DAGSTUHL
REPORTS

Volume 2, Issue 12, December 2012

Dagstuhl Reports, Vol. 2, Issue 12 ISSN 2192-5283

ISSN 2192-5283

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany.
Online available at http://www.dagstuhl.de/dagrep

Publication date
May, 2013

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at http://dnb.d-nb.de.

Aims and Scope
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e.g. summaries from panel discussions or open problem sessions.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de

Digital Object Identifier: 10.4230/DagRep.2.12.i

Report from Dagstuhl Seminar 12491

Interpreting Observed Action

Edited by

Susanne Biundo-Stephan[1], Hans Werner Guesgen[2], Joachim Hertzberg[3], and Stephen R. Marsland[4]

1 **Universität Ulm, DE**
2 **Massey University, NZ,** h.w.guesgen@massey.ac.nz
3 **Universität Osnabrück and DFKI RIC, Osnabrück, DE,** joachim.hertzberg@uos.de
4 **Massey University, NZ,** s.r.marsland@massey.ac.nz

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12491 "Interpreting Observed Action". The aim of the seminar was to get a coherent picture, which transcends the borders of applications and disciplines, of existing approaches and problems in interpreting observed action in semantic terms – primarily action by humans, but action by artificial agents may play some role, too. The seminar brought together, on the one hand, researchers from the different camps of AI, robotics, and knowledge-based systems who are working on the various aspects and purposes of interpreting observed action by humans, or occasionally, other agents; on the other hand, it added some researchers from cognitive science (psychology, neurosciences) working on human perception of behaviour and action. The main outcome of the seminar were a set of guidelines for setting up a workbench, which can be used to explore and test methods and techniques related to interpreting observed action.

Seminar 2.–6. December, 2012 – www.dagstuhl.de/12491
1998 ACM Subject Classification I.2 Artificial Intelligence
Keywords and phrases action, knowledge representation, plan recognition, symbol grounding, perception, behavior interpretation
Digital Object Identifier 10.4230/DagRep.2.12.1

1 Executive Summary

Susanne Biundo-Stephan
Hans Werner Guesgen
Joachim Hertzberg
Stephen R. Marsland

For many applications of smart embedded software systems, the system should sense the footprint of a human or humans acting in the system's environment, interpret the sensor data in terms of some semantic model about what the human is doing, and respond appropriately in real time. Examples of such applications include smart homes, human-machine or human-robot interaction, assistance, surveillance, and tutoring systems; given the current trend towards ambient intelligence, ubiquitous computing, and sensor networks, the number of systems in these categories can certainly by expected to rise in the next ten years or so.

The problem shares many features with classical object recognition and scene reconstruction from sensor data in terms of a static scene model. Interpreting in semantic terms sensor data from the environment has a long tradition in AI – arguably, it has been one of the original core problems put forth by AI's founding fathers. However, the problem of interpreting observed action in the sense of this seminar differs in some aspects from what state-of-the-art AI or engineering approaches would allow to be tackled by routine:

- **Events in space-time** rather than static objects need to be characterized. This necessarily involves some representation and model of temporal and spatial data (e.g., the human put a saucepan *on* the cooker *and then* turned the cooker on).
- **Real-time processing** of the sensor data or percepts is required to keep track of what is happening. In fact, "real time" here is the pace of human action, i.e., relatively slow compared to CPU clock ticks. However, given a potentially rich stream of sensor data and a potentially large body of background knowledge, even this pace is demanding for the respective knowledge processing methods.
- **Willed human action**, be it planned, intended, or customary, is the domain of interpretation. In knowledge representation, this appears to be a relatively unexplored area, compared to, say, upper ontologies of household items, red wine, or pizza varieties.

Contemplating the three words that make up the title of this seminar ("interpret", "observe", and "act"), it becomes clear that there are a number of issues that need to be addressed in this context. Firstly, any interpretation is to some degree subjective and uses a particular repertoire of basic actions in its language. Secondly, an observation uses a particular type of sensor data and often is not possible without interpretation at the same time. Thirdly, there are issues around what actions are to be considered:

- Are only willed and physical action to be considered?
- Is avoidance an action?
- What constitutes an action in the first place?
- When does a particular action end?
- Is an unsuccessful action an action?

In summary, what precisely is observed action interpretation and what would be benchmark data for it?

To find an answer to this question, the participants of the seminar emerged themselves in a variety of activities: technical talks, working groups, plenary discussions, and a number of informal discussions. In the rest of this report, some of these activities and their results are discussed in more detail.

2 Table of Contents

3 Overview of Talks

3.1 Recognizing Users' Intentions – A Key Competence of Companion-Systems

Susanne Biundo-Stephan (Universität Ulm, DE)

Development of a Companion-Technology aims at enabling the realization of technical systems that provide their functionality in a completely individualized way. They adapt to a user's expertise, background, capabilities, and needs; furthermore, they take into account the current situation as well as the user's emotional state. The features that distinguish Companion-systems, namely individuality, adaptiveness, availability, cooperativeness, and trustworthiness, are realized by (the interplay of) cognitive processes. These include planning and decision making, interaction and dialog, and perception and recognition.

Intention recognition is crucial for Companion-systems. It serves three purposes: (1) monitor whether the user acts as expected; (2) detect and explain / interpret deviations from expected behavior; (3) initiate appropriate measures to avoid the break of communication. As Companion-systems are knowledge-based systems, intention recognition can rely upon various sources including background domain knowledge, interaction history, situation and action context. Intention recognition involves various cognitive levels. Elementary activities and emotional state of a user are recognized on the sensor data processing level. Having "perceived" basic actions of the planning level this way, higher-level plans of action can be identified to serve as hypotheses for the recognition of action strategies and goals that represent the user's intentions. Finally, predictions generated on the planning level are fed back into the sensor data processing to guide activity recognition.

3.2 How Do We Interact With the World? Objects, Spaces and Interactions

Martin V. Butz (Universität Tübingen, DE)

Joint work of Butz, Martin V.; Ehrenfeld, S.; Herbort, O.
Main reference S. Ehrenfeld, M.V. Butz, "The modular modality frame model: Continuous body state estimation and plausibility-weighted information fusion," Biological Cybernetics, Vol. 107, Issue 1, pp. 61–82, Springer, 2012.
URL http://dx.doi.org/10.1007/s00422-012-0526-2

Various results are put forward that we do not really interact optimally with our environment – and particularly with objects. Other factors can rather easily influence our interaction kinematics and dynamics. Moreover, a model is presented in which such interactions can unfold and which integrates in a highly modularized manner body state estimations, sensory information, and sensorimotor predictions. Finally, by means of results from an eye-tracking experiment, it is shown how anticipations about object interactions guide our information search task-dependently. I conclude that we interact with the world in a highly anticipatory fashion, continuously integrating interaction knowledge and other biases as well as incoming sources of information in a weighted, integrative manner.

3.3 Learning Relational Event Models From Video

Krishna Sandeep Reddy Dubba (University of Leeds, GB)

Main reference K.S.R. Dubba, A.G. Cohn, D.C. Hogg, "Learning Event Models From Complex Videos Using ILP," in Proc. of the 19th European Conf. on Articial Intelligence (ECAI'10), Frontiers in AI, Vol. 215, pp. 93–98, IOS Press, 2010.

