform well on body-worn accelerometer data [12].

The transformation from raw data to linear segments consists of two steps, the approximation of data on a sliding buffer window and filling the buffer with new sensor samples. The main approximation step is carried out by a Bottom-Up brute-force algorithm that produces the linear segments by merging cheapest adjacent segments until a preset threshold is reached. The leftmost segment is output as a result, and the buffer is filled with new data, whereby the modified version

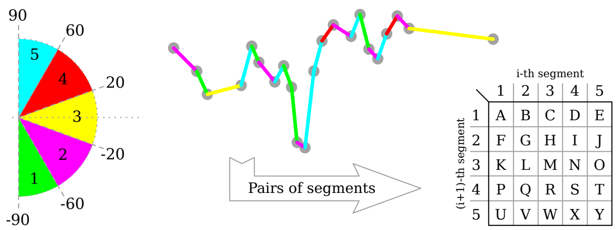

Figure 4. Mapping from linear segments to symbols: The slope range is divided into bins for a given number of separation points which are computed based on the training data segments' slopes histogram. The segments' corresponding bin numbers are then used as indices for the symbols matrix. Sliding through the segmented time series while considering two neighboring segments will thus result in a symbolic string.

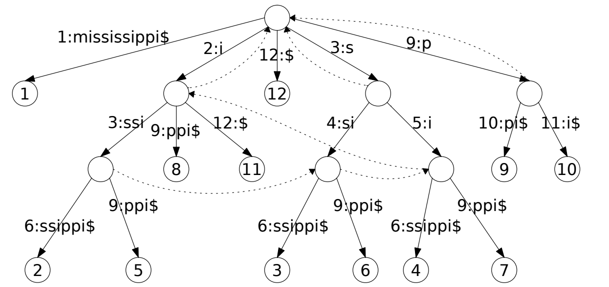

Figure 5. Generalized suffix tree for the string `mississippi` created by adding a unique terminator character $ to the original string. Suffix links are indicated with dotted edges, edge labels give the first occurrence position in the string of subsequent suffixes.

considers slope sign changes, and thus peaks in the signal, in incoming raw data. More details and a study on its efficiency can be found in [11]. The two plots in Figure 3 show a 10-second accelerometer time series and the resulting piecewise linear approximation produced by the modified SWAB algorithm. The approximation was performed per acceleration axis for implementation reasons, although it is also possible to approximate multidimensional time series.

Mapping to Discrete Symbols

After abstracting the raw acceleration data to linear segments, a discretization step is used to obtain a symbolic string representation of the original time series. This abstraction step is first and foremost required to enable fast discovery of motifs, but also for finding matches between motifs.

First, we evaluated two degrees of freedom per segment, considering the length and the slope, and mapping the resulting segments onto symbols in a similar way as was done with the SAX approach by Lin et al. in [14]. Our approach also discretizes the feature value space based on the distribution of the values. The main difference to SAX is that the first abstraction step produces constant segments of fixed length, thus having only one degree of freedom, while SWAB produces linear segments with individual slope and length. With this approach, our initial test showed that very long segments became over-represented in the motif discovery. This is due to inherent properties of accelerometer data, with long segments with a slope close to zero being over-represented, particularly during the night time and sedentary tasks, where little or no changes are present in the signal.

Being interested in mainly short and characteristic gestures, focus went to the slopes of two neighboring segments, whereby we use the angular representation of the slope defined as $\theta = \arctan(m)$. To achieve discretization, the slope range from -90 to 90 degrees was divided into bins, whereby the borders (quantiles) are selected on the basis of representative data in a histogram during the training phase. To avoid overrepresentation of non-motion motifs, segments with a slope close or equal to zero were not considered. The rest, where we do not assume Gaussian distribution, is used to compute the quantiles for a given number of bins (which was found to produce the best results when set to 5).

Mapping the linear segments to discrete symbols is realized by sliding through the time series, considering the slopes of two neighboring segments at a time, and converting them to one character (cf. Figure 4) using a 2 dimensional matrix. Converting an approximated time series using this approach will result in a long symbolic string, as shown in Figure 3, that can be parsed for motifs with the help of suffix trees. The advantages of this approach are two-fold: First, the length of a linear segment is not constrained to a fixed value, as it is the case with the SAX approach, and common errors where symbols afterwards would need to be merged are avoided. Secondly, by not taking the length of a segment into account when mapping the segments to symbols, more importance is placed on patterns in the data where strong peak sequences occur. With a symbolic representation of the time series now completed, the next section will discuss the method for the finding of motifs.

Extracting Motifs by means of Suffix Trees

Having mapped the raw acceleration data to a symbol sequence, motif discovery can now be done by finding substrings that occur multiple times in the target class. This is above all an efficiency problem: searching for all occurrences of every motif in a long string in an exhaustive fashion will result in a slow discovery process that is not scalable, as large sets of motifs are expected to be present.

To significantly speed up this search procedure of motifs, a technique is applied that transforms the string from a long array of symbols to a tree representation. This data structure, called *suffix tree*, generally requires more storage space than the string array, but in return allows searching for all substrings up to a certain length in linear time. Furthermore, suffix trees can be constructed in linear time using an algorithm by Ukkonen [25]. The motifs are found by checking the number of leaves for all suffixes up to a certain depth in the tree, which then corresponds to the number of the substring's occurrences in the data.

Figure 5 depicts the generalized suffix tree for the string `mississippi`. The generalized suffix tree is produced by adding a unique terminating character (such as commonly used $ or #) to the original string. With a generalized suffix tree created, this structure can be used for a multitude of different applications. The most common application is searching for query substring occurrences, for example, those of the substring `issi` in the example above: First, verifying whether the query is present in the original string at all can be answered by traversing the edges `2:i` and `3:ssi` of the tree from its root. The fact that this path can be traversed, means that the query is present in the original string. In this case, the places and number of occurrences are found by counting the leaf nodes in the sub-tree and looking at the leaf node indices: 2 occurrences with positions 2 and 5 are found after traversing the edges `6:ssippi$` and `9:ppi$`.

Suffix trees are used for the discovery of motifs that are likely descriptors for a target activity class. Motifs are found by searching the suffix tree to a certain depth, and accumulating those motifs that occur at least two times and have a minimum length (a trade-off evaluation on our data set showed a minimum length of 5 to still produce sufficient motifs for the bag-of-words classifier). The most discriminant motifs are then selected for classification, from all discovered ones by removing those that appear frequently in the background data provided during the training. After thus finding a set of motifs that tend to represent a particular activity class, these can be used as weak detectors in classification by evaluating the density of their occurrences.

Bag-of-Words Classification

Using the most discriminant motifs during a training phase as described in the previous paragraphs, classification is performed by local evidence of all motifs that support an activity. This straightforward bag-of-words classifier uses a sliding time window over the time series and accumulates local evidence by counting occurrences of motifs. As the activities tend to last at least 30 minutes and up to an hour and a half, a window size of 10 minutes was chosen.

EVALUATION

The entire approach as described in the previous section is tested in this section under conditions from the motivation scenario of psychiatric monitoring. After presenting the wrist-worn sensor prototype, as well as the test subjects and the chosen activities, a comparison of the proposed approach is given with two common activity recognition techniques that have been chosen as a benchmark. Finally, we are discussing the performance results of our and the other methods.

Participants and Target Activities

The data used in the following experiment comes from a group of volunteers who have no known psychiatric disorders and for whom a leisure activity was known before the recording phase (a key leisure activity, regularly performed as it would be chosen by a psychiatrist), which they would do once each day, for a full working week. For most, this turned out to be a leisure activity, for some a daily activity that was part of their regular schedule. Table 1 gives an overview of all participants, specifying their gender, age and their personally chosen target activity which will be used for testing detection accuracy. Additionally, the amount of raw sensor data as well as the total number of segments used in the evaluation are given. Here, the first data reduction step (modified

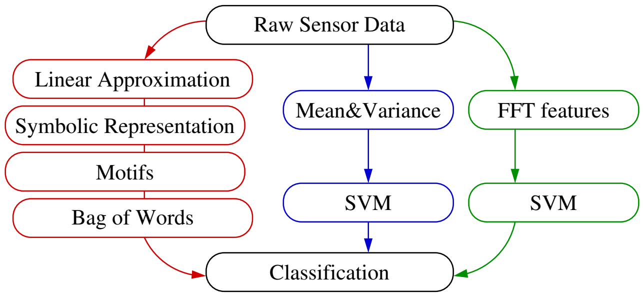

Figure 6. Overview of the detection evaluation: The dense motif classifier (red) is compared with two strong classifiers that rely on mean and variance features (blue) and FFT-derived features (green) respectively.

Table 1. The list of participants, specifying their gender, age, and the leisure activity they performed once a day, along with the number of recorded data samples and the number of linear segments produced from raw data by the segmentation process.

subject	gender	age	target activity	3D samples	segments
1	female	30	zumba	26.927.159	2.011.826
2	male	35	cycling	42.841.897	2.259.414
3	male	30	badminton	44.244.417	2.758.001
4	female	27	guitar	43.230.164	2.825.311
5	male	28	gym	34.822.499	2.480.707
6	female	26	flamenco	43.101.537	2.980.562

SWAB algorithm executed with approximation threshold of 10 and buffer size of 80) is also shown to have a significant effect, resulting in more than thirteen times less data points.

The data set from each participant was split into separate blocks of about a full day (24 hours ± 50 minutes) each to facilitate 5-fold cross validation. Each activity instance lasted approximately one hour. The target activity thus holds $\pm 5\%$ of the entire fold, with the rest being other daily activities.

Benchmarking the Performance

To evaluate the classification performance of our approach, a comparison to two standard activity recognition techniques was done. For the latter, several classifiers were identified, with the Support Vector Machine (SVM) as the best performing, as well as different feature sets to abstract the raw data. Due to their coverage in the activity recognition community, e.g. in [9] or [29], mean and variance were identified as one combination. An additional set of features based on Fast Fourier Transform (FFT) coefficients were chosen as another: the 16 FFT features that have been suggested and evaluated in [29] consist of the absolute, real valued FFT coefficients grouped into 4 logarithmic bands, 10 Cepstral Coefficients, the spectral entropy and energy of the signal.

One imbalance in this comparison is illustrated in Figure 6: Since the dense motif approach aims at extracting characteristic motion patterns for target activities from the symbolic representation of the original sensor data, more resources are spent on pre-processing the sensor data, and less on the classification. Although Figure 6 details just the required steps, and not their time complexity, it is clear that the approaches differ significantly in how the processing steps are weighted.

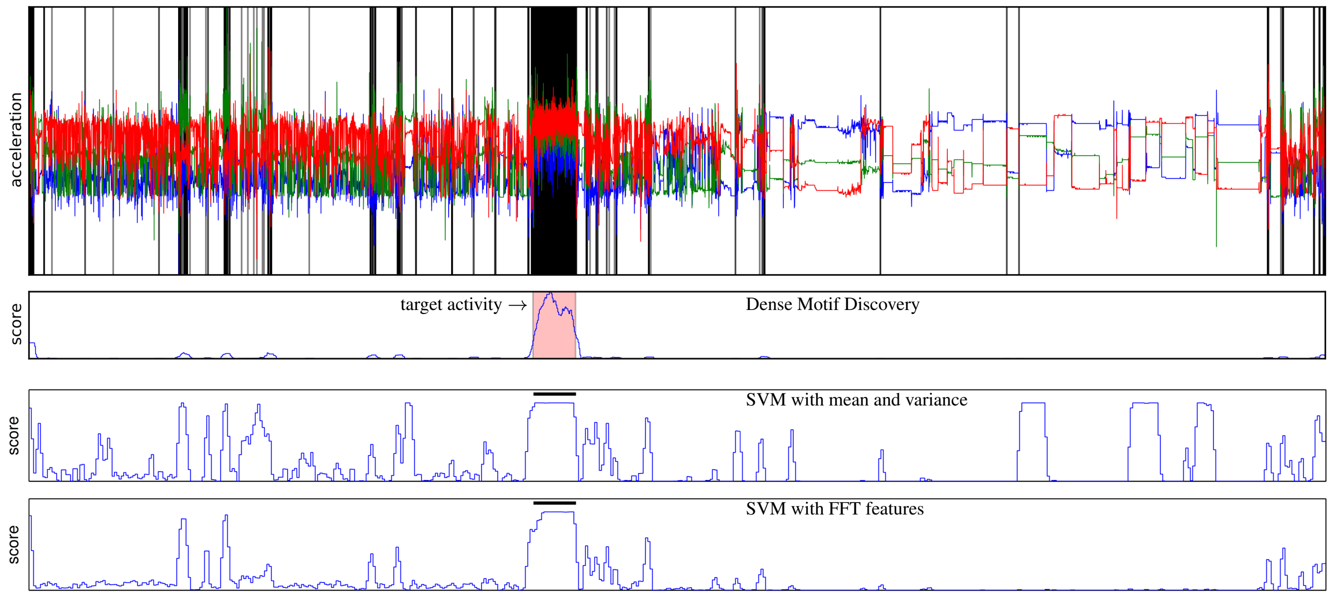

Figure 8. One day of the experiment data in which one of the participants cycled for about one hour. The topmost plot shows the original 3D sensor data, along with motif occurrences highlighted by black markers. The three plots below give the corresponding score plots produced by the different classification approaches during the evaluation: The first plot shows aggregated motif occurrences, while the two plots below show the smoothened SVM classification for mean and variance and FFT-based features respectively, with all three approaches using the same sliding window length. After combining all such results for all participants' data, the precision and recall figures show overall performance of the three approaches in Figure 9.

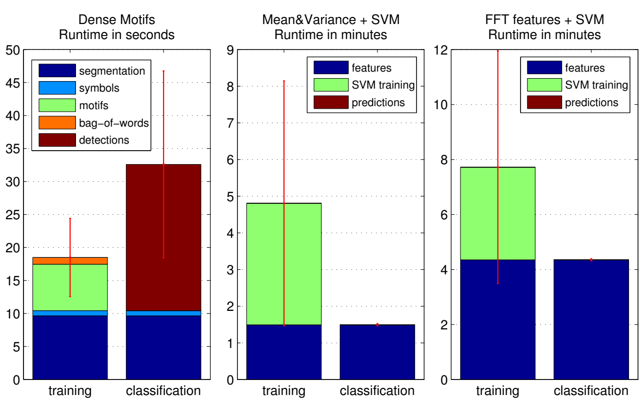

Figure 7. Mean execution times for training and detecting on 24 hours with the three approaches, with upper and lower quartiles (red lines): The dense motif method is especially faster in training, with segmentation and discovery of motif occurrences taking up most of the time. For the SVM-based approaches, most of the time is spent on calculating the features on the sliding window, with the classification done in a few seconds. Used parameters are the same as in the classification analysis.

Figure 7 shows the average times gathered during the 5-fold cross validations with our dense motifs approach for the best-performing set of parameters (as also shown in later evaluation plots, approximation threshold and buffer size: 10/80; symbols mapping to 5 bins). For one day worth of data the two abstraction steps (producing segments and converting them to symbols) require about 10 seconds. Depending on the activity, the time required for extracting the characteristic motifs from the training part of the dataset ranges from 3 up to 12 seconds. Obtaining motif occurrences for the classification on the fifth part of the dataset and computing the score needs from 18 up to 46 seconds, using a standard laptop setup and with the source code written in Python.

Both mean and variance, as well as the FFT-based features, are computed on a sliding window over the raw data, with window sizes varying from 1 to 30 seconds. For classification, the `svmtrain` and `svmclassify` methods from the Matlab Bioinformatics Toolbox were used. The performance of the features with the SVM classifier was evaluated by the same 5-fold cross validation as for the dense motifs approach. The detections produced by the SVM classifier are smoothened by a sliding window of 10 minutes to filter out outlier false detections, resulting in a score. At this stage, by evaluating the obtained classification versus the the ground truth annotations, precision and recall are computed for our dense motifs as well as for the features with SVM approaches. Figure 8 shows an example illustrating how the different classification techniques performed on the third day of the cycling dataset during the evaluation phase. The score plots below the raw data show the aggregated motif occurrences for the dense motif method, and the normalized results of the windowed filter after SVM.

Experiment Results and Discussion

This section presents the experiment results for the leave-one-day-out 5-fold cross validations: For every activity one day is left out for testing purposes, while training (obtaining the motifs that tend to represent the activity) on the other four days. Since the evaluation considered a wide range of possible parameter combinations (abstraction thresholds, buffer lengths, window sizes, etc.), only a few prolific figures are shown to discuss the experiment results.

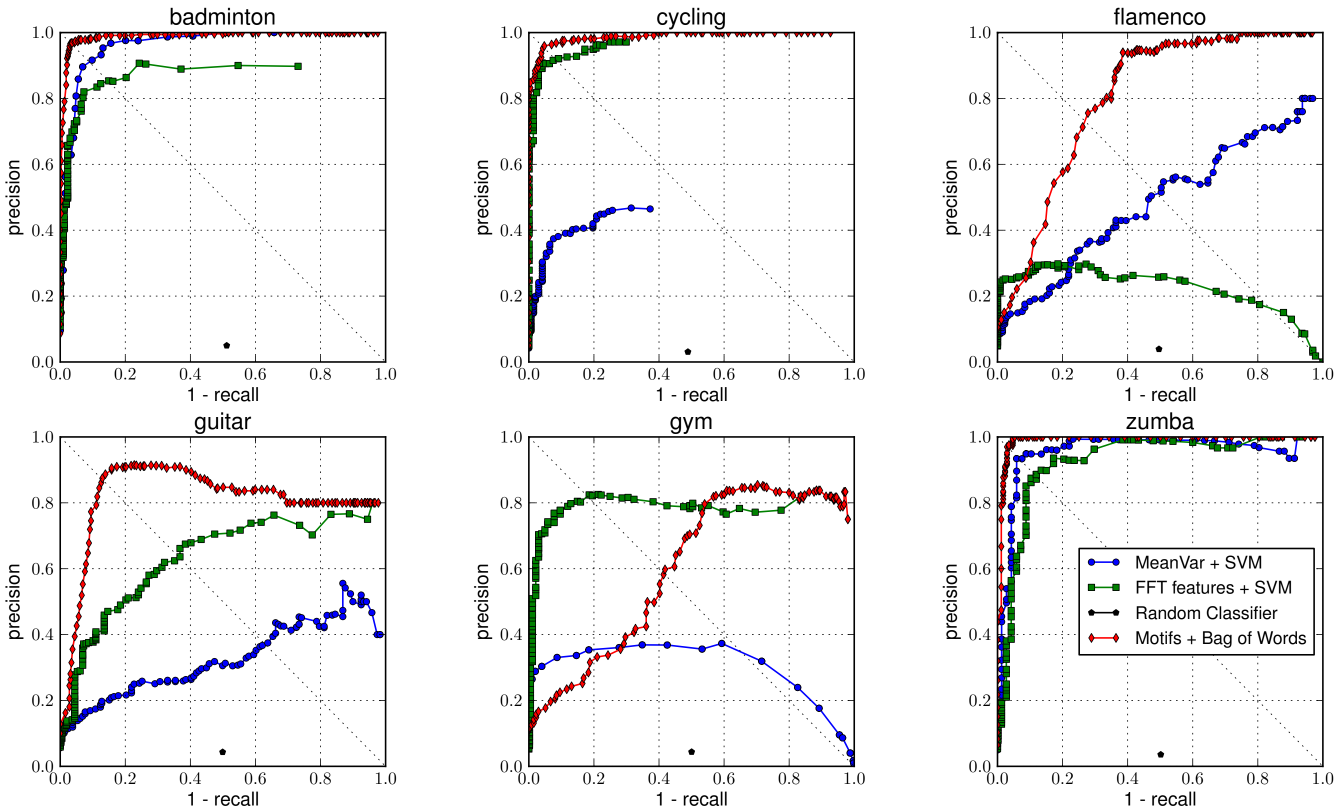

Figure 9. Precision and recall performance results on the six different activities obtained through the leave-one-day-out 5-fold cross validation, averaged over the number of folds. The dense motif approach outperforms the SVM classifier trained with mean & variance or FFT features on five out of six activities, while performing significantly worse to the FFT features trained SVM classifier on the gym dataset (fifth plot from the top-left).

Figure 9 shows a comparison of the best performing average precision and recall figures of our approach and the SVM classifier that has been trained with mean and variance, and the FFT features. Additionally, the performance of a random 'guessing' classifier is depicted for completeness. Precision and recall are averaged over the number of folds, while for each activity and classification method the choice of parameters with the best classification performance is chosen.

The SVM classifier trained with mean and variance features performs well on activities that involve a lot more motion, with especially the variance of the signal playing a significant role, as can be seen by comparing the activities badminton or zumba with gym, cycling or flamenco. While the first two activities exhibit very high accelerations due to sharp hand motions, the three latter activities lack such high accelerations. The dense motif approach is in many cases the best-performing, in some cases even significantly. To illustrate its strengths, the performance on the badminton data is shown in Figure 10, zooming in on a short time span of 50 seconds with motifs occurrences matching the underlying characteristic motion patterns. Motifs here often overlap, with their dense occurrences making the detection of the activity more reliable.

When investigating the impact of the abstraction, we noticed that the parameters that control the first abstraction step for

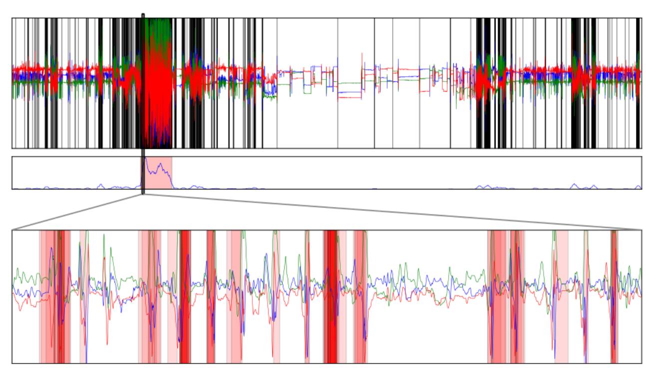

Figure 10. Dense motifs performance on one fold of the badminton data over one day (upper plot, with badminton activity marked in red), and a 50-seconds fraction thereof with motif occurrences (lower plot). Characteristic motions such as forehand, backhand, smashes, are often marked by motifs, while areas in between tend to be left out.

the badminton and zumba data have almost no impact on the classification performance. This can be explained by the importance of high acceleration peaks in the signal that are preserved even with a coarse grained approximation. The approximation threshold plays a more important role for the flamenco or cycling datasets, or generally for activities without extreme accelerations where coarse grained approximation results in a worsened classification. With the best per-

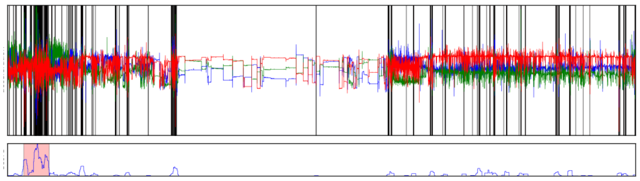

Figure 11. Example of dense motifs on a day with the flamenco activity (highlighted in red). Note the drops in the score for the target class.

forming parameters, the dense motif approach reaches over 95% in precision and recall on the badminton, zumba or cycling datasets. For the other datasets, namely flamenco, guitar and gym, the dense motifs approach needs a more detailed discussion.

The flamenco dataset (see also Figure 11) shows a wide performance difference between the chosen approaches. While mean and variance features give average results, FFT features perform surprisingly poorly. This could be explained by the fact that flamenco dancing is an activity with lots of irregular motions that are often complex. The dense motifs approach benefits the most from the characteristically short motions, which are not equally distributed over the whole activity, hence the slight drop in the recall. The equal error rate reaches 73% for the flamenco data set.

The dense motifs approach on the gym activity data was surprisingly low. The exercises consisted of different weightlifting workouts, with gestures being much slower compared to the other activities in this evaluation. Figure 12 shows the dense motif performance on 24 hours and a sub-sequence lasting for about 2 minutes. Inspection of different folds during the evaluation and comparison to other activities shows that the initial number of motifs is not very high in the first place, and is heavily reduced as motifs that appear in the background data are discarded (equal error rate of 60%). The slow motions are also the reason why the mean & variance trained SVM classifier fails at classifying the activity correctly. The FFT-based features, on the other hand, computed on a window of 5 seconds, were able to profit from the frequency domain characteristics of the gym exercise activity: The SVM classifier trained with these features performs considerably good reaching over 80% in equal error rate.

The playing guitar data turned out to be captured well with motifs, with the approach gaining a significant advantage over a classifier on the traditional features. While different ways to play were observed, including hitting or plucking the strings depending on musical genre, the dense motif approach still detected much of the activity, reaching 87% in equal error rate. Figure 13 illustrates one day, where the participant took a short break in activity. Such reduced performances due to the different ways to play the guitar, as well as breaks of varying durations, might be of particular interest to the psychiatric analysis and might provide additional hints regarding their patients' mood. Implementing detection for such events requires further investigations though and are left as future work.

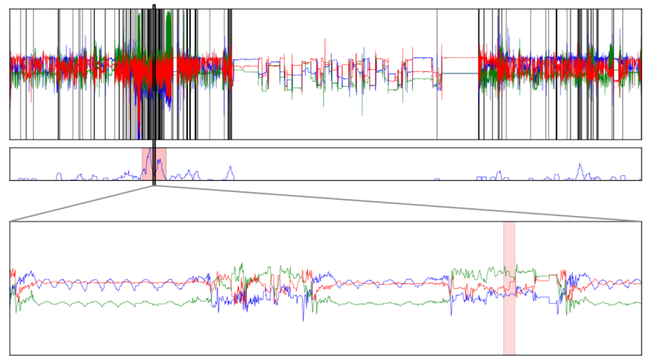

Figure 12. Dense motifs for a day of the gym data (upper plot, with gym activity marked in red in the middle plot) and a sub-sequence lasting 2 minutes (lower plot, with motifs marked in red). Two of the exercises can be recognized by periodically signals in the lower plot.

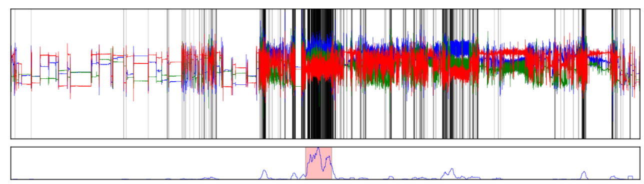

Figure 13. Dense motifs on one day with playing guitar for an hour: the gap in the middle of the activity (see bottom plot) was found to be due to a short bathroom break of the participant.

CONCLUSIONS AND FUTURE WORK

This paper presented a practical activity detection system to spot leisure activities in long-term datasets, that is based on finding parts in the data that contain frequent matches with a set of motifs. These dense motifs are discovered in exemplar training data by finding the most descriptive motifs for the activity against the large amount of background data. The approach has been designed for continuous deployment in psychiatry monitoring, and was evaluated on a data set with similar constraints, containing more than a month of data taken from a custom-built wrist-worn sensor unit that records 3D accelerometer data at 100Hz.

Experiments show that the approach is able to detect many physical activities, on par with standard approaches, reaching an equal error rate performance of 95% for 3 of the 6 activities, and only being significantly outperformed on one. It was demonstrated to be able to work on large and long-term sets of inertial data and can, unlike many traditional approaches, be expected to be scalable for weeks or months of such data. A remaining weakness identified is the method's reliance on short gestures, such that slower movements (such as weight lifting exercises) are not always picked as motifs.

While this work aimed at efficiency and therefore focused on extracting characteristic features, future work is underway to replace the bag-of-words classifier by more powerful models. The data set and code used in this paper are publicly available [2], to encourage reproduction of these results.

[2] at http://www.ess.tu-darmstadt.de, or by contacting the first author

ACKNOWLEDGEMENTS

This work was funded by the German Research Foundation (DFG) as part of the Emmy Noether Programme.