Learning event models from videos has applications ranging from abnormal event detection to content based video retrieval. When multiple agents are involved in the events, characterizing events naturally suggests encoding interactions as relations. This can be realized by tracking the objects using computer vision algorithms and encoding the interactions using qualitative spatial and temporal relations. Learning event models from this kind of relational spatio-temporal data is particularly challenging because of the presence of multiple objects, uncertainty from the tracking and especially the time component as this increases the size of the relational data (the number of temporal relational facts is quadratically proportional to the number of intervals present).

Relational learning techniques such as Inductive Logic Programming (ILP) hold promise for building models from this kind of data, but have not been successfully applied to the very large datasets which result from video data. In this thesis, we present a novel supervised learning framework to learn relational event models from large video datasets (several million frames) using ILP. Efficiency is achieved via the learning from interpretations setting and using a typing system that exploits the type hierarchy of objects in a domain.

Positive and negative examples are extracted using domain experts' minimal event annotations (termed deictic supervision) which are used for learning relational event models. These models can be used for recognizing events from unseen videos. If the input data is from sensors, it is prone to noise and to handle this, we present extensions to the original framework by integrating abduction as well as extending the framework based on Markov Logic Networks to obtain robust probabilistic models that improve the event recognition performance.

The experimental results on video data from two challenging real world domains (an airport domain which has events such as loading, unloading, passengerbridge parking etc. and a verbs domain which has verbs like exchange, pick-up etc.) suggest that the techniques are suitable to real world scenarios.

3.4 STRANDS: Spatial-Temporal Representations and Activities for Cognitive Control in Long-Term Scenarios

Tom Duckett (University of Lincoln, UA)

Joint work of Hawes, Nick (STRANDS project Coordinator) and colleagues in the STRANDS consortium

"STRANDS" (Spatial-Temporal Representations and Activities for Cognitive Control in Long-Term Scenarios) is a new FP7 IP Project, which will run from April 2013 to March 2017, involving six academic institutes and two industrial partners across four European countries.

The project aims to enable mobile service robots to achieve robust and intelligent behaviour in human environments through adaptation to, and the exploitation of, long-term experience. Our approach is based on understanding 3D space and how it changes over time, from milliseconds to months. We will develop novel approaches to extract quantitative and qualitative spatio-temporal structure from sensor data gathered during months of autonomous operation. Extracted structure will include recurring geometric primitives, objects, people, and models of activity. We will also develop control mechanisms which exploit these structures to yield adaptive behaviour in highly demanding, real-world security and care scenarios.

3.5 Visual Attention for Mobile Systems

Simone Frintrop (Universität Bonn, DE)

Visual attention is one of the concepts of human perception that focuses the processing capabilities on the regions of a scene that are most promising. Such a mechanism is not only valuable for humans, but also for computer vision and robotic systems. Especially robots that act in an unknown, complex environment, have to prioritize which aspect of the sensory input to process first. Here in Dagstuhl, I presented an overview of our research on computationally modeling visual attention as well as some applications for intelligent vision system. For example, I introduced our work on saliency detection based on multivariate probability distributions and our current approach for detecting unknown objects in 3D scenes.

3.6 Internal Simulations for Behaviour Recognition

Verena V. Hafner (HU Berlin, DE)

Joint work of Schillaci, G.; Lara, B.; Hafner, Verena V.
Main reference G. Schillaci,B. Lara, V.V. Hafner, "Internal Simulations for Behaviour Selection and Recognition," in Proc. of the 3rd Int'l Workshop on Human Behaviour Understanding (HBU'12), LNCs, Vol. 7559, pp. 148–160, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-34014-7_13

Humans are experts at recognising and identifying the behaviour of others. It is believed that they run internal simulations in order to simulate certain (sensorimotor) actions internally when the (visual) signal is noisy or delayed which is the case most of the times. In this study, we provide a computational model consisting of pairs of learned inverse and forward models on a humanoid robot. This allows the robot to chose its actions and recognise the reaching actions of a human.

3.7 Mutual Understanding of Humans and Robots

Alexandra Kirsch (Universität Tübingen, DE)

Joint work of Kirsch, Alexandra; Michael Karg; Christina Lichtenthäler; Thibault Kruse

When two partners interact, they have to understand the actions of the other. For a robot interacting with a human this means that it has to

- understand the situation and in particular the actions of the human partner,
- select appropriate actions and execute them in a legible way, and
- interpret the user's reaction towards the robot actions as feedback.

I present examples from our work on all three aspects. In the area of plan recognition we try to identify everyday activities based on locations, durations of standing at the locations, and objects. These observations are relatively easy to obtain and can already lead to decent plan recognition and useful predictions, even without accurately recognizing single user actions. When choosing actions the robot must take into account the user's needs and social conventions. In the area of robot navigation we have seen that typical robot navigation approaches can lead to illegible behavior. When explicitly taking into account the human approach direction and distance, a more natural robot behavior can be generated. There are currently no generally accepted metrics and test procedures for human-aware behavior. We attempt to develop objective measures that can be used to evaluate robot behavior, and have additionally the potential to be observed by the robot itself, so that the robot can improve its behavior based on implicit user feedback.

3.8 Interpreting Observed Action in Dynamic Human-Robot Teams under Asymmetric Agency and Social Sentience

Geert-Jan M. Kruijff (DFKI – Saarbrücken, DE)

Main reference H. Zender, M. Janicek, G.-J. Kruijff, "Situated Communication for Joint Activity in Human-Robot Teams," in IEEE Intelligent Systems, Vol. 27, Issue 2, IEEE CS, 2012.
URL http://dx.doi.org/10.1109/MIS.2012.8

The talk considers human-robot teaming, particularly the aspect of how a robot can recognize, and decide how to act, when things go wrong. As they inevitably will. The talk introduces the notions of asymmetric agency and social sentience, to discuss how we could model robots as team members.

3.9 Model-free Behaviour Recognition

Stephen R. Marsland (Massey University, NZ)

Joint work of Marsland, Stephen; Guesgen, Hans W.; the MUSE group
Main reference H.W. Guesgen, S. Marsland, (Eds.) "Human Behavior Recognition Technologies: Intelligent Applications for Monitoring and Security," ISBN 9781466636828, IGI Global, 2013.
URL http://dx.doi.org/10.4018/978-1-4666-3682-8
URL http://www.igi-global.com/book/human-behavior-recognition-technologies/72160

An overview of our approach to behaviour recognition in smart homes, including our unsupervised approach, which is based on the concept that behaviours are activities that are repeated, probably with minor variations.

3.10 Activity Recognition with SCENIOR

Bernd Neumann (Universität Hamburg, DE)

Joint work of Bohlken, Wilfried; Hotz, Lothar; Koopmann, Patrick; Neumann, Bernd;
Main reference W. Bohlken, B. Neumann, L. Hotz, P. Koopmann, "Ontology-Based Realtime Activity Monitoring Using Beam Search," in Proc. of the 8th Int'l Conf. on Computer Vision Systems, LNCS, Vol. 6962, pp. 112-121, Springer, 2011.
URL http://dx.doi.org/10.1007/978-3-642-23968-7_12

SCENIOR (SCENe Interpretation by Ontology-based Rules) is an implemented system for realtime recognition of multi-object activities in real-world scenarios. SCENIOR has been developed for monitoring aircraft turnarounds at Blagnac Airport in Toulouse, for example aircraft arrival preparation, unloading, loading, refuelling and other service operations. SCENIOR expects time-marked information about individual object locations as input and delivers activity descriptions as output.

Activity recognition is based on declarative models specified in the standardised ontology language OWL and the Semantic Web Rule Language SWRL. The recognition system is compiled automatically from this model base. Thus models can be modified, or new models can be added, without reprogramming the recognition process.

Input data can be obtained by several means, for example by object-centered GPS location transmission, or by visual tracking using cameras. In the prototypical implementation at Blagnac Airport, objects have been tracked by 6 cameras firmly installed at the border of the apron. The tracked object locations were used to generate elementary events relating the objects to fixed zones, such as "Loader-Positioned-In-Right-FWD-Loader-Zone" or "Tanker-Leaves-Left-Tanker-Zone". These were transmitted as input to SCENIOR and interpreted in realtime as meaningful activities (or unrelated events).

Activity models are structured in a compositional hierarchy. This way, it is possible to recognise sub-activities contributing to a higher-level activity and generate predictions about the completion of a turnaround.

The system also comprises uncertainty management based on probability distributions for the temporal properties of all parts of a turnaround. This way, the system can cope with imperfect or partial data, and possible alternative interpretations can be given a preference rating.

Viewed as a support system for aircraft servicing, the following benefits can be expected:

1. The system gives a realtime account of completed service activities, thus providing progress control.
2. The system allows estimates about the completion of remaining activities, thus facilitating further scheduling.
3. The system also recognises unscheduled or unusual activities, and may thus provide security information.

Viewed as a technological framework, SCENIOR permits scene interpretation for diverse domains. When adapting SCENIOR to another domain, the main tasks are: (i) Developing sensor analysis up to primitive event recognition, and (ii) modelling higher-level concepts in an OWL ontology.

3.11 Neural Mechanisms for the Analysis of Articulated Motion Sequences

Heiko Neumann (Universität Ulm, DE)

Joint work of Neumann, Heiko; Layher, Georg; Giese, Martin A.
Main reference G. Layher, M.A. Giese, H. Neumann, "Learning representations for animated motion sequence and implied motion recognition," in Proc. of the 22nd Int'l Conf. on Artificial Neural Networks – Part I (ICANN'12), LNCS, Vol. 7552, pp. 288–295, Springer, 2012.

The detection and categorization of articulated, or biological, motion is a crucial task underlying action recognition. Neural representations of perceived animate objects are built in STS (superior temporal sulcus) which is a region of convergent input from intermediate level form and motion representations in primate cortex. STS cell sub-populations are selectively responsive to specific action sequences. It is still largely unknown how and to which extent form and motion information contribute to the generation of representations specific to biological motion and what kind of mechanisms are involved in the learning processes.

A model architecture is proposed for the unsupervised learning of task specific articulated motion sequence representations. The processing builds upon two mainly separated pathways akin of the dorsal and the ventral streams in the visual cortex of primates. Along the model dorsal pathway image motion is processed and patterns of motion are represented [1, 4]. Respectively, contour and form information is processed and represented along the model ventral stream [2, 5]. Distinctive global level motion and form category representations are learned in independent pathways. Unsupervised Hebbian learning is employed to build such categorical representations which serve at input stage to the feedforward convergent processing at the level of model STS. How does the model automatically select significant motion patterns as well as meaningful static snapshot categories (keyposes) from video inputs? Such keyposes correspond to articulated postures which are particularly characteristic for a given motion sequence and thus decrease the ambiguity of a temporal prediction. It is shown how sequence selective representations are learned in STS by fusing form and motion input from the segregated bottom-up driving input streams [3]. We also emphasize the role of feedback signals propagated backwards along descending processing channels to make predictions about future input as anticipations generated by sequence-selective neurons. Network simulations demonstrate the computational capacity of the proposed model by reproducing several experimental findings from neurosciences and recent behavioral data.

References

1 C. Beck, H. Neumann. *Interactions of motion and form in visual cortex – A neural model*. Journal of Physiology Paris 104: 61–70, 2010

2 T. Hansen T, H. Neumann. *Neural mechanisms for the robust representation of junctions*. Neural Computation 16: 1013–1037, 2004.

3 G. Layher, M.A. Giese, H. Neumann. *Learning representations for animated motion sequence and implied motion recognition*. Proc. Int'l Conf. on Artificial Neural Networks, Part I, ICANN 2012, A.E.P. Villa et al. (Eds.), LNCS 7552, Springer, pp. 288–295, 2012

4 F. Raudies, E. Mingolla, H. Neumann. *A neural model for transparent motion processing*. Neural Computation 23: 2868–2914, 2011

5 U. Weidenbacher, H. Neumann. *Extraction of surface-related features in a recurrent model of V1-V2 interactions*. PLoS ONE 4(6): e5909, 2009

3.12 Observing and Modeling the Embodiment of Attention

Lucas Paletta (Joanneum Research – Graz, AT)

Joint work of Paletta, Lucas; Santner, Katrin; Fritz, gerald; Mayer, Heinz
Main reference Visual recovery of saliency maps from human attention in 3D environments. Proc. IEEE International Conference on Robotics and Automation, ICRA 2013, Karlsruhe, Germany, May, 2013, in print.

Computational modeling of visual attention has recently been emerging as an important field of computer science and Artificial Intelligence. From human attention we know that many brain areas, including those processing motor signals, are involved in the computation of saliency as an indicator for which information a next action should consider. The concept of embodied attention understands attention processing as a meaningful system component within a perception-action cycle of autonomous systems where saliency computation should be operated according to the task at hand. Previous work (Paletta et al., 2005) developed a model for eye movements and belief aggregation for the task of object recognition where sequential attention strategies are adjusted in the frame of reinforcement learning. Furthermore, contextual rules (Perko et al., 2009) may prime the location of attention processing in the visual information. Current work targets at including physical actions such as body posture and position dynamics into the framework. In a first step, we extract ground truth data from human studies using eye tracking glasses and a tuned framework of SLAM (simultaneous localisation and mapping) that allows to map human gaze and integrated saliency measures directly onto the acquired three dimensional model of the environment, with high precision, with wearable interfaces that enable natural behaviours and without the use artificial markers (Paletta et al., 2013). Future work will use human ground truth to learn extended models of embodied attention from human behaviour.

3.13 Hybrid planning and plan recognition

Bernd Schattenberg (Universität Ulm, DE)

This presentation introduces plan recognition in the context of hybrid planning. The hybrid planning approach integrates the notion of action abstraction from hierarchical task network planning with the notion of means-end reasoning from partial-order causal-link planning into a common framework. We introduce the formal framework and show how search in the space of plan refinements generates solutions.

The talk also motivates central issues in recognizing plans from the point of view of hybrid planning for human users: plans provide context to the observed actions and hence any deviation of a user from a committed plan raises questions of whether the deviation compromises the overall causal structure wrt. the goals, whether the deviation can be interpreted as an ad-hoc (or: improvised) alternative task implementation, and the like.