REFERENCES

1. O. Amft, H. Junker, and G. Tröster. Detection of eating and drinking arm gestures using inertial body-worn sensors. In *ISWC '05*, pages 160–163. IEEE, 2005.
2. P. Bech, T. Bolwig, P. Kramp, and O. Rafaelsen. The bech-rafaelsen mania scale and the hamilton depression scale. *Acta Psych. Scand.*, 59(4):420–430, 1979.
3. K.-H. Chang, M. Chen, and J. Canny. Tracking free-weight exercises. In *UbiComp '07*, pages 19–37, Berlin, Heidelberg, 2007. Springer.
4. B. Chiu, E. Keogh, and S. Lonardi. Probabilistic discovery of time series motifs. In *KDD '03*, pages 493–498, New York, NY, USA, 2003. ACM.
5. T. Choudhury, G. Borriello, S. Consolvo, D. Haehnel, B. Harrison, B. Hemingway, J. Hightower, P. P. Klasnja, K. Koscher, A. LaMarca, J. A. Landay, L. LeGrand, J. Lester, A. Rahimi, A. Rea, and D. Wyatt. The mobile sensing platform: An embedded activity recognition system. *IEEE Pervasive Computing*, 7:32–41, 2008.
6. P. Corkum, R. Tannock, H. Moldofsky, S. Hogg-Johnson, and T. Humphries. Actigraphy and parental ratings of sleep in children with adhd. *Sleep*, 24(3):303, 2001.
7. R. Hamid, S. Maddi, A. Bobick, and M. Essa. Structure from statistics - unsupervised activity analysis using suffix trees. In *ICCV '07.*, pages 1–8. IEEE, 2007.
8. T. Holleczek, J. Schoch, B. Arnrich, and G. Tröster. Recognizing turns and other snowboarding activities with a gyroscope. In *ISWC '10*, pages 75–82, 2010.
9. T. Huynh and B. Schiele. Analyzing features for activity recognition. In *sOc-EUSAI '05*, pages 159–163, New York, NY, USA, 2005. ACM.
10. S. S. Intille, K. Larson, E. M. Tapia, J. S. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson. Using a Live-In Laboratory for Ubiquitous Computing Research. In *Pervasive*, 2006.
11. E. J. Keogh, S. Chu, D. Hart, and M. J. Pazzani. An online algorithm for segmenting time series. In *ICDM '01*, pages 289–296. IEEE, 2001.
12. K. V. Laerhoven, E. Berlin, and B. Schiele. Enabling efficient time series analysis for wearable activity data. In *ICMLA '09*, pages 392–397. IEEE, 2009.
13. J. Lester, T. Choudhury, N. Kern, G. Borriello, and B. Hannaford. A hybrid discriminative/generative approach for modeling human activities. In *IJCAI '05*, pages 766–772. Morgan Kaufmann Publishers, 2005.
14. J. Lin, E. Keogh, S. Lonardi, and B. Chiu. A symbolic representation of time series, with implications for streaming algorithms. In *DMKD '03*, pages 2–11, 2003.
15. M. Mahdaviani and T. Choudhury. Fast and scalable training of semi-supervised CRFs with application to activity recognition. In *NIPS*, December 2007.
16. D. Minnen, T. Starner, I. A. Essa, and C. L. I. Jr. Discovering characteristic actions from on-body sensor data. In *ISWC '06*, pages 11–18. IEEE, 2006.
17. N. Oliver and E. Horvitz. A comparison of HMMs and dynamic bayesian networks for recognizing office activities. In *UM 2005*, volume 3538, page 199, 2005.
18. J. Partridge, K.; Begole. Activity-based advertising: Techniques and challenges. In *Workshop on Pervasive Advertising 2009*, 2009.
19. D. Patterson, D. Fox, H. Kautz, and M. Philipose. Fine-grained activity recognition by aggregating abstract object usage. In *ISWC*, pages 44–51, 2005.
20. M. Philipose, K. Fishkin, M. Perkowitz, D. Patterson, D. Fox, and H. Kautz. Inferring Activities from Interactions with Objects. *IEEE Pervasive Computing*, 3(4):50–57, October 2004.
21. T. Stiefmeier, D. Roggen, G. Ogris, P. Lukowicz, and G. Tröster. Wearable activity tracking in car manufacturing. *IEEE Pervasive Computing*, 7:42–50, 2008.
22. T. Stiefmeier, D. Roggen, and G. Troster. Fusion of string-matched templates for continuous activity recognition. In *ISWC '07*, pages 41–44. IEEE, 2007.
23. D. Tacconi, O. Mayora, P. Lukowicz, B. Arnrich, G. Tröster, and C. Haring. On the feasibility of using activity recognition and context aware interaction to support early diagnosis of bipolar disorder. In *Ubiwell Workshop 2007*, pages 206–209, 2007.
24. M. H. Teicher. Actigraphy and motion analysis: New tools for psychiatry. *Harvard Review of Psychiatry*, 3(1):18–35, 1995.
25. E. Ukkonen. Constructing suffix trees on-line in linear time. In *IFIP World Computer Congress on Algorithms, Software, Architecture - Information Processing '92*, pages 484–492. North-Holland Publishing Co., 1992.
26. A. Vahdatpour, N. Amini, and M. Sarrafzadeh. Toward unsupervised activity discovery using multi-dimensional motif detection in time series. In *IJCAI '09*, pages 1261–1266. Morgan Kaufmann Publishers Inc., 2009.
27. K. Van Laerhoven, H. Gellersen, and Y. Malliaris. Long term activity monitoring with a wearable sensor node. In *BSN 2006*, pages 171–174. IEEE, 2006.
28. F. H. Wilhelm, M. C. Pfaltz, and P. Grossman. Continuous electronic data capture of physiology, behavior and experience in real life. *Interacting with Computers*, 18:171–186, March 2006.
29. A. Zinnen, U. Blanke, and B. Schiele. An analysis of sensor-oriented vs. model-based activity recognition. In *ISWC '09*, pages 93–100. IEEE, 2009.

Sensor Networks for Railway Monitoring: Detecting Trains from their Distributed Vibration Footprints

Eugen Berlin and Kristof Van Laerhoven
Department of Computer Science
Technische Universität Darmstadt
{berlin,laerhoven}@ess.tu-darmstadt.de

***Abstract*—We report in this paper on a wireless sensor network deployment at railway tracks to monitor and analyze the vibration patterns caused by trains passing by. We investigate in particular a system that relies on having a distributed network of sensor nodes that individually contain efficient feature extraction algorithms and classifiers that fit the restricted hardware resources, rather than using few complex and specialized sensors. A feasibility study is described on the raw data obtained from a real-world deployment on one of Europe's busiest railroad sections, which was annotated with the help of video footage and contains vibration patterns of 186 trains. These trains were classified in 6 types by various methods, the best performing at an accuracy of 97%. The trains' length in wagons was estimated with a mean-squared error of 3.98. Visual inspection of the data shows further opportunities in the estimation of train speed and detection of worn-out cargo wheels.**

Keywords*—*feature extraction, sensor data abstraction, event classification, railway monitoring, wireless sensor networks

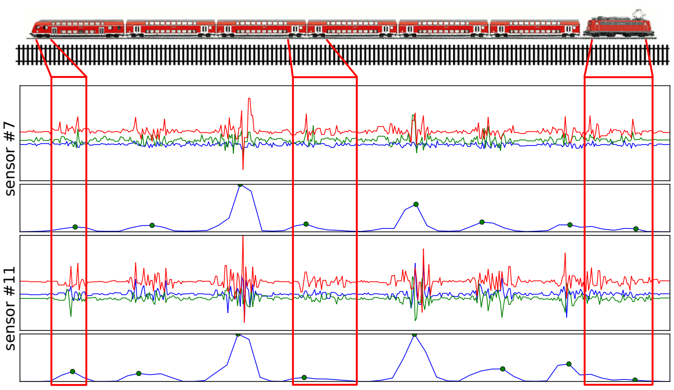

Fig. 1: Miniature sensor nodes attached to the railway tracks capture the vibrations caused by passing trains. From the raw 3D acceleration data of these events (upper plots), features can be extracted that are characteristic enough to be be used for on-sensor train classification (bottom plots). Using a network of such nodes makes the detection more robust and allows additional analysis, such as estimation of the train's speed by using time delays between sensors.

I. Introduction

Sensor networks have become a popular tool for various applications, due to being able to cover and monitor large areas and drastically reduce the intrusion into existing environments as well as disturbance of its inhabitants. The ability of wireless sensors to span a sensing and communication network with minimal resources by using small, robust, power-efficient and inexpensive hardware, highly benefits large-scale monitoring application. Such applications traditionally aim at periodic sampling of sensor values for long time periods, in order to obtain a detailed overview on physical phenomena in the environment. Many sensor network applications focus hereby on collective observation of slowly changing physical values, including temperature, humidity, gas concentrations in the air or particle concentration in the water. Hereby, the sensor nodes have to periodically wake up from low-power state in order to sample their sensors and disseminate the information through the network.

Other popular sensor network applications aim at detecting sporadic events, such as abrupt rising or falling of temperature and humidity, extremely high or hazardous concentrations of gas or pollutants in the air. The ability of the sensor nodes to detect such events directly at the source is of great advantage to the whole network, allowing to significantly reduce the amount of wireless communications within the network, thus preserving the limited power supplies.

Advances in software and hardware technology during the past decades made wireless sensor networks (WSN) more scalable, allowing the sensor nodes to be deployed for much longer time spans, in part due to the introduction of power-saving idle and sleep modes of the hardware components. With these advances, sensor networks have been increasingly deployed in scenarios with the aim to detect, monitor and report on more complex critical phenomena, such as seismic activity [1], disaster detection [2] or emergency scenarios [3], [4]. These applications require high-fidelity sensor data that preferably should be analyzed in or close to real-time, which conflicts with the fact that wireless sensor nodes within a network tend to be heavily constrained by their hardware capabilities and resources. Wireless communication is known to be one of the most energy-expensive operations, so that transmitting raw sensor data to a remote base station through the network will deplete the limited power supply in a short amount of time, resulting in sensor nodes and thus the network running the risk to become in-operational.

Our work focuses specifically on sensor data abstraction in an application where the sensors have been sampled at relatively high frequencies. Hereby, sampling rates range from hundreds of Hertz, for acceleration sensors and gyroscopes, up to thousands of Hertz for microphones. Using efficient and easy to compute features such as mean, variance, signal amplitude, and similar, abstraction of such sensor data is possible directly on the sensor nodes, even with their limited hardware resources. For applications where events require large amounts (i.e. multiple hundreds or thousands) of sensor readings to be adequately captured, computing such abstractions significantly reduces the amount of data in comparison to the original signal.

In this work, this approach of evaluating local features on a distributed set of wireless nodes is applied on a railway monitoring scenario. Figure 1 depicts one event from a data set recorded for the case study: The data set was obtained by deploying a network of sensor nodes that are equipped with sensitive inertial sensors at the railway tracks, capturing the vibrations caused by passing trains. We show that from the raw sensor data captured by the sensors in such a network, we are able to classify the type of train as well as the train's length. To achieve this in a realistic setting, we limit our approach to a set of sufficiently efficient features that can be implemented on the sensor node in an on-line fashion, thus allowing on-sensor event detection, train type classification and length estimation.

The remainder of this paper is structured as follows: Section II presents relevant related work. Section III is dedicated to our sensor hardware, the deployment location and the obtained data set. In section IV we propose and evaluate a set of features, both for train type classification and train length estimation. We discuss our results and highlight interesting findings as motivation for future work in section V. Finally, we list our conclusions made in this paper in section VI.

II. Related Work

Many research scenarios motivate the deployment of wireless sensor networks through the suitability of small sensors to densely monitor infrastructure, such as buildings, bridges, roads, rails. The huge diversity in sensors hardware, types of applications, deployment procedures, methodical and implementation approaches is astounding. One of the driving forces for monitoring structures and detecting relevant events is the goal for improving safety and organizing maintenance tasks. Various scenarios with alarm or control systems also motivate the detection and monitoring of critical events. This section will therefore present several application scenarios and frame our train monitoring study amid these related work.

A multitude of research, including [5], [6], [7] or [8] describe different application scenarios for wireless sensor networks, where detection and classification of rare events is of particular interest. The sensor networks, equipped with vision-based, acoustic, seismic, magnetic and infrared sensors, facilitate distributed observation of an area, aiming first and foremost at spotting and classifying ground vehicles or humans. While the scale of these deployments varies a lot, the need for energy-efficient sensors accounts for the features to be relatively simple to compute. For the car toll system application [8], the authors follow a similar approach to our work by choosing simple features (vehicle length, the average observed energy and peak-patterns in the signal) to detect and classify various ground vehicles, such as cars, pickup trucks, vans, buses and motorcycles.

Vibration sensors are often used for monitoring and ensuring infrastructure safety, such as in [9] or [10], where particular vibration frequencies in the raw data were considered as well as various complex features were utilized.

Railway safety and train detection plays an important role both from practical as well as research point of view, spawning several application scenarios utilizing different types of sensors: In [11], the authors present a system for short-term deployments that uses accelerometers to detect arriving trains in order to warn maintenance personnel working on tracks. Detecting and classifying train events by means of an acceleration sensor was suggested by [12]. To enhance railway safety, [13] deployed a vibration sensor on running trains, aiming at detecting rail deformations during motion. Electromagnetic sensor array can be used to detect and count wheels, as shown in [14]. Railroad operation monitoring through a wireless sensor network is presented in [15], aiming at more safety and improved efficiency of railway maintenance.

Reducing the energy consumption in wireless sensor network deployments, is often very crucial for the lifetime of the deployment. Data compression approaches, such as in [16], aim at reducing wireless communication payload. In the wearable and mobile sensor research domains, limited power supplies require more data processing directly at the device, avoiding energy-consuming transmission or storing to local memory. Reducing the communication payload by detecting activities directly on a mobile device, as presented in [17], will also extend the lifetime of the sensors.

Our work focuses on a wireless sensor network application scenario, where observed events can not be detected with simple threshold approaches and require on-line data processing. The aim is to classify train types and estimate train lengths by means of their vibration footprint directly on the sensor nodes. The events in this scenario tend to occur sporadically and last for only a short time period. The ability to detect, classify and monitor such events with sensor nodes deployed along railway tracks would allow to deploy such a network for various long-term railway applications.

III. Deployment

This section deals with the deployment of our sensor network that gathered the data set, the placement on the railway tracks as well as the underlying hardware choices.

A. Hardware

Since we aim at a long-term deployment of the sensor network at a given railway track of interest (Figure 2), the environment can be expected to be rough. Even though the sensor nodes used for the deployment are still research prototypes, their hardware and plastic enclosures especially manufactured for this purpose have been designed to withstand harsh outdoor environments. Further practical considerations that have been taken into account include protection of the sensor nodes against rain, snow damage, dust accumulation, and exposure to the sun, as well as attachment methods to the metal rails. As in many typical wireless sensor network deployments, we cannot assume power to be readily available next to the railroad and our deployment therefore has to rely on local batteries.

The sensor nodes for this paper's deployment and evaluation were designed to be small, robust and inexpensive enough to be left at the railway tracks surviving varying weather conditions. The sensor's main board features a Microchip PIC18F46J50 microcontroller as the main processing unit, an

Fig. 2: Our system's concept: A sensor network deployed along railway tracks captures vibrations caused by passing trains. Immediately computing efficient features from streaming sensor data allows train type classification and counting wagons on the sensor nodes. In future deployments, these can be used to estimate train speed and detect worn-down cargo wheels. As such, this system is envisioned as a flexible instrument to assist in railway monitoring and maintenance tasks.

Fig. 3: One of the sensor nodes attached to the rails during the experimental deployment. The sensor node is wrapped up in a plastic bag to protect the sensor from dust and humidity.

Analog Devices 3D MEMS accelerometer for capturing the vibrations and a micro-SD card connected via SPI-bus for local storage. Furthermore, the main board contains interfaces for reprogramming, wireless extension, and additional sensors such as light or temperature. A mini-USB port is used for connecting the sensor node to a computer, which allows configuring the sensor, initiating the logging process and accessing the recorded sensor data afterwards. A plastic case (overall dimensions: 37x33x15mm) manufactured through 3D printing holds the components and provides basic protection.

The PIC18F46J50 microcontroller features 65528 bytes of program memory and 3776 bytes or random access memory and is equipped with a real-time clock which is driven by a precise $32.768kHz$ Abracon crystal with a frequency tolerance of $\pm20ppm$. The real-time clock drift amounts to few milliseconds per day, allowing exact time-stamping of the occurring events. Time synchronization in the network therefore needs to happen infrequently, allowing synchronized monitoring of passing trains along the mounted sensors and speed estimation based upon event delay. This particular microcontroller also supports USB communication and is able to swiftly change into low-power modes depending on the tasks at hand.

The 3D ADXL345 accelerometer is configured to sample its data at $100Hz$ and transfer the readings in bursts of 32 samples to the microcontroller. The time span between the bursts is long enough allowing the microcontroller to process the previous burst of sensor data and to switch into a low-power idle mode to preserve limited battery power. All sensor nodes were configured to sense vibrations in the $\pm4g$ range and at 10-bit resolution. A tiny lithium polymer rechargeable battery with a capacity of $180mAh$ was chosen as the power source, allowing a single sensor node to run for approximately two weeks while logging all data to the micro-SD card.

The choice for a sensor node design with a low-power microcontroller makes the entire module small and cheap to produce, but also results in another significant challenge: The limited amount of program and operating memory as well as the lack of a floating-point unit poses a harsh limit on the used algorithms and their implementation. Thus, the proposed feature extraction routines have to work under extreme memory constraints and should avoid the use of larger functions (as they are for example used in Fourier analysis).

Choosing inexpensive off-the-shelf MEMS accelerometer sensors to detect and characterize trains by their vibration footprint results first and foremost in nodes that can easily be built in large quantities, allowing deployment of large-scale sensor networks. A second benefit lies in the accelerometer packages occupying very little space and generating sensor data at a bandwidth and resolution that can be processed directly at the sensor node with available processing capabilities. On the other hand, one can expect the sensor data quality to be less accurate than that of specifically-designed and more expensive vibration sensors, making the extraction of distinctive and characteristic features more important.