3.14 Can spatial partitioning help with interpreting observed action?

Sabine Timpf (Universität Augsburg, DE)

In this talk I explore the notion of spatial partitioning for subdividing an activity such as navigating from a perron to another perron in a train station into several actions. Actions, essential actions and operations as subdivisions of an activity are introduced. The strategy for partitioning space is taken from earlier work on schematic geometry especially discussing the notion of a gateway as a place where one scene is left and another scene is entered. Wayfinding works by walking along a sequence of scenes through gateways. In order to implement this notion in an agent-based system, a change of perspective to an immersive, bottom-up one is necessary. The problem of implementing how a human agent would walk around an obstacle and the necessary parameters for this operation are discussed. The agent's movements can be interpreted as operations in the activity, subdividing the actions, but not exactly as a hierarchical subdivision.

3.15 Towards Learning Activities From Kinect Data

Thomas Wiemann (Universität Osnabrück, DE)

In this talk we present an approach to recognize activities from Kinect data based on clustering in polygonal reconstructions of the sensed scene. The idea behind this research is to identify changes in a reference scene. If an activity is going on new objects will appear or disappear and a person's movement will cause shadows. The sequences of added clusters and detected shadows are logged over time. Our aim is to use machine learning techniques to identify activities in the logged data.

3.16 Qualitative Spatial Reasoning for Interpreting Action

Diedrich Wolter (Universität Bremen, DE)

The claim of this work can be summarized as follows: qualitative representation and reasoning techniques provide good means to represent knowledge about actions and qualitative reasoning supports recognition and interpretation of observed actions.

Qualitative representations of space and time are acknowledged for their ability to capture cognitive concepts that underly spatial and temporal knowledge. With their according reasoning techniques, qualitative approaches provide a symbolic approach to representation and reasoning with spatial and temporal knowledge. Their primary feature is to bridge the gap between low-level data and conceptual knowledge. A drawback of qualitative approaches is that they are domain-dependent. Different applications require different representations and different reasoning algorithms. This can be a severe burden for application developers. To overcome this problem, we develop a versatile reasoning toolbox SparQ (http://www.sfbtr8.uni-bremen.de/project/r3/sparq/) that aims at making the various reasoning algorithms easily accessible to application developers.

More recently, qualitative approaches have been combined with expressive general logics. In case of combining a spatial representation with a logic, the term spatial logic has been coined. We use the term qualitative spatial logic to stress that a qualitative spatial representation is used. Our work revealed that the combination of a modal logics with a qualitative spatial representation is particularly attractive. Modal logics offer the expressivity to express temporal knowledge (linear temporal logic (LTL), for example) which allows purely spatial representation to be extended to represent spatio-temporal knowledge as required to represent actions. Good reasoning characteristics of the modal logic can be conveyed to the extended logic. Developing qualitative spatial logics has several interesting implications. Firstly, action representations based on qualitative spatial logics are comprehensible. Since the formulas are based on the qualitative concepts that relate to cognitive concepts, such formalism is well-suited for knowledge engineering. Secondly, suppose we are given a formula representing an action. Given observation data, one can approach action recognition as a model checking task. Sophisticated software tools are available that can easily cope with large data sets, making the overall approach efficient. Thirdly, this process can be generalized to reason about unknowns. For example, reasoning tools can be used to sensibly supplement missing information. Reasoning can be used to infer which pieces of information, if adjoined to the observations made, would allow an action to be identified. For example, by reasoning one can yield an interpretation stating that the observed actions suit the pattern of 'laying the dining table', given that an unidentified object X is plate. The ability to supplement missing pieces of information sensibly shows that qualitative spatial logics provide adequate means to tackle interpreting actions.

4 Working Groups

Most of Tuesday was dedicated to working groups, which were supposed to develop sets of criteria for benchmarks that can be used for research in formalisms, methods, techniques and systems related to interpreting observed action. The working groups took into consideration that there might be a need for different benchmarks for different sensor setups and that the benchmarks should allow for different scenarios (or classes of scenarios), should include humans, robots, and animals, should be suitable for academic exercises as well as real-world applications, and should allow for comparing different approaches against each other.

The three working groups looked at three classes of benchmark scenarios: a plan recognition scenario (like helping users of technical systems or recognising that someone is having breakfast), a state monitoring scenario (like a robot waiter monitoring coarse-grained states of its customers or an electronic caretaker looking after an elderly person) and a scenario for activity recognition in dynamic environments (like road traffic or crowd dynamics as in airports, open-air concerts, and soccer stadiums).

The working group focussing on plan recognition first looked at what plan recognition means and came to the conclusion that this is an active rather than a passive process. The group then looked at characteristics of plans, which include causal structure, execution of plans, multiple plans, hierarchical structure, and the number of agents involved. Potential tasks in this context are identifying and anticipating activities, rating the normality of an activity, annotating video data streams, and generalisiation and conceptualisation of activities.

The discussion in the working group on state monitoring focussed on getting a better understanding of what exactly a state is in the context of interpreting observed action. In this context, a state is predefined and specific, and often defined through a collection of cases. State monitoring might involve multiple indicators, and benchmarking might need to use artificial data.

The working group dealing with activity recognition in dynamic environments started with enumerating various relevant scenarios, which include road traffic and accident surveillance, autonomous vehicles and observation of the environment, soccer game observations, sea vessel observations, recognising abnormal behaviour, pedestrian/crowd behaviour recognition, collaborative activities (garbage collectors, rescue teams), and potentially gesture recognition. They emphasised that there is a difference between interpretation and recognition, where the latter (in contrast to the first) uses fixed sets of interpretations and learned concepts.

5 Plenary Discussion

The plenary discussion on the last day of the seminar focussed on two overarching questions that resulted from the working groups:

1. What are the structural dimensions of a benchmark scenario?
2. What should a benchmark website include?

One of the most important dimension of a benchmark scenario turned out to be the quality of the data. Are the data noisy and/or do they contain errors (e.g. produced by noisy channels)? Are data coming from different sources, and what are these sources (e.g. sensors, cameras, etc.)? Are the data complete or do they represent only partial information? Under which circumstances was the data obtained (e.g. night vs. day in the case of video data).

A second dimension revolves around issues related to the observed activities. Are the activities single activities or repetitive activities? Are we dealing with hierarchies of activities? Are multiple actors involved in the activity? Are actions performed in parallel? Is there diversity in the actions? Are there multiple scenes or just a single scene, and if the first is the case, is there diversity in the scenes?

Issues around time and space form another dimension of the benchmark scenarios. Are we dealing with a small or large scale space (in respect to the scene). Are there multiple timescales, and what is the length of a typical episode? Are the data continuous or do they represent snapshots?

Last but not least we need to consider the complexity of the data as well as its availability. Does the benchmark provide various levels of complexity? Are the data real-world data (indoor or outdoor) or artificially generated data? Are data available for different domains?

As far as a potential benchmark website is concerned, there is a need for tools that enable users to upload new data, reports on results achieved with the data and experiences gained with them, software used to achieve the results and experiences, sets of use cases, and collections of success stories. It is also desirable to have a toolbox for manipulating data sets so that they can be inspected easily, transformed into different formats, etc.

In respect to the data sets, they should be available at different levels of granularity and abstraction, and potentially in different representations. The data sets should be provided with evaluation guidelines, annotations, and clear explanations of their aims. Metadata and ground truths should be available at least for some of the data sets.