The sensors' raw acceleration logging routine requires roughly a fifth of program and random access memory. The proposed feature extraction algorithms would therefore easily fit on these sensor nodes, allowing on-line computation and forwarding of features or classification results through the sensor network instead of the much larger amount of raw sensor data. In order to evaluate these features, the following subsections present an experimental deployment for obtaining real-world vibration patterns.

B. *Location*

Before being permitted to deploy the sensor nodes on highly busy railway tracks, railway company officials were involved in the planning of this paper's experiments. The location for the deployment of the sensor nodes was suggested by the railway company officials themselves, specifically due to the variety of train types and their maximum possible speeds. The particular spot featured four tracks running completely straight for multiple kilometers, thus allowing train speeds of up to 250 kilometers per hour.

Two out of these four railway tracks are general purpose high-speed approved tracks for different train types, including regional and high-speed passenger trains from two European countries, and cargo trains. The other two tracks, due to their technical characteristics, were used by low-speed passenger trains connecting nearby cities.

During the deployment, each sensor was additionally wrapped up in a plastic vacuum bag in order to protect it from dust, humidity and rain, and attached to the rails using double-sided adhesive tape, as shown in Figure 3. Attaching the sensors to the rails was carried out under the supervision

Fig. 4: Snapshots from the video recordings showing in this case a single locomotive passing by on the 2nd high-speed track. These videos were used as ground truth for the evaluation of type of train and train composition (wagon count).

ID	Class	Description	Count
1	Regio	passenger trains connecting cities in a region	63
2	CityRail	trains service city center, suburbs and nearby cities	15
3	Cargo	various cargo trains	39
4	Loc	single locomotives being transferred	10
5	Thalys	French inter-city high-speed passenger train	5
6	ICE	German InterCity Express high-speed passenger train	9

TABLE I: From all the train events in raw sensor data, 141 could be annotated and used for the evaluation. Here, different train types and their count in the data set are shown.

of a railroad maintenance crew, whereby the tracks had to be officially closed for train traffic for a short period of time. For the deployment, six sensor nodes have been deployed along both high-speed and low-speed tracks, with a distance of 10 meters between them. At configuration time, the nodes' realtime clock units were synchronized with a camera setup registering both audio and video from passing trains, so that the data could easily be synchronized afterwards.

C. The Data Set: A Visual Inspection

This subsection describes the variety of trains present in the data set, consisting of train events captured by the sensors, and the video footage recorded for annotation purposes.

In total, the sensors have captured 186 train events, of which 141 could be annotated based on video footage recorded during the deployment. Figure 4 illustrates the video data with a series of frames from video footage capturing a single locomotive passing by. Table I provides an overview on the six different train types: four different passenger train classes – two types of high-speed trains, regional passenger trains and city rail trains – along with a cargo and locomotive classes.

Thalys, a French high-speed passenger train, typically consists of head and tail locomotives and 8 passenger wagons which are connected to a single continuous unit, resulting in 10 wagons in total. The last-generation InterCity Express (ICE) is a German high-speed passenger train which in our experiment typically contained 8 railmotor wagons (i.e., no locomotive). Figure 5 depicts models and truck constellations for these two train types, while Figure 6 shows examples of the captured vibration footprints and window variance used for feature extraction, revealing tiny differences: Thalys' windowed variance peaks in the middle of the train correspond to single trucks (often referred to as "Jacobs bogies") between the wagons; ICE peaks correspond to two adjacent trucks, thus resulting in slightly wider variance peaks.

The Regio class contains the regional passenger trains that connect nearby cities within a region, but do not stop at stations

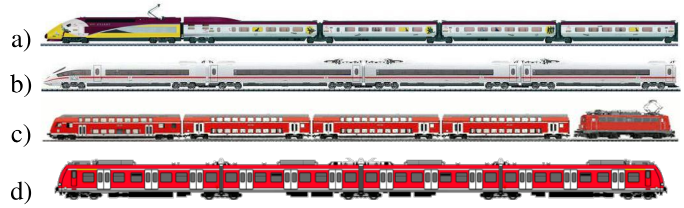

Fig. 5: The four passenger train types in the data set, from above: a) Thalys, b) ICE, c) Regio, d) CityRail. Note that the trucks' locations differ among the train types, with two trucks for a wagon (e.g. ICE) or one between them (e.g. CityRail), resulting in characteristic vibration footprints.

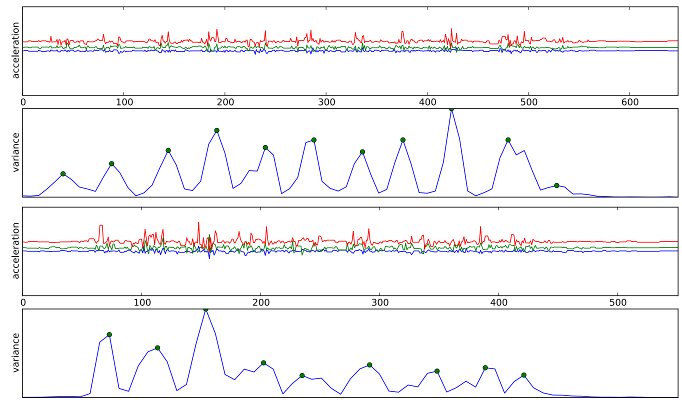

Fig. 6: Example plots showing Thalys (above) and ICE (below) high-speed passenger trains raw sensor data and windowed variance with extracted peaks. These plots show how different constellation of trucks produce distinctive vibration footprints (cf. train models in Figure 5).

in between. Regional trains consist of a locomotive pulling or pushing a number of bi-level wagons (as shown in Figure 2 and the corresponding model in Figure 5c). In our experiment, these trains' lengths were 3, 5, 6 and 7 wagons in total.

The Cargo class has proven to be rather versatile, with one characteristic feature that all cargo trains have in common: at least one locomotive is pulling a highly varying number of wagons. Both the locomotives as well as the wagons themselves can be of different types (e.g., tanks, containers, car- or freight wagons), as well as different lengths and truck constellations. In our experiment, cargo trains had mostly one, sometimes two, locomotives with a total number of wagons ranging from 13 up to 43.

The locomotives class was added due to single locomotives being transferred to another station. In the experiment, 10 such events have been captured, whereby both single as well as two connected locomotives have been observed.

The city railway trains (CityRail) connect larger cities with its suburbs and other smaller towns nearby. The train typically consists of two electrical units with 4 wagons each (Figure 5d). The CityRail trains in our experiment were running exclusively on the separate low-speed tracks.

With this data set, we can now perform an extensive evaluation, in which particular focus is given to finding a set of efficient-to-calculate features that can be implemented directly on the sensor nodes for on-line train type classification. A second objective that has been identified as valuable information to automatically detect by the sensor network is the estimation of train length. The following section will present the proposed features, the classification performance and train length estimation results.

IV. Evaluation

This section is divided into three parts. First we present features that can be efficiently extracted on the sensor nodes from raw 3D accelerometer data. The second part deals with the train type classification performance of the proposed features, specifically aiming at finding the optimal parameters as well as the best performing combinations of features. The third part evaluates how well the train length can be estimated.

A. Features

Aiming at a railway monitoring application which relies on detecting and classifying passing trains directly at the sensor nodes, a set of characteristic features that are able to describe and distinguish various train events is necessary. In this scenario, train events occur sparsely over the course of time, so that most sensor data acquired by the sensor nodes can be discarded as not relevant, when no trains are passing by. These flat signal sections between train events can be accurately segmented by utilizing the standard deviation in a sliding window over signal.

When a train passes by, the sensor node will be able to detect this event by the changing acceleration values, and due to the real-time clock also the exact time when the event begins and ends. From this, the duration of an event can already be derived as the first feature. When considering the whole event, the total amount of vibration caused by the train can easily be computed in an on-line fashion and thus used as a feature.

Visual inspection of the vibration footprints has led to the assumption that single trucks (containing one or multiple axles) can be found in the signal. To achieve this, a small window buffer is used to compute windowed variance from the raw sensor data (cf. Figure 1) or 6), which can then be used to derive further features. The number of peaks in the windowed variance plot is such a feature, which captures trucks and can be used to count the number of wagons. Since the peaks' detection highly depends on the width of this sliding window, it is being considered a parameter in the following evaluation.

After computing the peaks, more information can be extracted: the amount of vibrations of the trucks through maximum and average of the amplitudes, truck distances through the average distance between peaks, variety of wagon lengths or trucks constellations via variance of peak distances. Additionally, the overall area under the variance curve, as well as the average area per peak will be considered. For the offline evaluation we compute this feature using the Python `scipy.integrate` library. Table II summarizes the proposed features that were used in this study.

With this set of features at hand, our interest lies in finding the appropriate parametrization for the sliding window and from that the best performing feature combination.

ID	Feature	Description
0	duration	event duration (vibration exceeding a threshold)
1	variance	total amount of vibration caused by the train
2	peaks	number of peaks extracted from windowed variance
3	max. amplitude	maximum peak value
4	avg. peaks	average distance between peaks
5	avg. amplitude	average peak amplitude
6	area	total area under curve
7	avg. area	average peak area under curve
8	var. peaks	variance of peak distances

TABLE II: Overview of all features considered for train type classification. During the 5-fold cross validation on the data set using an SVM classifier, all possible feature combinations (with a minimum number of three) have been tested, whereby the features were also computed with a varying window size.

B. Train Type Classification

This section discusses the feature extraction parametrization, mainly depending on the window size over which the windowed variance of the raw signal is being computed. Since the sensor nodes were set to a sampling frequency of $100Hz$, the following range of window sizes was found to be of interest for evaluation: 12, 14, 16, 18, 20 data points.

For the train type prediction, the versatile support vector machine (SVM) classifier has been chosen. The feature space was normalized before being used for the evaluation.

The performance evaluation was carried out through a stratified 5-fold cross-validation, whereby the size proportion of the six classes was preserved. The classifier was trained on 4 parts, while one part was left out for testing. Hereby, all possible feature combinations have been tested (with three as a minimum features set size), resulting in 466 combinations. Multiplying this with the range of window sizes, we end up with 2330 cross-validation runs. Here, we present the cheapest (regarding the number of features required) best performing feature combinations.

The classification results obtained during the 5-fold cross validation were accumulated, and confusion matrices were computed averaged over the number of folds. Figure 7 shows the four most illustrative confusion matrices, along with their parameters, the window size and the set of features.

For the first evaluations with the first three features only (duration, total variance and number of peaks), the SVM classifier was able to reach an overall accuracy of 90.78% for the window size of 18 data points. Adding the maximum amplitude to the feature set led to an increase of total accuracy 93.62% for the window size of 16. Adding more features to the set or interchanging them would improve the accuracy in very little steps, such that many combinations would reach a classification performance with 136 out of 141 train types correctly identified (96.45% accuracy). Figures 7a, 7b and 7c show three confusion matrices with corresponding feature sets that were able to achieve this high classification performance.

The feature set consisting of feature IDs 1, 4, 5, 6, 7 and 8 (computed from the windowed variance with a window size of 16 data points) has reached the maximum possible accuracy of 97.16%, with 137 of 141 train events being correctly identified. Figure 7d depicts the confusion matrix for this best performing feature set and window size.

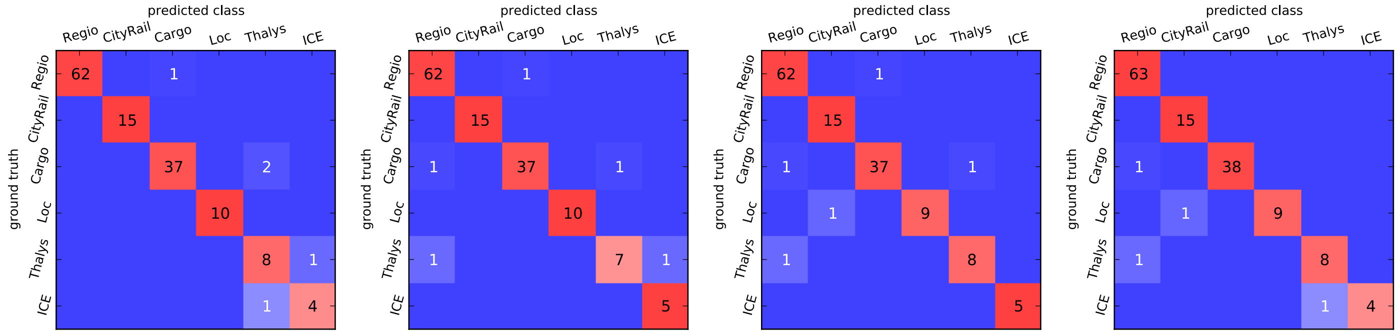

(a) ww=14; ftrs=[2, 3, 4, 5] (b) ww=16; ftrs=[2, 4, 5, 6, 7] (c) ww=16; ftrs=[1, 2, 4, 6, 7, 8] (d) ww=16; ftrs=[1, 4, 5, 6, 7, 8]

Fig. 7: Exemplary selection of confusion matrices obtained during the 5-fold cross-validation with the SVM classifier. The first three matrices show an accuracy of 96.45%, for different sets of features computed on different window sizes. The right-most confusion matrix shows the feature set performing best, reaching an accuracy of 97.16%.

These classification performance on our data set suggests that training an SVM classifier offline and implementing it on the sensor nodes would allow to detect train events, compute features and predict train types directly at the signal source with high accuracy.

Predicting the types of passing trains with reaching accuracies up to 96% and 97%, would be sufficiently promising for several applications. Following our scenario, deploying a sensor network with such a SVM classifier implemented on each sensor node, it is still possible to improve on the classification performance. This can be achieved by utilizing the sensor network's communication capabilities and let neighboring sensor nodes decide upon the train type by a voting mechanism amongst classifiers.

After evaluating the train type classification performance with features, the following section will give insight on how well the train length estimation worked.

C. Train Length Estimation

To estimate the train length, we primarily rely on counting the number of wagons in the trains. This can be achieved by using the already introduced feature "number of peaks" as a basis. Additionally, using the train type obtained from the previous estimation step is used as a prior. The wagon count can furthermore be improved by a comparison and voting procedure amongst neighboring sensor nodes on the same railroad track.

Besides using the raw signal and computed features, it is useful to include inherent model knowledge about the train type constellations: The ICE and Thalys high-speed passenger trains as well as the CityRail trains consist of specific wagons only (locomotives are built-in or the wagons are motorized themselves). Regio trains consist of varying amount of wagons with a separate locomotive. While with these trains the axles constellations are fixed due to defined sets of wagons, the cargo train class poses a much higher variety: wagons with single, double and triple axles per truck, wagons of different lengths, and varying load are possible.

Using the annotations from the video footage as ground truth (number of wagons) and the number of peaks extracted from the windowed variance, we use their difference for performance analysis. Since the peaks correspond to the trucks, the number of peaks usually is by 1 more than the amount of wagons in the train. This deviation of 1 can be visually recognized in the box plots shown in Figure 8. In addition, we compute the mean-squared error for the whole data set as well as for each individual class (see Table III), whereby the deviation has been accordingly taken into account.