Other aspects of the website include:

- Wiki and forum for discussing technical issues and results
- Tutorials for the material available
- Links to relevant conferences
- Repository with papers using the benchmarks

Overall, the value of a benchmark website for interpreting observed action seems to be without question. However, it is yet to be determined who is in charge of the website and what the timeframe is for setting it up.

Participants

- Sven Albrecht
Universität Osnabrück, DE
- Mehul Bhatt
Universität Bremen, DE
- Susanne Biundo-Stephan
Universität Ulm, DE
- Martin V. Butz
Universität Tübingen, DE
- Amedeo Cesta
ISTC-CNR – Rome, IT
- Krishna Sandeep Reddy Dubba
University of Leeds, GB
- Tom Duckett
University of Lincoln, UA
- Frank Dylla
Universität Bremen, DE
- Simone Frintrop
Universität Bonn, DE
- Martin A. Giese
Universitätsklinikum Tübingen, DE
- Hans Werner Guesgen
Massey University, NZ
- Verena V. Hafner
HU Berlin, DE
- Joachim Hertzberg
Universität Osnabrück, DE
- Alexandra Kirsch
Universität Tübingen, DE
- Geert-Jan M. Kruijff
DFKI – Saarbrücken, DE
- Stephen R. Marsland
Massey University, NZ
- Bernd Neumann
Universität Hamburg, DE
- Heiko Neumann
Universität Ulm, DE
- Lucas Paletta
Joanneum Research – Graz, AT
- Bernd Schattenberg
Universität Ulm, DE
- Aryana Tavanai
University of Leeds, GB
- Sabine Timpf
Universität Augsburg, DE
- Thomas Wiemann
Universität Osnabrück, DE
- Diedrich Wolter
Universität Bremen, DE

Report from Dagstuhl Seminar 12492

Human Activity Recognition in Smart Environments

Edited by

James L. Crowley[1], Kai Kunze[2], Paul Lukowicz[3], and Albrecht Schmidt[4]

1 **INRIA Rhône-Alpes, FR,** james.crowley@imag.fr
2 **Osaka Prefecture University, JP** kai.kunze@gmail.com
3 **DFKI Kaiserslautern, DE,** Paul.Lukowicz@dfki.de
4 **Universität Stuttgart, DE,** schmidtke@cmu.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12492 "Human Activity Recognition in Smart Environments". We established the basis for a scientific community surrounding "activity recognition" by involving researchers from a broad range of related research fields. 30 academic and industry researchers from US, Europe and Asia participated from diverse fields including pervasive computing, over network analysis and computer vision to human computer interaction. The major results of this Seminar are the creation of a activity recognition repository to share information, code, publications and the start of an activity recognition book aimed to serve as a scientific introduction to the field. In the following, we go into more detail about the structure of the seminar, discuss the major outcomes and give an overview about discussions and talks given during the seminar.

Seminar 02.–07. December, 2012 – www.dagstuhl.de/12492
1998 ACM Subject Classification I.5 Pattern Recognition; H.5 Information Interfaces and Presentation
Keywords and phrases Activity Recognition, Machine Learning, Context Awareness, Pattern Analysis, Sensor Networks, Wearable computing, Computer Vision
Digital Object Identifier 10.4230/DagRep.2.12.17

1 Executive Summary

James Lawrence Crowley
Paul Lukowicz
Albrecht Schmidt

Today, commercial systems have become popular that utilize a broad range of sensors to facilitate gesture and motion-based interaction. Examples range from multi-touch surfaces, through tilt control common in mobile phone applications, and complex motion-based games controllers, e.g. Nintendo WII and Microsoft Kinect. While these systems are mainstream, the next basic research challenge is activity-driven, implicit interaction. Two key differences to existing systems are:

1. the interpretation of complex human activities, and
2. the extension of interaction from periods where a user consciously performs control gestures to permanent monitoring of user activity.

Conceptually, activity-driven interaction builds on the vision of context awareness developed since the 1990 [1, 2, 3]. Applications range from sports, through mobile gaming, information retrieval, personal healthcare to industrial work support [4, 5, 6]. For example, monitoring certain activities can support therapy in areas ranging from cardiovascular diseases to psychiatric disorders and cognitive disabilities. Activity based support (automatically showing correct manual pages, pointing out mistakes) can speed up industrial maintenance tasks by up to 30%.

Despite demonstrated potential, currently only very simple activity based applications such as physical activity monitoring have managed to go beyond early stage lab demonstrations. From the scientific point of view the question is how to map information available from unreliable, often simple sources onto complex human activities. The main challenges stem from the combination of three factors:

- In every day situations sensor choice, placement and configuration is often dictated by practicability, usability, and user acceptance constraints rather by the requirements of the recognition system. In addition, the system configuration may dynamically change [7, 8].
- The diversity of human behavior. Even the simplest activities can be performed in a multitude of ways differing not only between people, but also between individual execution instances of a single person. (e.g. using different arms, different hand positions, or even the hip to close a drawer)
- The complexity of human behavior. Relevant human actions are seldom atomic and independent. Instead, a complex hierarchy of actions that may be executed in parallel, overlap and interleave is to be considered by the recognition system.

Beyond the technological challenges involved in the recognition system, there are additional unsolved problems including application design, usability, user acceptance, and business models for commercialization.

The field lacks also definitions for many commonly used terms including "action," "sensor," "evidence," and even "activity" itself, leading to ambiguity in scientific discourse. The conceptual grounding provided by Nardi and Kaptelinin's definition of Activity Theory are perfectly understandable to a human [9]. Yet, they are not easily codified into machine programmable constructs. The theory recognizes that elements of "activity" such as goal and motive are socially constructed, depending on the perspectives of the actors in the system. Despite the complexities of "activity" at the human cognitive level, researchers demonstrated that some notions of activity can be utilized in computer systems, but meanings of terms differ among the various research groups. Currently, many different communities are involved in research related to activity recognition, including the core ubicomp community, human computer interaction, computer vision, cognitive science and artificial intelligence.

Privacy concerns are a critical barrier to adoption of activity-based technologies [10]. These concerns range from risk of criminal activities(e.g., stalking and identity theft), to social issues of managing personal relationships. Technological approaches to addressing the concerns must also be based on deep understanding of the psychological, sociological and political constraints under which people will operate activity-based systems.

The top level objective of the workshop was to define and establish the scientific community and associated research questions/methodologies related to the broad area of activity recognition. The major tangible outcomes are the start of the creation of an activity recognition repository accessible under http://activity-recognition.github.com and the plan of writing a standard book about activity recognition.

References

1 A Schmidt, K Aidoo, and A Takaluoma. Advanced interaction in context. *Handheld and Mobile Computing*, 1999.

2 T. Starner, B. Schiele, and A. Pentland. Visual contextual awareness in wearable computing. In *Proc. of the Second Int'l Symp. on Wearable Computing. Pittsburgh, October*, 1998.

3 Albrecht Schmidt. Implicit human computer interaction through context. *Personal and Ubiquitous Computing*, 4(2):191–199, 2000.

4 Jesse Hoey, Thomas Plötz, Dan Jackson, Andrew Monk, Cuong Pham, and Patrick Olivier. Rapid specification and automated generation of prompting systems to assist people with dementia. *Pervasive and Mobile Computing*, 7(3):299–318, 2011.