Window	Overall	Regio	CityRail	Cargo	Loc	Thalys	ICE
12	28.14	11.78	1.87	78.95	7.30	4.00	2.00
14	10.33	4.65	1.00	27.95	2.60	3.56	0.20
16	**4.02**	2.84	0.27	**8.74**	2.20	2.33	**0.00**
18	**3.98**	2.89	0.33	8.87	1.20	**1.44**	0.60
20	5.62	**2.43**	**0.20**	14.28	**1.00**	6.56	2.00

TABLE III: Mean-squared error of the estimated train lengths for the whole data set and per class. The window size has a huge impact on the quality of the estimation. Overall, window size of 18 performs slightly better than a window size of 16.

For a more concrete example, let us consider a regional passenger train with 7 wagons (including the locomotive). For this train, the peak detection algorithm extracts 8 distinctive peaks (cf. Figure 1). Hereby we observe that the first peak belongs to the first trucks (double axles) of the first wagon, the following 5 peaks belong to adjacent wagon trucks (two times double axles for passenger wagons), and the last two peaks represent the last wagon and the locomotive (which has triple axle trucks that can not be distinguished in the signal with the fixed window size). Due to Regio trains' variety in length (3, 5, 6 and 7 wagons including one or even two locomotives) and their varying speed when passing by the sensors, the relation of wagons to the number of peaks tends to highly deviate as well as show lots of outliers (Figure 8a).

The window size to compute the windowed variance from raw sensor data has a significant impact on the peak detection. On the other hand, leaving the window size fixed at the best performing size of 16 data points (for classification and length estimation) would deteriorate the system's performance, as a fixed window size for a fixed sampling rate of the sensor node leads to the issue of not being speed independent.

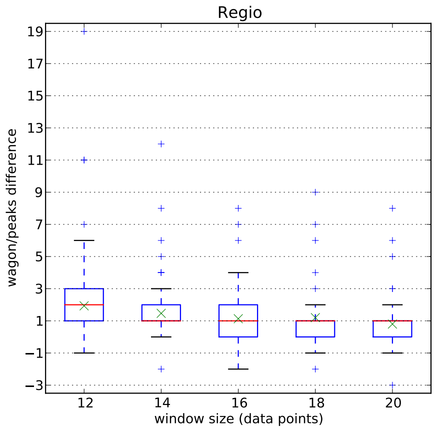

(a) Regio trains with varying length and speeds result in high deviations and outliers.

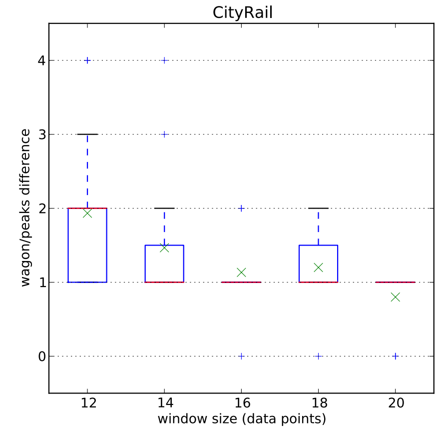

(b) CityRail trains' trucks can be quite accurately detected with window size of 16 & 20.

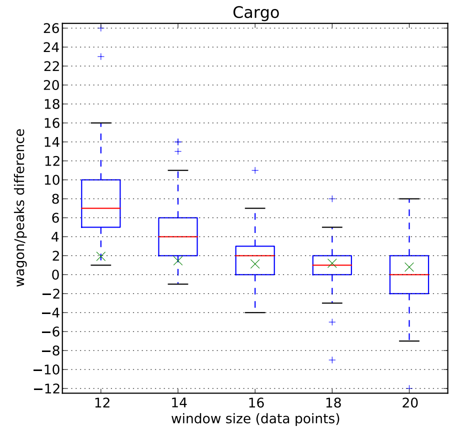

(c) Cargo trains' variety in number and types of wagons results in high peak count variance.

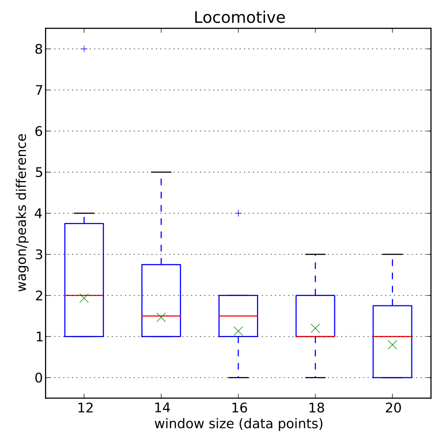

(d) High variance for single locomotives due to locomotive type, axles configuration, speed.

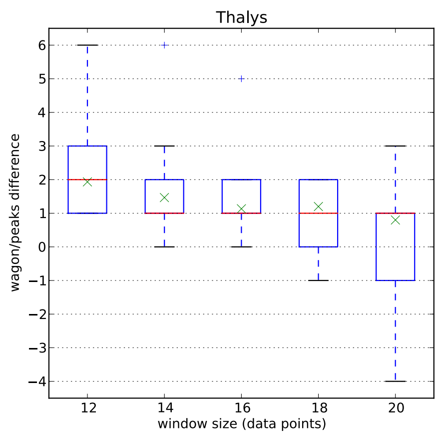

(e) Thalys trains trucks between wagons can be best detected with a window size of 18.

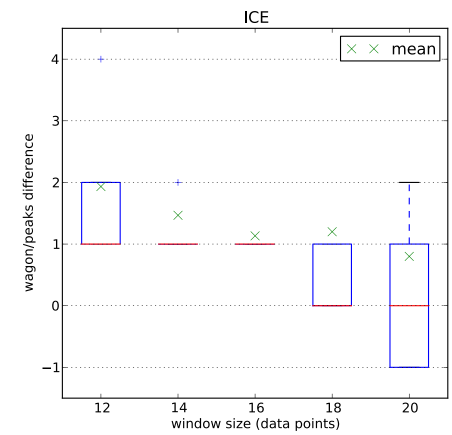

(f) ICE trains double-axles trucks per wagon are optimally detected with a window of 16.

Fig. 8: Differences between the real number of wagons and number of peaks computed from the windowed variance for each of the six train classes. The number of peaks and therefore the accuracy of the wagons count depends on the size of the sliding window and the train speed. From these results, window sizes of 16 and 18 data points are performing best for counting wagons. This is verified by the overall minimum squared-mean error of approximately 4.0 shown in Table III.

V. Discussion and Outlook

This section discusses the evaluation results and will point out particularly interesting findings for the underlying scenario.

First, good train type classification results (up to 97% accuracy) could be achieved on our data set with proposed features. During the evaluation, suitable window sizes (16 data points) both for type prediction and train length estimation could be found. Better performance in this regard can be achieved through implementing distributed voting among neighboring sensor nodes inside the sensor network. This would allow the sensor nodes to compare decisions and remove outliers.

Estimating train length with a fixed window size bears the problem of not being speed independent. In case of a very slowly moving train, which is very likely to happen and has also been observed in our data set, the fixed window will lead to detecting separate axles instead of trucks, resulting in a completely misleading peak count. One possible approach to tackle this issue would be the introduction of a variable window size. This would, on the other hand, lead to a more complex feature extraction routine and result in more computation on the sensor node and therefore in a higher power consumption.

Besides the problems addressed in this work, the recorded data set allows to extract much more information useful for the railway monitoring scenario. Estimating train speeds belongs to this category of very useful details and can be achieved with multiple sensors placed at a predefined distance which, in our experiment, was 10 meters. Detecting a passing train on two sensors and then computing the time delay between event arrival seems to be an easy and reasonable approach. For this, time synchronization inside the sensor network is important, but can be nowadays considered as a solved problem. An example for the feasibility is shown in Figure 9: vibrations caused by a passing regional passenger train are captured by two sensors in 10 meters distance from each other. By aligning these raw sensor data in time, the delay became visible.

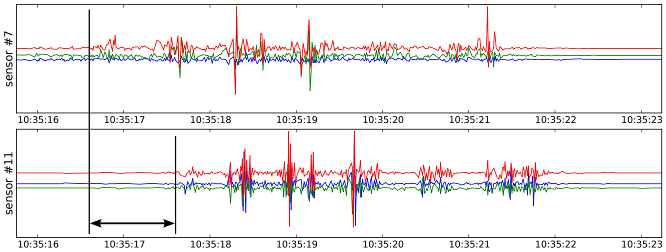

Fig. 9: Two sensor nodes showing the vibrations footprint caused by a passing regional passenger train. Knowing the distance between the sensor nodes ($10m$ in our deployment) and the time delay of the event between two sensors (markers in the plots) will allow estimating the train speed.

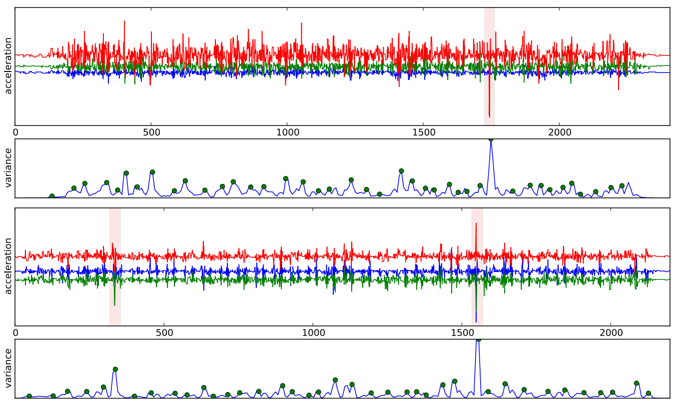

Fig. 10: The sensor nodes have picked up extreme impacts (large peaks in raw data and variance plots) from a passing cargo train, caused by a defect wheel that is not completely balanced due to wear during braking. Since these wheels could cause damage to rails, rail bed, and the wagons, detecting such events automatically would be of significant interest as well.

Another very promising application for such a sensor network would be the detection of worn-out or defect wheels. Figure 10 shows data from two sensors which have picked up extreme accelerations caused by a passing cargo train. These extreme amplitude peaks in the raw sensor data are most likely caused by worn wheels (having lost their roundness due to abrasion caused by blocking when the train breaks). These worn-down wheels could cause damage to rails or rail bed, as well as the wagons themselves, thus making the detection of such events particularly interesting.

VI. Conclusion

We evaluated the suitability of a sensor network consisting of tiny, inexpensive sensor nodes for a train monitoring application. It relies on sensor data from 3D MEMS accelerometers that are able to capture vibrations caused by running trains. To enable in-network event detection and train type classification, we have proposed a set of features that can be computed efficiently and in an on-line fashion directly on the sensor nodes. After deploying the sensors at railway tracks, recording a real-wold data set and video footage for annotation purposes, we conducted an evaluation of the proposed features. Using an SVM classifier, the feature set (1, 4, 5, 6, 7, 8) and the window size of 16 data points were found to produce optimal results for this data set. The SVM classification performance reached 97% accuracy. The length estimation performance accounted to 3.98 mean-squared error for the whole data set.

Acknowledgment

We would like to thank Hans Gabler, Hubert Schmühl and Harald Behr from the Deutsche Bahn Netz AG who made this real-world deployment possible. This work was sponsored by the Project "Long-Term Activity Recognition with Wearable Sensors" (LA 2758/1-1) from the German Research Foundation (DFG).

References

[1] G. Werner-Allen, K. Lorincz, M. Ruiz, O. Marcillo, J. Johnson, J. Lees, and M. Welsh, "Deploying a wireless sensor network on an active volcano," *Internet Computing, IEEE*, vol. 10, no. 2, pp. 18–25, 2006.

[2] M. Ramesh, S. Kumar, P. Rangan, and A. Vidyapeetham, "Wireless sensor network for landslide detection," in *International Conference on Wireless Networks (ICWN 2009)*, 2009, pp. 89–95.

[3] T. Gao *et al.*, "The advanced health and disaster aid network: A light-weight wireless medical system for triage," *IEEE Trans. on Biomedical Circuits and Systems*, vol. 1, no. 3, pp. 203–216, September 2007.

[4] T. Gao, C. Pesto *et al.*, "Wireless medical sensor networks in emergency response: Implementation and pilot results," in *Conference on Technologies for Homeland Security (HST 2008)*, 2008, pp. 187–192.

[5] G. P. Mazarakis and J. N. Avaritsiotis, "Vehicle classification in sensor networks using time-domain signal processing and neural networks," *Microprocess. Microsyst.*, vol. 31, pp. 381–392, 2007.

[6] S. Gupte, O. Masoud, R. Martin, and N. Papanikolopoulos, "Detection and classification of vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 3, no. 1, pp. 37–47, March 2002.

[7] P. Dutta *et al.*, "Design of a wireless sensor network platform for detecting rare, random, and ephemeral events," in *Int. Symp. on Information Processing in Sensor Networks (IPSN)*, 2005, pp. 497–502.

[8] S. Keawkamnerd, J. Chinrungrueng, and C. Jaruchart, "Vehicle classification with low computation magnetic sensor," in *Int. Conf. on ITS Telecommunications (ITST 2008)*, 2008, pp. 164–169.

[9] H. Jingqiu *et al.*, "On vibration signal analysis in bridge health monitoring system by using independent component analysis," in *Society of Instrument and Control Engineers (SICE)*, 2010, pp. 2122–2125.

[10] K. Sukun *et al.*, "Health monitoring of civil infrastructures using wireless sensor networks," in *Int. Conf. on Information Processing in Sensor Networks (IPSN 2007)*. ACM, 2007, pp. 254–263.

[11] L. Angrisani, D. Grillo, R. Moriello, and G. Filo, "Automatic detection of train arrival through an accelerometer," in *Instrumentation and Measurement Technology Conference (I2MTC)*, 2010, pp. 898–902.

[12] E. Berlin and K. Van Laerhoven, "Trainspotting: Combining fast features to enable detection on resource-constrained sensing devices," in *Int. Conf. on Networked Sensing Systems (INSS)*, 2012, pp. 1–8.

[13] C. Wang *et al.*, "On-line vibration source detection of running trains based on acceleration measurement," in *Int. Conf. on Intelligent Robots and Systems (RSJ 2006)*, 2006, pp. 4411–4416.

[14] P. Donato, J. Urena, M. Mazo, J. Garcia, and F. Alvarez, "Electromagnetic sensor array for train wheel detection," in *Sensor Array and Multichannel Signal Processing Workshop*, 2004, pp. 206–210.

[15] E. Aboelela, W. Edberg, C. Papakonstantinou, and V. Vokkarane, "Wireless sensor network based model for secure railway operations," in *IPCCC 2006*, 2006, pp. 624–628.

[16] A. Reinhardt, D. Christin, M. Hollick, and R. Steinmetz, "On the energy efficiency of lossless data compression in wireless sensor networks," in *Conf. on Local Computer Networks (LCN 2009)*, 2009, pp. 873–880.

[17] F.-T. Sun, C. Kuo, and M. Griss, "Pear: Power efficiency through activity recognition (for ecg-based sensing)," in *Int. Conf. on Perv. Comp. Techn. for Healthcare (PervasiveHealth)*, 2011, pp. 115–122.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JBHI.2014.2326604, IEEE Journal of Biomedical and Health Informatics

MyHealthAssistant: An Event-driven Middleware for Multiple Medical Applications on a Smartphone-mediated Body Sensor Network

Christian Seeger, Kristof Van Laerhoven, and Alejandro Buchmann

J-BHI Special Issue on "Service Science for e-Health"

***Abstract*—An ever-growing range of wireless sensors for medical monitoring has shown that there is significant interest in monitoring patients in their everyday surroundings. It however remains a challenge to merge information from several wireless sensors and applications are commonly built from scratch. This paper presents a middleware targeted for medical applications on smartphone-like platforms that relies on an event-based design to enable flexible coupling with changing sets of wireless sensor units, while posing only a minor overhead on the resources and battery capacity of the interconnected devices. We illustrate the requirements for such middleware with three different healthcare applications that were deployed with our middleware solution, and characterize the performance with energy consumption, overhead caused for the smartphone, and processing time under real-world circumstances. Results show that with sensing-intensive applications our solution only minimally impacts the phone's resources, with an added CPU utilization of 3% and a memory usage under 7 MB. Furthermore, for a minimum message delivery ratio of 99.9%, up to 12 sensor readings per second are guaranteed to be handled, regardless of the number of applications using our middleware.**

***Index Terms*—Body sensor networks, Cybercare, Medical services, Middleware, Performance analysis, Wireless sensor networks**

I. Introduction

An ageing population and low birth rates are leading to a demographic change in the West that significantly challenges its health care systems [1], [2]. In addition to this, the World Health Organization predicts that chronic diseases will become the most expensive problem faced by current health care systems and sees the integration of prevention into health care as the main solution for this problem [3]. A paradigm shift towards integrated and preventive health care, as well as equipping patients with information, motivation, and skills in prevention and self-management, are described as essential elements for solving these problems. Systems that collect information from a network of on-body and ambient sensors are a promising tool for such solutions: As body sensor network systems are capable of continuously monitoring a person's physiological and physical state [4]–[6], they can provide patients with the required information and motivation. Combined with the additional information of the user's surroundings via ambient sensors, full-fledged Body and Ambient Sensor Network (BASN) health monitoring solutions can be built to face these upcoming challenges in health care systems.

We focus in this paper especially on the need for supporting multiple sensor constellations and the integration in the user's environment as significant features for many medical BASN applications. For a patient with a cardiac disorder for instance, monitoring of blood pressure, ECG, and physical activities would be preferred. As monitoring progresses, data from additional respiration and blood oxygen saturation sensors might become relevant to observe a developing sleep apnea. Additionally, information about the user's environment is important for a correct interpretation of many vital parameters: Coffee consumption before taking a blood pressure reading can for instance influence the results [7]; Readings from ambient sensors help determining such contextual information.

MyHealthAssistant is proposed here as a middleware designed for managing body and ambient sensor networks for user-centric health monitoring. From a systems perspective, it 1) is able to cope with interchanging sets of sensor units, 2) is fast to deploy on a user's personal phone, and 3) supports ambulant, day-long monitoring. The event-based design contains dedicated modules that translate sensor data to events to support adapting the system's functionality, extending the sensor set, and to cope seamlessly with changing sensors. The middleware is furthermore designed to run on a smartphone, making use of its connectedness, processing power, and user acceptance. Additional services for detecting sensing artifacts and a worsening system status aim to support application developers.

This paper is structured as follows: after discussing differences between our work and related work, we describe the design choices for our middleware supporting body and ambient sensor network applications in Section III. Section IV presents three case study applications that were built on top of the middleware, all making use of different sets of sensing units. Section V discusses the middleware's performance in detail with focus on energy consumption and efficiency of information routing through the system.

II. Related Work

Many body sensor network (BSN)-based projects in health care focus on monitoring of a particular disease or set of phys-

Supported by the German Research Foundation (DFG) within research training group 1362: Cooperative, Adaptive, and Responsive Monitoring in Mixed Mode Environment. Manuscript received December 21, 2013; revised May 21, 2014.

C. Seeger, K. Van Laerhoven and A. Buchmann are with the Department of Computer Science, Technische Universität Darmstadt, Germany e-mail: {cseeger,buchmann}@dvs.tu-darmstadt.de, kristof@ess.tu-darmstadt.de.

iological signals [4]–[6]. They benefit from the independence from stationary in-hospital observations, allowing patients to freely move and live their daily life while being monitored over longer times and under more realistic conditions. In the Partnership for the Heart project [8], 710 patients with cardiac disorders were equipped with a stationary scale, ECG, SpO2, blood pressure sensors, a hip-worn activity sensor as well as a PDA for transmitting the daily measurements to a remote health care provider. Less hospital stays, an increased quality of life, and a faster reaction to health changes are promising results of this study. The system contained a fixed set of sensors and a relatively sparse monitoring technique.

In contrast to such a fixed setup, middleware for sensor networks such as MiLAN [9] allow a more flexible combination of sensors and applications. MiLAN allows applications to define their QoS requirements over time and how to meet these requirements using different sensor combinations. Based on different priority levels and information about the available sensors and their status, MiLAN continuously adapts the network configuration to meet the application's needs while maximizing the system's lifetime. In order to provide a proper management of the sensor network, MiLAN requires a tight integration with the sensors and protocols.

The Self-Managed Cell (SMC) [10] is a middleware which consists of a policy-based architecture that supports autonomic management and self-configuration for BSNs. Policies define how the system should adapt in response to specific events and an event bus provides content-based subscriptions.

Waluyo et al. [11] propose a middleware for medical BSNs that supports multiple sensors and applications, plug and play features and resource management. Similar to the use case presented in our introduction, they consider vital parameter monitoring and behavior monitoring as two applications running on the same sensor network. In that project however, parts of the middleware and the applications reside on a PC.

In [12], an Android based body area network for telemedical systems is presented. The authors discuss two sensor network setups that consists of either wired or wireless connections from the sensors to a gateway node, which transmits sensor readings to an Android phone for further processing and data forwarding. Several challenges such as data acquisition, visualization, data storage, and safety are discussed.

The work presented so far has a particular focus on optimizing the sensors and the sensor network itself. Further related projects are DexterNet [13], SIXTH [14], Lifeware [15], MobiSense [16], VITRUVIUS [17], and Kamal et al. [18].

In contrast, our middleware approach focuses on providing sensor information to multiple mobile and personal health applications running on a mobile device such as a smart phone. It operates with off-the-shelf sensors from different manufacturers and it does not need an adaptation of the sensors. Therefore, the aim of our middleware is to mediate between sensors from different manufacturers and multiple health-related applications running on a single mobile device.

Jones et al. [19] present the MobiHealth project which consists of a generic BSN for health care as well as a generic mobile health service platform. This system also provides sensor data to multiple applications running on a mobile device with focusing mainly on the network infrastructure among a patient's BSN and health care provider.

Morón et al. propose a smart phone-based telecare system using a body area network [20]. The system consists of commercial off-the-shelf Bluetooth sensors that measure vital parameter and a conventional smart phone which allows using the system without any hardware or software modifications. A target application of that work is the efficient monitoring and management of chronic diseases. The particular focus is on analyzing the impact of different implementations (i.e., Java vs. Python) and different message forwarding techniques between the smart phone and the back-end system.

For health care applications, event-based systems have been presented mainly in areas with high amounts of data. Examples are intensive care solutions, real-time sleep analysis, and solutions for establishing large health care networks. All solutions benefit from the efficient data processing provided by event-based systems. The following describes two examples.

Intensive care units are equipped with numerous devices for monitoring a patient's health parameters. Many of them are stand-alone devices with individual alarming systems triggering their own alarm event. Guerra et al. [21] propose an event-based system that combines the events from individual sensors. It integrates in one place historical data, events, rules, and data mining models and it is highly customizable. In addition, the system performs data mining for identifying possible future risks (e.g. cardiac arrests).

Besides patient monitoring in intensive care units, Singh et al. [22] propose an event-based middleware for patient monitoring in their home environments. The home care system sends monitoring reports and state changes to health care providers and triggers alarms in case of emergencies. By characterizing such a scenario as highly data-driven, the authors chose an event-based system. The particular focus of this work lays on enabling data security by adding dissemination control.

III. ARCHITECTURE

In [23], the requirements for applications running on a phone-based medical body sensor network were analyzed and an event-driven, layered middleware architecture was proposed. Event-driven systems fit the nature of both body and ambient sensor networks because: a) sensor constellations and running applications change over time, b) most body and ambient sensors send their readings in an event-driven manner (e.g., as alarms when thresholds are exceeded), and c) modular sensor units are agnostic on which applications use their readings. Event hierarchies allow establishing a common, extensible data abstraction on which various applications can be built and by introducing event transformations (e.g., ACTrESS [24]) or ontologies (e.g., CONNECT [25], [26]) to the system, a comprehensive interoperability and integration in other event-based systems would be provided.

A. Event-driven Architecture

Fig. 1 depicts the architecture of MyHealthAssistant, annotated with implementation details as an Android `Remote Service`. On the left, wireless sensor units can be connected

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JBHI.2014.2326604, IEEE Journal of Biomedical and Health Informatics

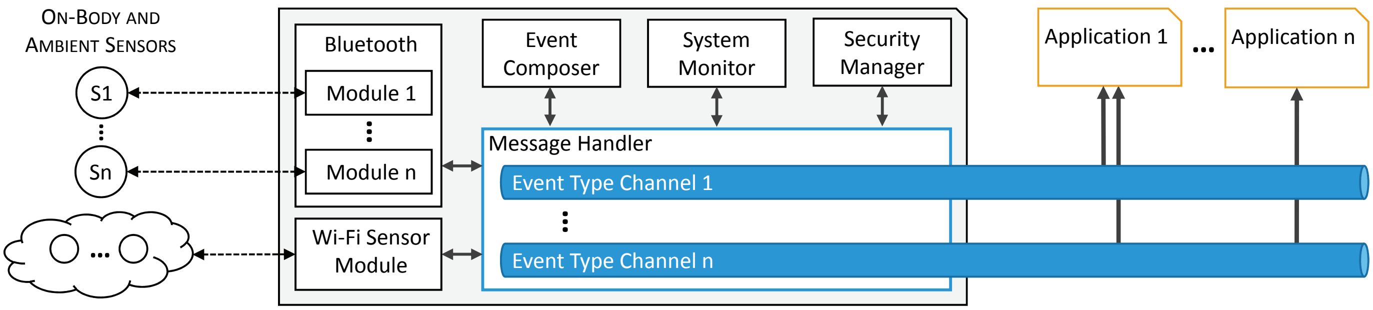

Fig. 1. Event-driven architecture mediating between on-body and ambient sensors and multiple applications. The system's *message handler* utilizes Android `BroadcastReceivers` for inter-application communication. *Sensor modules* translate raw sensor data to sensor events and the *event composer* provides additional sensor fidelity information. The *system monitor* monitors the overall system status and the *security manager* provides information for access control.

to the system: In our implementation, ambient sensors are connected using Wi-Fi, while on-body sensors (*S1*, *Sn*) use the Bluetooth (BT) protocol as it is supported by most Android phones, though also other communication protocols such as ZigBee can be introduced. A *sensor module* handles the sensor communication and creates a corresponding sensor event, which is forwarded to the *message handler*. The latter injects the received event to broadcast channels with respect to the event hierarchy, thus informing subscribed applications about new sensors readings.

The *event composer* in Fig. 1 interprets incoming events, identifies general situations on which the system has to react and creates a corresponding derived event. Events from our case study's heart rate sensor for instance contain also the current battery level: Upon receiving a sensor reading indicating low battery power, an alarm event is created and sent to the *message handler*. The *event composer* also allows to check for inaccurate or invalid sensor data and emits events enriched with fidelity information [27]. This allows applications to only consider sensor readings reaching a certain data fidelity threshold, and throw away low-fidelity data such as heart rate readings with sensing artefacts or blood pressure readings taken after exertion or with a wrong arm posture during the measurement. The *system monitor* measures the overall system status and detects critical situations such as a low battery level. In order to monitor the phone's overall liveliness, heartbeat messages containing status information are periodically sent to a server, allowing a remote detection of a crashed phone or a bad connection to the network carrier. Figure 2 depicts the Android implementation of myHealthAssistant.

B. Broadcast Channels

Applications are usually interested in sensor readings of a certain type only. Sensor readings are therefore published to broadcast channels according to their type, allowing applications to simply subscribe to the reading types of their interest. When a hierarchy of reading types is available, applications can also subscribe to a group of sensor types (e.g., cardiovascular readings) to receive all information within this group (e.g., heart rate and blood pressure readings). Management information such as battery alarms or changes in sensor connections are propagated in broadcast channels and structured in an event hierarchy as well. To save redundant

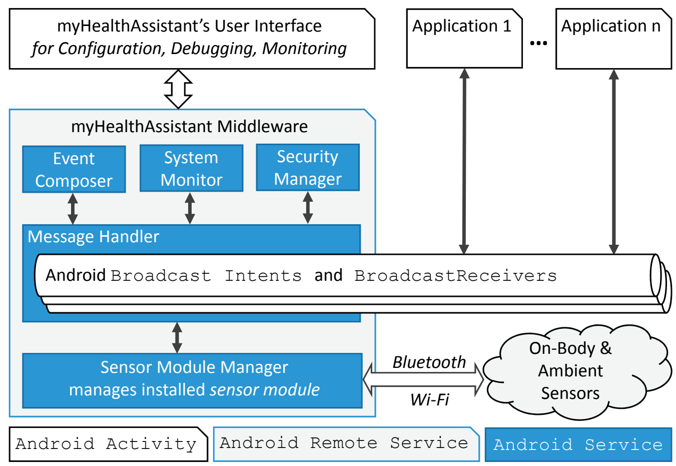

Fig. 2. Android implementation of myHealthAssistant: the middleware is started as an Android `Remote Service` which administers Android `Services` implementing the middleware components. Events sent to applications are encapsulated in `Broadcast Intents` that are received via `BroadcastReceivers`.

and unnecessary calculations, the system also provides a mechanism that allows applications to exchange events.

From an implementation perspective, applications in Android are running in Dalvik process virtual machines that are strictly separated from each other. In order to support inter-process communication, Android provides `BroadcastReceivers` which allow applications to receive information from other applications and the Android system. For receiving information, a BroadcastReceiver needs to be instantiated and registered to the desired broadcast channel (cp. Listing 1). Information is exchanged via Android `Intents` which can contain data of simple data types and `Parcelable` objects. Since events in our system are implemented as Parcelable objects, they are encapsulated in Intents and sent via the Android broadcast service. For an application to inject an event to the system, it thus needs to add this event to an Intent and send it to a receiver using the Android `sendBroadcast()` method (cp. Listing 2). Upon receiving this event, MyHealthAssistant distributes it to the different channels with respect to its event type.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JBHI.2014.2326604, IEEE Journal of Biomedical and Health Informatics

Listing 1. Registering an Android `BroadcastReceiver` (ReadingEventReceiver) for receiving events of type *blood pressure*.

```
mReadingEventReceiver = new ReadingEventReceiver();            // create event receiver
registerReceiver(mReadingEventReceiver,                        // register listener for
    new IntentFilter(SensorReadingEvent.BLOOD_PRESSURE);       // blood pressure readings
```

Listing 2. Sending a reading event (myReadinngEvent) to myHealthAssistant using the Android `sendBroadcast()` method.

```
Intent i = new Intent();                                       // create Intent
i.putExtra(Event.PARCELABLE_EXTRA_EVENT, myReadingEvent);      // add event
i.setAction(MyHealthAssistant.RECEIVER_CHANNEL);               // set channel
getApplicationContext.sendBroadcast(i);                        // send Intent
```

C. Application Development

Our middleware solution handles all sensor communication and provides the information in a common data abstraction for the applications to use. For receiving sensor readings, an application needs to subscribe to the broadcast channel of interest. The communication with sensor units, dealing with reading artifacts, and monitoring the liveliness of the system are responsibilities of MyHealthAssistant. Having a middleware as a layer between sensors and applications furthermore allows the multiplexing of sensor information which could otherwise only have been accessed by one application. The event hierarchy of our current implementation is implemented with a tree structure along which the events are distributed in channels.