5 J Lester, T Choudhury, and G Borriello. A practical approach to recognizing physical activities. *Proc. of Pervasive*, Jan 2006.

6 Kai Kunze, Florian Wagner, Ersun Kartal, Ernesto Morales Kluge, and Paul Lukowicz. Does context matter?-a quantitative evaluation in a real world maintenance scenario. *Pervasive Computing*, pp. 372–389, 2009.

7 Kai Kunze and Paul Lukowicz. Dealing with sensor displacement in motion-based onbody activity recognition systems. In *Proc. of the 10th Int'l Conf. on Ubiquitous Computing*, pages 20–29. ACM, 2008.

8 K. Kunze and P. Lukowicz. Symbolic Object Localization Through Active Sampling of Acceleration and Sound Signatures. *Lecture Notes in Computer Science*, 4717:163, 2007.

9 V Kaptelinin and BA Nardi. Activity theory in a nutshell. *Acting with Technology: Activity Theory and Interaction Design*, pp. 29–72, 2006.

10 Stephan Sigg. Context-based security: State of the art, open research topics and a case study. In *Proc. of the 5th ACM Int'l Workshop on Context-Awareness for Self-Managing Systems*, pp. 17–23. ACM, 2011.

2 Table of Contents

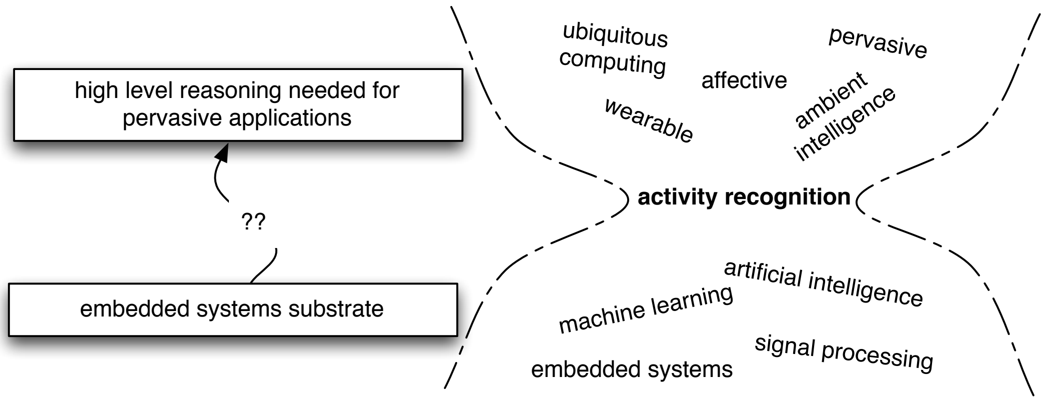

Figure 1 Activity recognition – the link between our sensor rich environments and high level reasoning.

3 Outcomes

There was broad agreement that activity recognition is the link between our sensor rich environments and high level reasoning about a user's actions (see Fig. 1). However, the definitions for "activity", "action", "sensor" and "activity recognition" depend highly on the research direction. Based on the presentations and discussions during the seminar we can abstract 5 major research directions each with different relations towards activity recognition. They are listed subsequently:

- Sensing – researchers explore novel sensing modalities to detect specific types of activities [2]. Here the focus is on novel sensors for activity recognition, building new hardware prototypes. The activities recognized are usually relatively simple (no complex time series etc.) to emphasize the usefulness of the sensors. The research itself is very practical and evidence based, with an emphasis on elaborate experimental setups. Therefore it's often difficult to reproduce the results, since building identical sensors involve significant effort and many papers do not contain sufficient details on the hardware.
- Systems – Building activity recognition systems, usually focused on a specific application scenario (e.g. health care, maintenance etc.). This field is very similar to sensing, as it emphasizes the usefulness of a system to a particular application field, also empirical and difficult to reproduce as the systems use often prototype hardware and are fine tuned for a specific environment.
- Machine Learning – researchers work on novel machine learning algorithms and see the field of activity recognition as a suitable application field for the novel algorithms [1]. Although the inference is based on sensors, the hardware itself is secondary. The emphasis of the research is on the type of algorithms and the improvements they bring over other inference approaches. The research focuses is often more theoretical. As long as the datasets and algorithms the research is based on are available, it's easier to reproduce.
- Middleware – researchers combine different isolated sensing, systems and machine learning solutions to build integrated systems [3]. This field is more on the theoretical side as it tries to build standards to use for exchanging activity recognition information.
- High Level Modeling – to apply higher level reasoning towards lower level activity inference [4]. This research direction deals mainly with applying artificial intelligence reasoning techniques to activity recognition problems. It also looks to cognitive science

for models of how complex activities can be composed of simple actions. Work is often theory focused and can be reproduced more easily, as long as algorithms and datasets are open.

Depending on the research direction, the general understanding of the term "activity" might differ and the interests are very diverse. However, there are two major agreements:

1. Every researcher utilizes a recognition chain for their work (ranging from sensor inputs towards the higher level activity recognition), although the actual implementations and components differ widely.
2. Some sort of hierarchical structure exists for activities in nearly all scenarios (atomic motions/actions towards compositions).

To consolidate the different view points towards activity recognition and facilitate a more closely integrated research community, we decided to work on two distinct projects, which are the major outcomes of the seminar:

1. The creation of a activity recognition repository to enable a single access point for anybody who wants to learn about activity recognition and to provide an ongoing exchange about research methods, code, datasets and other types of information.
2. Planning of an "Activity Recognition" book to introduce interested Master-level students and researcher from other fields to activity recognition.

3.1 Activity Recognition Book

To build the foundation for a new research community around activity recognition, we started with a book project. The goal of the book is to provide a broad introduction into the field of activity recognition. The tentative audience is advanced Bachelor or Master students with some background in Computer Science. The topics brainstormed that should be included are shown in Fig. 2.

During the seminar we succeeded to come up with a general outline (Fig. 3), starting with a introduction that describes the vision of the field and gives the reader already some hands-on applications (to illustrate the usefulness and power of activity recognition). The foundations chapter gives general definitions and an overview about the recognition architectures and the focuses of the different research directions discussed beforehand. The Tools and Systems Chapter is a more practical part, providing the reader with essential background in the recognition chains used. Finally, information sources touches the sensing part of activity recognition, discussing a broad variety of sensors. The problem types mention the issues one can tackle with activity recognition techniques. Finally, the Applications Chapter explains how the new methods learned can be applied using again detailed examples. As activity recognition lives from detailed, extensive experimental setups and empirical evaluations a Research Methodology Chapter is also crucial (as some readers might have a more theoretical Computer Science background).

3.2 Activity Recognition Repository

During the workshop we have setup a public repository about information, code and datasets related to activity recognition.