There are two ways for adapting to new requirements following ever-changing network protocols and technologies: The first, middleware-centric method is adding a new or modified *sensor module* to the system. For adaptations to a new protocol based on already implemented technology, a single method that implements the new protocol needs to be written. For adapting the system to a new network technology, connection handling and packet retrieval have to be implemented as well. The second method would involve an additional application which manages the sensor communication and then injects the resulting event to the middleware as described above. Since the latter method would result in a lack of control over the application-based sensor module, the middleware-centric method is identified as the preferred one.

IV. Applications

Three applications were built as case studies on top of MyHealthAssistant, inspired by different application areas: Fitness support, tele-monitoring, and elderly care. They illustrate the type of applications that are targeted by MyHealthAssistant and were used in the evaluation.

A. Fitness Support Application

Studies [28], [29] have shown that an Internet and phone-based user motivation system can significantly increase and maintain the level of physical activity. We developed an application that captures the user's activities and monitors the heart rate throughout a day [30]. It consists of two sensor network setups: one for daily activity detection and another setup for capturing exercises during a gym visit (see Fig. 3c). In the first setup, the user wears a single accelerometer (cp.

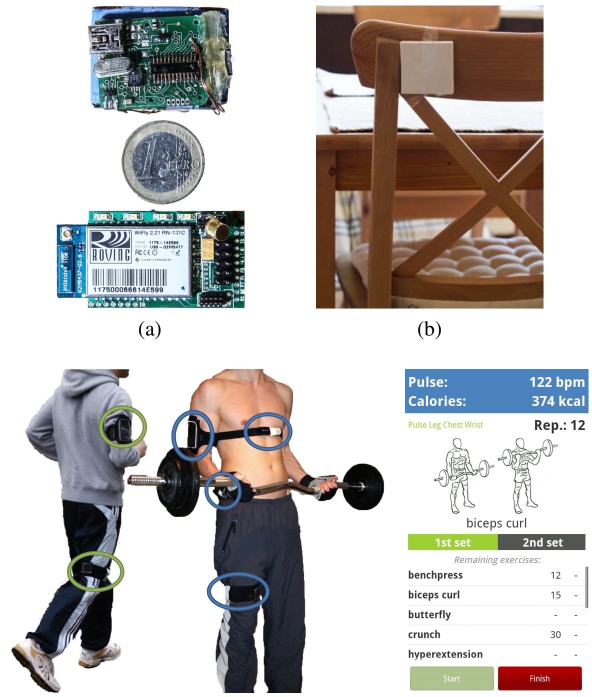

(a) (b) (c) (d)

Fig. 3. Hardware modules used in our case studies: (a) Bluetooth accelerometer at the top and a Wi-Fi module with a ball-in-tube sensor attached at the bottom, (b) Wi-Fi module attached to a chair, (c) two sensor configurations of the fitness application, and (d) a screenshot of the fitness application shown during a weight lifting exercise.

Fig. 3a upper sensor) and a distinction is made between being idle (i.e. sitting or standing), doing moderate movements (i.e. walking, cycling) and doing sports (i.e. running). In the second setup, two extra sensors embedded in weightlifting gloves and chest strap allow the recognition between 16 gym workouts as well as counting exercise repetitions. Fig. 3d shows a screen-shot during a biceps curl exercise. Detailed evaluation of the system [30] showed that the application's recognition performance matches that of state-of-the-art methods, while being capable of reliable activity and heart rate monitoring with real-time user feedback, for at least 12 hours a day.

B. Telemonitoring Application

Telemedicine is a promising application area for body sensor networks, in which patients' health parameters are collected and transferred to a remote healthcare provider, to observe patients over long periods or remotely detect

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JBHI.2014.2326604, IEEE Journal of Biomedical and Health Informatics

dangerous circumstances. Our second application case study utilizes different sensor units to gather the user's weight, blood pressure, heart rate, ECG, and daily activities. The detected activities are correlated with the user's heart rate, allowing an alarm to be sent whenever the heart rate becomes a-typical for the current activity. Sensor readings are stored in a database and transferred to an existing telemedical platform. Reminders are displayed whenever user-initiated measurements (such as those taken by a scale fitted with bluetooth) are due.

C. Elderly Care Application

The elderly care assistance application monitors both vital parameters and user interactions with the environment, which are interpreted and monitored as activities of daily living. If the user misses to do a specific activity (e.g. tooth brushing) within a given time slot, the system reminds the user. In addition, it creates a list of performed activities at the end of a day which can be used in order to detect changes in daily habits.

In addition to the sensor units mentioned in the previous application, the assistance application makes use of ambient sensors, which are low-power Wi-Fi modules equipped with either a ball-in-tube sensor (cp. Fig. 3a at the bottom), a reed switch, or a passive infrared sensor. Details about the modules are presented in [31]. Upon recognizing a movement, the sensor units send HTTP POST messages to the phone and these events are used together with rules to derive the activities. If for instance the chair in the dining area (cp. Fig. 3b) is used within 45 minutes after the cutlery drawer and the fridge were opened, the application assumes that the person is eating. The system was evaluated in different apartments for activities such as tooth brushing, showering, airing, eating, entering and leaving the apartment, cooking, and desk work, as will be described in Section V. The system makes use of the existing Wi-Fi network infrastructure preserving the phone's Wi-Fi connectivity.

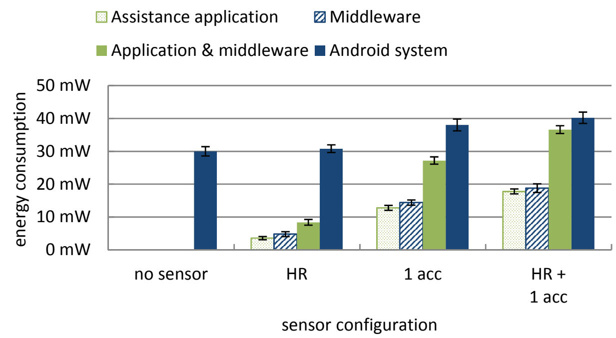

(a) assistance application

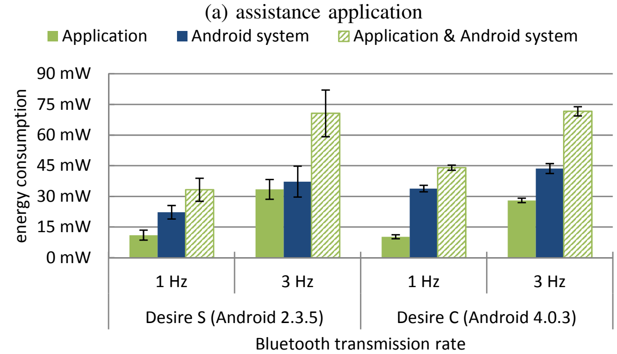

(b) impact of different transmission rates

Fig. 4. (a) Shows the energy consumption of the assistance application for different sensor constellations. (b) Depicts the energy consumption for two different Bluetooth transmission rates (1 Hz vs. 3 Hz). Increasing the transmission rate drastically increases the overall energy consumption.

Start	Stop	Duration	Remaining battery power
9:45	22:15	12 hours 30 minutes	30%
7:20	21:20	14 hours 00 minutes	20%
9:45	22:00	12 hours 15 minutes	50%
6:10	22:20	16 hours 10 minutes	15%
6:10	21:20	15 hours 10 minutes	20%

TABLE I
THE REMAINING BATTERY LEVELS AFTER EACH DAY OF USING MYHEALTHASSISTANT WITH 11 SENSOR UNITS FOR PHONE #2 .

V. SYSTEM EVALUATION

The main task of our BASN middleware is to collect measurements from the system's sensors and to direct this information to the applications. An important aspect is that the whole system should last for at least a day without re-charging any of the components. We will therefore analyze and discuss in this section the properties of our system with respect to three requirements: energy consumption, performance for different sensor constellations, and execution speed performance for an increasing number of subscribed applications. We chose two representatives of the current smartphone market: **(#1)** HTC Desire S, 1 GHz CPU, 768 MB RAM, 1450 mAh battery, Android version 2.3.5 and **(#2)** HTC Desire C, 600 MHz CPU, 512 MB RAM, 1230 mAh battery, Android version 4.0.3.

A. Energy Consumption

The operating time of a BASN system is critical for its usability and user acceptance. We analyzed the energy consumption of the the third case study application (elderly care), as it is the most power-consuming. Additionally, impact of increased transmission rates are analyzed and energy consumption figures are measured using PowerTutor [32].

Fig. 4(a) depicts the energy consumption of four different sensor configurations and our assistance application running on phone #1. The heart rate sensor sends 59 bytes per second whereas the accelerometer is sending only 6 bytes per second, but triggers the activity recognition which explains the higher energy consumption. It can be observed that with an increasing message workload, the energy consumption of the application, MyHealthAssistant, as well as the Android system increases. The energy overhead for Android is about 30% for the additional communication tasks.

Since the energy consumption of single processes does not indicate how long a device can operate, we measured the remaining battery level at the end of each day. The setup consisted of 6 ambient sensors for detecting interactions with the environment, a heart rate sensor, a scale, a blood pressure sensor, and an accelerometer for continuous activity detection. Weight and blood pressure readings were taken in the morning and in the evening and the accelerometer was worn continusouly, while the heart rate sensor was worn sporadically. Table I shows the results after one week of monitoring with phone #2: With a significant processing overhead of performing near real-time activity detection, the phone's battery level remained at 15% after 16 hours of operating. On day three, the accelerometer ran out of power, hence the 50% remaining battery level compared to the 30% of day one.

Wireless communication has the biggest impact for the system's energy consumption; MyHealthAssistant largest contribution here comes from avoiding redundant communication between multiple applications with the same sensor unit. Fig. 4 furthermore illustrates the importance of avoiding duplicate or redundant transmission in a comparison of the energy consumption between having a Bluetooth sensor sending measurements at 1 Hz and at 3 Hz for phones #1 and #2.

B. Performance

We will first analyze MyHealthAssistant's performance under an increasing workload: This is done by instantiating a growing amount of event generators, each injecting events with a 1 Hz frequency. The second case study application described in Section IV-B is subscribed to these events and logs the incoming events in order to calculate the delivery ratio. Since our event generators do not require communication with sensors, we also analyze the system's performance in an hybrid approach consisting of real sensors together with event generators. After we discussed the system's behavior on events containing only a single value as payload, we will increase the payload size of the events. In a last test, we will observe the system's behavior on multiple applications being subscribed to it. All results presented in this section are averages over three or more test runs.

1) Number of Events: As a first step, we will observe how the system behaves for an increasing amount of events per second. Since the Bluetooth protocol is limited to seven active connections, event generators were used for this analysis that run in the same process as the middleware and operate in the same way as the *sensor modules* described in Section III. Each event generator injects events once a second, consisting of following fields: `ID`, `type`, `timestamp`, `producer ID`, `sensor type`, `time of measurement`, and an integer value as the `payload`.

Fig. 5 depicts the results for phone #1 and #2, showing that CPU utilization of phone #1 fluctuates around 2% for the whole test: MyHealthAssistant is marginally affected by the increasing workload, as the actual event distribution is done by the Android system while our middleware solution only decides on which channels the event is sent. The throughput is thus limited by the Android system. In our case, the event delivery ratio drops below 99.9% for more than 12 events per second and reaches 98.7% for a workload of 30 events/s. In contrast, the slower phone #2 running Android 4.0.3 shows a different behavior: The CPU utilization increases with an increasing workload, which is likely due to Android 4 performing some tasks within the application process. Most of the processing is done by the operating system, however, and the delivery ratio rapidly drops for more than 20 events/s. The system's memory usage of 5.7 MB/7.4 MB is relatively modest compared to the phone's RAM of 512/768 MB (the slight increase is due to the growing amount of event generators).

In summary, since communication among Android applications goes via the operating system, the maximum number of events per second is limited by the capabilities of Android. With an increasing workload, the delivery ratio starts dropping. For a minimum delivery ratio of 99.9% up to 12 events/s are handled. We believe that this is sufficient for most BASN applications since the energy consumption of wireless communication is the more limiting factor. Furthermore, the phones we used for our analysis are relatively slow compared to current phones with multi-core processors. Faster hardware is expected to speed up this inter-process communication.

2) Hybrid: Bluetooth Sensors and Event Generators: Fig. 6 depicts the test results for a hybrid setup compared to a setup with event generators only. For the hybrid setup, we connected a heart rate sensor and an accelerometer to phone #1. Furthermore, we tested the impact of activity recognition being activated. Having Bluetooth sensors connected to the system increases both the CPU utilization as well as the memory usage due to the additional overhead evoked by the Bluetooth communication, while the impact on the delivery ratio is marginal. An enabled activity recognition leads to a slight impact, mostly because the detected activities are sent to the middleware, thus resulting in an increased workload.

3) Event Size: The events injected so far were consisted of either one or six integer values (the accelerometer sends 6 values/s). Fig. 7 depicts the system's behavior for bigger event sizes. We increased the payload size of the injected events from 1 to 200 integer values. The CPU utilization is not affected by the event size since MyHealthAssistant decides to which channel an event has to be sent only based on the event type and does not inspect the payload. The delivery ratio decreases slightly, which can be explained by the higher amount of transferred data which leads to a slight increase in memory usage: Increasing event sizes have only a marginal impact on the system.

4) Applications: The support of multiple applications is a big advantage of having a middleware solution like MyHealtAssistant as a layer between applications and sensors. We therefore tested our system for four different workload setups and up to eleven Android `BroadcastReceivers` subscribed to the events and running in different processes. As Fig. 8 shows, the amount of subscribed applications has little impact on the system: CPU utilization and message delivery ratio remain while the memory usage increases slightly.

VI. Conclusions and Future Work

A growing variety of on-body and ambient sensor units is leading to new and promising ways to monitor patients in their natural surroundings. These sensors also result in several challenges to application developers, however, as the increasing heterogeneity of data formats, protocols, and communication channels hinders them in a swift application development. As the currently-available platforms are embedded systems that operate on batteries, an additional challenge is the limitation of available resources. We propose a middleware solution,

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JBHI.2014.2326604, IEEE Journal of Biomedical and Health Informatics

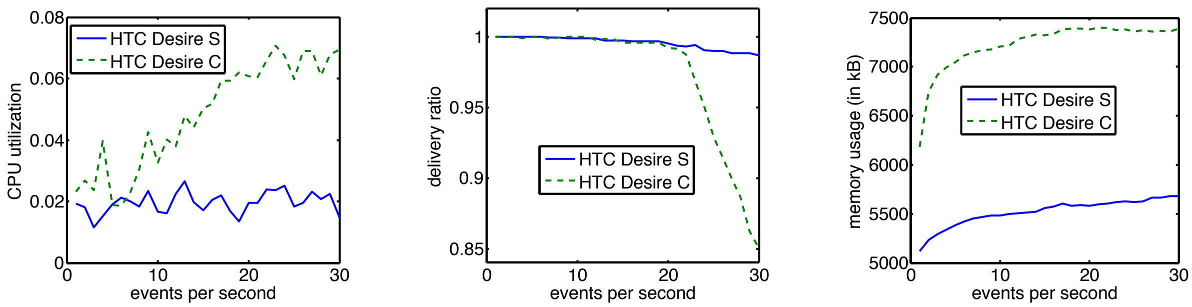

(a) CPU utilization (b) event delivery ratio (c) memory usage

Fig. 5. Performance analysis for an increasing amount of event generators. For a low-end phone (HTC Desire C, 600 MHz CPU, 515 MB RAM) and more than 20 events per second, the delivery ratio starts dropping rapidly.