So far, information about activity recognition was distributed among the different research groups and most definitions and explanations focused on a particular type of activity

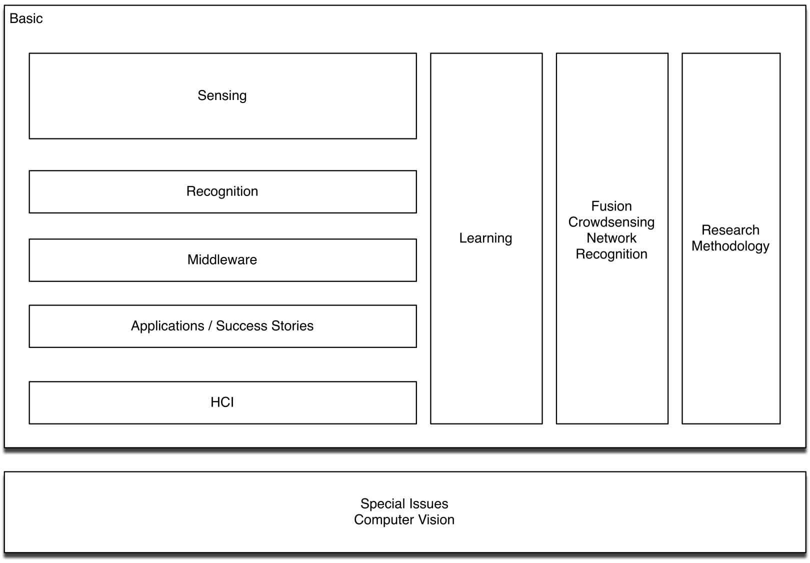

Figure 2 Topic brainstorming results for the activity recognition book.

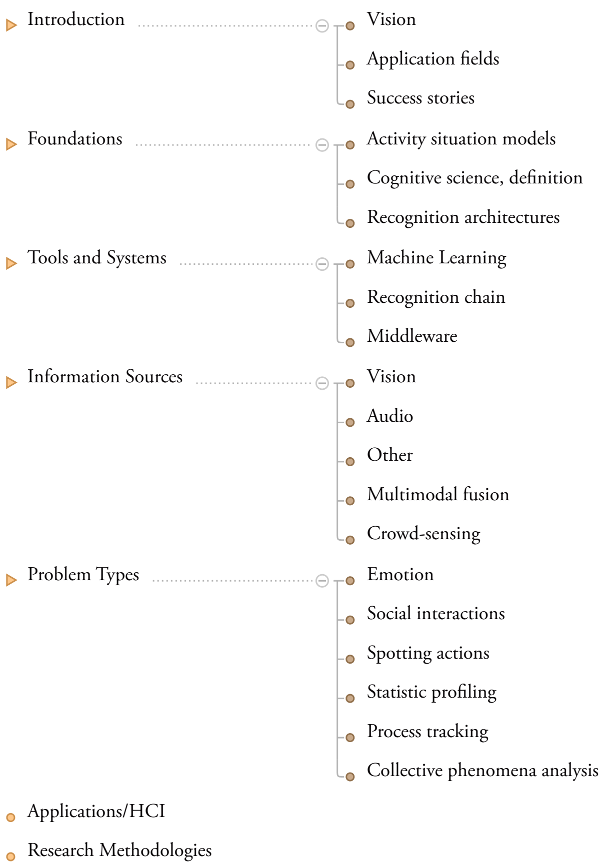

Figure 3 Tentative outline of the activity recognition book.

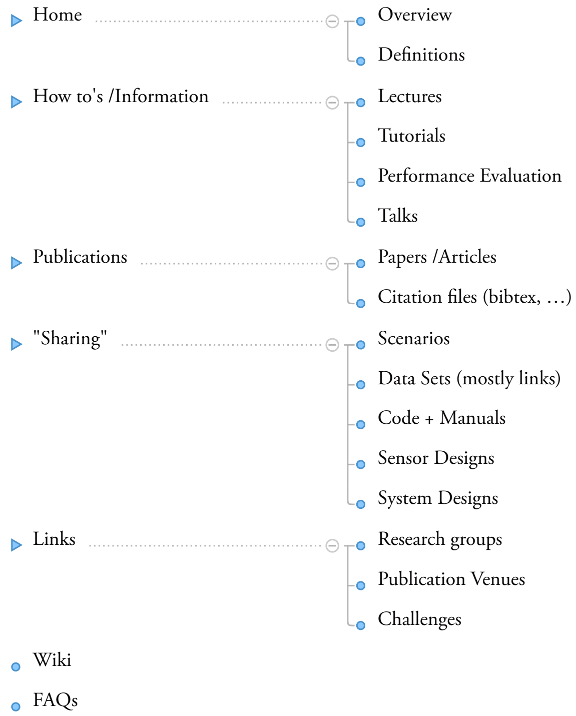

Figure 4 Tentative outline of the activity repository.

recognition. To overcome this problem we decided to create a central repository where we can gather all types of information about the topic. A first version of the repository website, as seen in Fig. 5 is available under: http://activity-recognition.github.com The tentative outline of the repository as concluded during the seminar is shown in Fig. 4. The home page starts with an introduction to the field and the definitions. It contains also a information page with lectures, tutorials and talks centering around activity recognition. The publications page holds all relevant papers and articles. In the sharing section, links to datasets, codes and schematics are included. Followed by information about research groups and publication venues.

The repository is hosted at github, see Fig. 6 (all the source-code is available under a free license), this is done so it's independent from any individual research institution contributing and give people easy access to all information contained on the repository.

Some only partially resolved issues center around how to make the datasets accessible and how to best use social media. The sections should not be to broad at the beginning to gather enough meaningful data and not have a sparsely-populated site, that might not be helpful.

References

1 M. Andriluka, S. Roth, and B. Schiele. People-tracking-by-detection and people-detection-by-tracking. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conf. on*, pp. 1 –8, june 2008.

2 Jingyuan Cheng, Oliver Amft, and Paul Lukowicz. Active capacitive sensing: Exploring a new wearable sensing modality for activity recognition. *Pervasive Computing*, pp. 319–336, 2010.

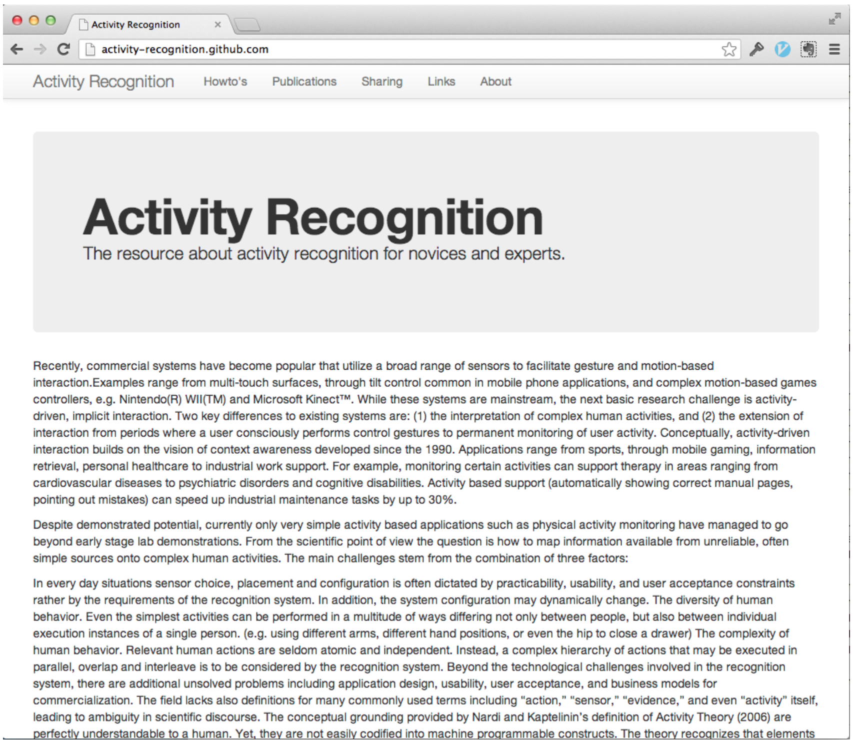

Figure 5 Initial version of the activity repository is already online under http://activity-recognition.github.com.