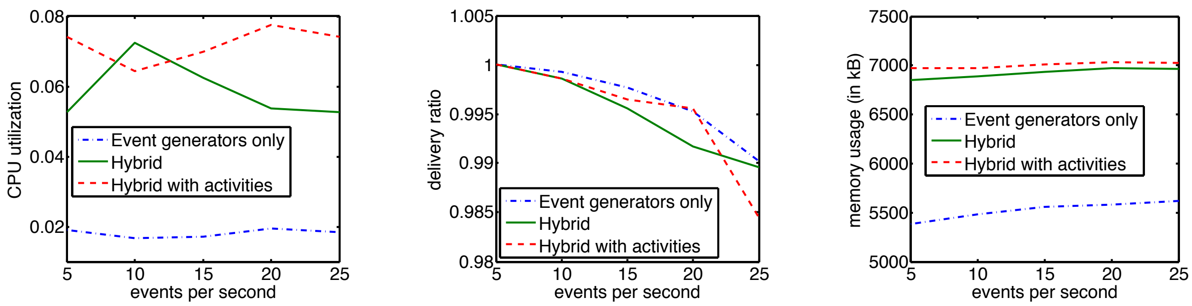

(a) CPU utilization (b) event delivery ratio (c) memory usage

Fig. 6. Comparison of the system's performance for: I) event generators, II) a hybrid setup with an increasing amount of event generators from a Bluetooth heart rate sensor and a Bluetooth accelerometer, and III) the hybrid setup plus activity recognition as described in Section IV-C.

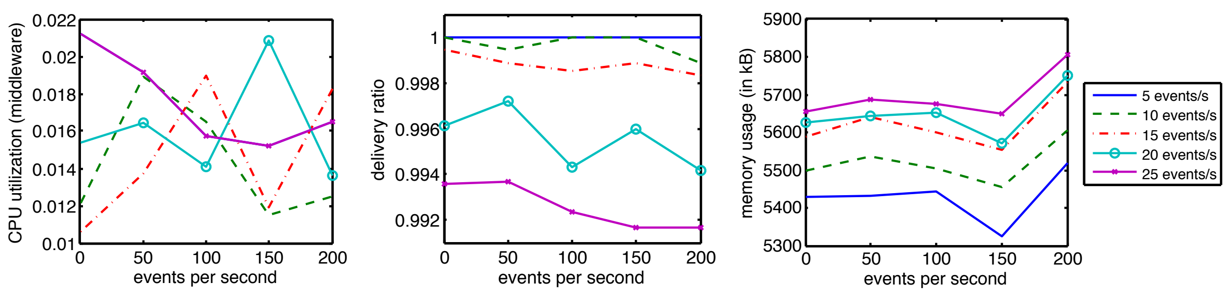

(a) CPU utilization (b) event delivery ratio (c) memory usage

Fig. 7. Increasing event size from 1 to 200 integer values per event and different throughput configurations.

MyHealthAssistant, which is designed for phone-based deployment and focuses on the efficient management of wireless sensor data for multiple healthcare applications.

The event-driven middleware architecture is designed to aggregate and provides information from both body sensor networks and ambient sensor networks to subscribed applications via broadcast channels. The liveliness of the phone as well as individual sensors is monitored and an event composer module provides the calculation of fidelity levels for sensor readings. A back-channel allows applications to share events to avoid redundant processing.

We evaluated MyHealthAssistant with respect to energy consumption, message throughput and served applications. It was shown that the burden of hosting the middleware solution as well as a monitoring application including activity recognition on a usual Android phone is low enough for at least 16 hours of monitoring. Additionally, using a case study analysis of three target applications implemented on our system, we have shown that performance is sufficient to enable many applications on current phone models, assuming that the system is charged overnight. For high requirements on the event delivery ratio (at least 99.9%) on a single-core phone, the maximum message throughput is limited to 12 events per second whereas the event size does not impact the delivery ratio. We believe that this is sufficient for many current applications since most sensors aggregate their data before transmission. Our ECG sensor, for instance, sends readings twice a second including 107 values per message. We had up to 11 applications subscribed to our middleware with no impact on its performance.

The accelerometer modules using in the case studies are open-source (both hardware and firmware) and publicly available for download[1]. The proposed middleware architecture's

[1] HedgeHog Project: http://www.ess.tu-darmstadt.de/hedgehog

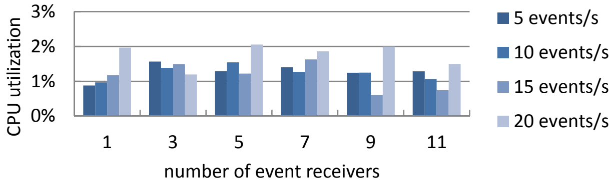

(a) CPU utilization

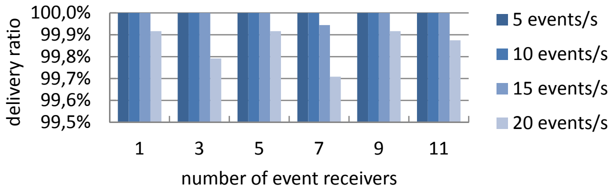

(b) event delivery ratio

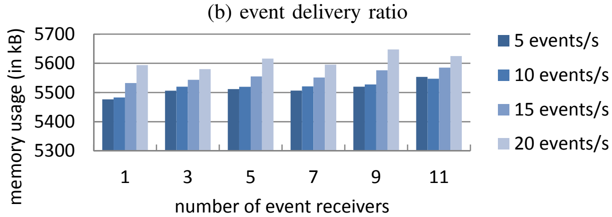

(c) memory usage

Fig. 8. System behavior on an increasing amount of subscribed applications: no impact for (a) and (b); a marginal impact for (c).

Android implementation will be available as open-source software[2]. The telemedicine platform is a prototype kindly provided to us for this project by Robert Bosch GmbH.

REFERENCES

[1] European Commision Staff, *Demography report 2010*. European Commission, 2011.

[2] L. B. Shrestha and E. J. Heisler, *Changing Demographic Profile of the United States*. Congressional Research Service, 2011.

[3] World Health Organization, "Integrating prevention into health care," http://www.who.int/chp/about/integrated_cd/en/, 2012, [03-23-2014].

[4] M. Chen, S. Gonzalez, A. Vasilakos, H. Cao, and V. C. Leung, "Body area networks: A survey," *Mobile Network Applications*, vol. 16, no. 2, pp. 171–193, Apr. 2011.

[5] P. Khan, A. Hussain, and K. S. Kwak, "Medical Applications of Wireless Body Area Networks," *International Journal of Digital Content Technology and its Applications*, vol. 3, no. 3, pp. 185–193, 2009.

[6] P. Neves, M. Stachyra, and J. Rodrigues, "Application of wireless sensor networks to healthcare promotion," *Journal of Communications Software and Systems (JCOMSS)*, vol. 4, no. 3, pp. 181–190, 2006.

[7] National Heart Lung and Blood Institue, "Tips for Having Your Blood Pressure Taken," http://www.nhlbi.nih.gov/hbp/detect/tips.htm, 2013.

[8] Partnership of the Heart, "Telemedical interventional monitoring," http://www.partnership-for-the-heart.de/en/, 2013.

[9] W. Heinzelman, A. Murphy, H. Carvalho, and M. Perillo, "Middleware to support sensor network applications," *Network, IEEE*, vol. 18, no. 1, pp. 6 – 14, jan/feb 2004.

[10] S. L. Keoh, N. Dulay, E. Lupu, K. Twidle, A. Schaeffer-Filho, M. Sloman, S. Heeps, S. Strowes, and J. Sventek, "Self-managed cell: A middleware for managing body-sensor networks," in *Intl. Conf. on Mobile and Ubiq. Systems: Networking Services*, 2007, pp. 1–5.

[11] A. Waluyo, S. Ying, I. Pek, and J. K. Wu, "Middleware for wireless medical body area network," in *Biomedical Circuits and Systems Conference, 2007. BIOCAS 2007. IEEE*, nov. 2007, pp. 183 –186.

[12] M. Wagner, B. Kuch, C. Cabrera, P. Enoksson, and A. Sieber, "Android based Body Area Network for the evaluation of medical parameters," in *Proceedings of the Tenth Workshop on Intelligent Solutions in Embedded Systems (WISES)*, 2012, pp. 33–38.

[2]The MyHealthAssistant Website: http://www.dvs.tu-darmstadt.de/research/myhealthassistant/

[13] P. Kuryloski, A. Giani, R. Giannantonio, K. Gilani, R. Gravina, V.-P. Seppa, E. Seto, V. Shia, C. Wang, P. Yan, A. Yang, J. Hyttinen, S. Sastry, S. Wicker, and R. Bajcsy, "Dexternet: An open platform for heterogeneous body sensor networks and its applications," in *Sixth International Workshop on Wearable and Implantable Body Sensor Networks (BSN)*, 2009, pp. 92–97.

[14] D. Carr, M. O'Grady, G. O'Hare, and R. Collier, "SIXTH: A Middleware for Supporting Ubiquitous Sensing in Personal Health Monitoring," in *International Workshop on Advances in Personalized Healthcare Services*. Springer, 2012.

[15] J. Rodríguez-Molina, J.-F. Martínez, P. Castillejo, and L. López, "Combining wireless sensor networks and semantic middleware for an Internet of Things-based sportsman/woman monitoring application." *Sensors (Basel, Switzerland)*, vol. 13, no. 2, pp. 1787–835, Jan. 2013.

[16] A. B. Waluyo, W.-S. Yeoh, I. Pek, Y. Yong, and X. Chen, "MobiSense: Mobile Body Sensor Network for Ambulatory Monitoring," *ACM Transactions on Embedded Computing Systems*, vol. 10, no. 1, pp. 1–30, Aug. 2010.

[17] V. T. Bui, R. Verhoeven, and J. J. Lukkien, "A Body Sensor Platform for concurrent applications," *IEEE Second International Conference on Consumer Electronics - Berlin (ICCE-Berlin)*, pp. 38–42, Sep. 2012.

[18] R. Kamal, N. Tran, and C. Hong, "Event-based middleware for healthcare applications," *IEEE Journal of Communications and Networks*, vol. 14, no. 3, pp. 296–309, 2012.

[19] V. Jones, A. Van Halteren, N. Dokovsky, G. Koprinkov, J. Peuscher, R. Bults, D. Konstantas, W. Ing, and R. Herzog, *MobiHealth: Mobile Services for Health Professionals, in M-Health: Emerging Mobile Health Systems*. Springer, 2006.

[20] M. J. Morón, A. Gómez-Jaime, J. R. Luque, and E. Casilari, "Development and Evaluation of a Python Telecare System Based on a Bluetooth Body Area Network," *EURASIP Journal on Wireless Communications and Networking*, vol. 2011, pp. 1–10, 2011.

[21] D. Guerra, U. Gawlick, and P. Bizarro, "An integrated data management approach to manage health care data," *DEBS*, pp. 40:1–40:2, 2009.

[22] J. Singh and J. Bacon, "Event-based data dissemination control in healthcare," *Electronic Healthcare*, pp. 167–174, 2009.

[23] C. Seeger, A. Buchmann, and K. Van Laerhoven, "An event-based bsn middleware that supports seamless switching between sensor configurations," in *2nd ACM SIGHIT Intl. Health Informatics Symposium*, 2012.

[24] T. Freudenreich, S. Appel, S. Frischbier, and A. Buchmann, "Actress - automatic context transformation in event-based software systems," in *Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems*. ACM, July 2012, pp. 179–190.

[25] G. S. Blair, A. Bennaceur, N. Georgantas, P. Grace, V. Issarny, V. Nundloll, and M. Paolucci, "The role of ontologies in emergent middleware: supporting interoperability in complex distributed systems," in *ACM/IFIP/USENIX Intl. Middleware Conf.* Springer, 2011, pp. 410–430.

[26] A. Bennaceur, G. Blair, F. Chauvel, H. Gang, N. Georgantas, P. Grace, F. Howar, P. Inverardi, V. Issarny, M. Paolucci, A. Pathak, R. Spalazzese, B. Steffen, and B. Souville, "Towards an architecture for runtime interoperability," in *4th intl. conf. on leveraging applications of formal methods, verification, and validation*. Berlin, Heidelberg: Springer, 2010, pp. 206–220.

[27] C. Seeger, A. Buchmann, and K. Van Laerhoven, "Wireless sensor networks in the wild: Three practical issues after a middleware deployment," in *the Sixth International Workshop on Middleware Tools, Services and Run-time Support for Networked Embedded Systems (MidSens 2011)*, ACM Press. Lisbon, Portugal: ACM Press, 12/2011 2011.

[28] R. Hurling, M. Catt, M. De Boni, W. B. Fairley, T. Hurst, P. Murray, A. Richardson, and S. J. Sodhi, "Using Internet and Mobile Phone Technology to Deliver an Automated Physical Activity Program: Randomized Controlled Trial," *J. of Med. Internet Research*, vol. 9, no. 2, Apr. 2007.

[29] D. F. Tate, R. R. Wing, and R. A. Winett, "Using Internet Technology to Deliver a Behavioral Weight Loss Program," *Journal of the American Medical Association*, vol. 285, no. 9, pp. 1172–1177, Mar. 2001.

[30] C. Seeger, A. Buchmann, and K. Van Laerhoven, "myhealthassistant: A phone-based body sensor network that captures the wearer's exercises throughout the day," in *The 6th International Conference on Body Area Networks*, ACM Press. Beijing, China: ACM Press, 11/2011 2011.

[31] B. Ostermaier, M. Kovatsch, and S. Santini, "Connecting things to the web using programmable low-power wifi modules," in *2nd Intl. Workshop on the Web of Things*. ACM, Jun. 2011.

[32] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R. P. Dick, Z. M. Mao, and L. Yang, "Accurate online power estimation and automatic battery behavior based power model generation for smartphones," in *8th IEEE/ACM/IFIP Intl. Conf. on Hardware/software Codesign and System Synthesis*. ACM, 2010, pp. 105–114.

Vrije Universiteit Brussel

FACULTEIT VAN DE WETENSCHAPPEN

DIPLOMA

WIJ, RECTOR VAN DE VRIJE UNIVERSITEIT BRUSSEL, VERLENEN NA BESLISSING VAN DE EXAMENCOMMISSIE, AAN

Kristof Van Laerhoven

GEBOREN TE KAPELLEN OP 16 NOVEMBER 1975

de academische graad van

LICENTIAAT IN DE INFORMATICA

BEHAALD MET GROTE ONDERSCHEIDING

Wij kennen dit diploma toe conform de bepalingen van het Decreet van 12 juni 1991 betreffende de universiteiten in de Vlaamse Gemeenschap

Brussel, 9 juli 1999

ELS WITTE
Rector

Bij dit diploma hoort een supplement dat de studiegegevens vermeldt.

051099135458526

THE UNIVERSITY OF LANCASTER

It is hereby certified that

Kristof Van Laerhoven

has been duly admitted to the degree of

Doctor of Philosophy

in

Computer Science

Vice-Chancellor

Secretary of the University

March, 2006

0165603-00035274-1