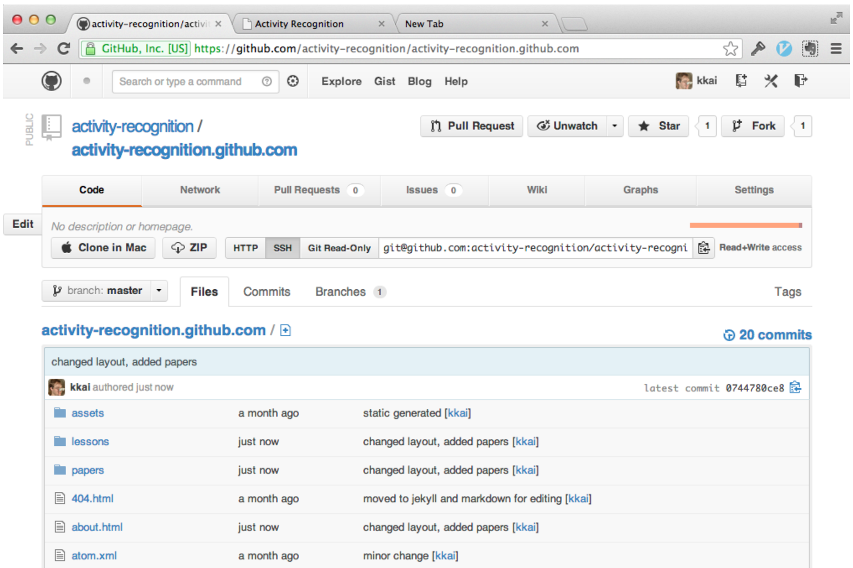

Figure 6 Github interface to manage the repository.

3 Alois Ferscha. Contextware: Bridging physical and virtual worlds. *Reliable Software Technologies–Ada-Europe 2002*, pp. 569–589, 2002.
4 Mehul Bhatt, Andrew Flahive, Carlo Wouters, Wenny Rahayu, David Taniar, and Tharam Dillon. A distributed approach to sub-ontology extraction. In *Advanced Information Networking and Applications, 2004. AINA 2004. 18th Int'l Conf. on*, vol. 1, pp. 636–641. IEEE, 2004.

4 Open Problems

In the following we discuss some open problems which could not be completely resolved and need to be tackled in turn by the community and the two projects – repository and book.

4.1 What is an activity?

Although we were able to exchange the definitions used by the different research directions related activity recognition, we did not succeed in defining all terms dealing with this novel research field. In particular, the question of an adequate hierarchical representation and of a taxonomy of relevant activities remained open.

4.2 Reproducibility

Given the nature of activity recognition, some research is hard to reproduce, as it centers around specific hardware (mostly prototype systems) etc. We need a community effort to create a research environment that enables better reproducibility (sharing of schematics, usage of rapid prototyping AR platforms etc.). The same holds for datasets and algorithms. Given there are invite types of "activities" and a wide variety of sensor modalities, we really lack datasets. We really need an openCV for activity recognition to move the filed forward. One of the bigger questions here, was how can we give incentives as a community to ensure reproducibility etc.

5 Keynote Talks

The seminar began with a series of keynote talks as a starting point and to provide common ground for later discussions.

5.1 Wearable and Pervasive Sensor-based Activity Recognition

Paul Lukowicz (DFKI Kaiserslautern, DE)

Activity Recognition in Smart Environments can be defined as the ability to extract high level information about human actions and complex real world situations from heterogeneous ensembles of often very simple, unreliable sensors. The talk started by discussing the foundations of such systems as well as issues related to sensing and empirical evaluation in complex real life applications. Such applications range from industrial process support through personal health care, sports and entertainment to the optimization of energy usage at home. Examples were given from divers EU and national project in which I currently participate.

The core of current work on context aware systems involves small groups of devices interacting with a single user or a small groups of users in system configurations specifically designed for often narrowly defined applications. On the other hand, smart phones, home automation devices, robots and other intelligent systems are becoming ubiquitous and are increasingly equipped with flexible networking capabilities. Thus, looking at the future of context aware systems we need to consider large collectives of such devices dynamically forming, cooperating, and interacting with large user populations over a broad range of spatial and temporal scales.

5.2 Analysis of human interactions through complex networks theory

Katharina A. Zweig (TU Kaiserslautern, DE)

Network analysis provides the uniform framework to analyze big data sets from disciplines as far apart as medicine, sociology, human complex problem solving, cancer biology, and archaeology. Once the problem is represented as a network, various centrality measures, clustering algorithms, and statistical models can be applied to find the most central nodes, the densest subgraphs, or the strongest motifs [1]. Most of us have heard sentences like this for so long, that we do not even question them anymore. In this talk has shown that while all methods are in principle applicable, they come with their own modeling assumptions that might yield their results useless for a specific network. To prove the talk looked at airports with a low degree and a high betweenness centrality are not necessarily anomalous in an economic sense. It also solved the riddle why market basket analysis assumes that Pretty Woman is a good recommendation for somebody who loves Star Wars V. Understanding when to use which method is what I call "network analysis literacy". To make us all more literate, I finally advocate for a more transparent description of network analytic methods and their implicit modeling assumptions to enable an informed choice of the best measure for a given network analytic question.

References

1 Sudarshan Iyengar, Katharina A Zweig, Abhiram Natarajan, and CE Veni Madhavan. A network analysis approach to understand human-wayfinding problem. CogSci, 2011.

5.3 Socio-inspired ICT

Alois Ferscha (Universität Linz, AT)

A yet underexplored impact of modern ICT (like e.g. pervasive and ubiquitous computing) relates to services exploiting the "social context" of individuals towards the provision of quality-of-life technologies that aim at the wellbeing of individuals and the welfare of societies". This talk therefore is concerned with the intersection of social behavior and modern ICT, creating or recreating social conventions and social contexts through the use of pervasive, omnipresent and participative technologies. An explosive growth of social computing applications such as blogs, email, instant messaging, social networking (Facebook, MySpace, Twitter, LinkedIn etc.), wikis and social bookmarking is observed, profoundly impacting social behavior and life style of human beings while at the same time pushing the boundaries of ICT simultaneously. We aim to investigate interface technologies for one important phenomenon in humans, namely that of "social awareness". We aim at providing human environment interfaces, which allow individuals and groups to sense, explore and understand their social context.

5.4 Visual Perception of Activities

James L. Crowley (INRIA Rhóne-Alpes, FR)

This talk focuses on vision based activity recognition. Computer vision is a mature research field that had to face similar problems as activity recognition does today (with the limitation of just using one sensor). The talk introduces different problems in computer vision related to activities, from tracking over situation models.

5.5 Recognizing Reading Activities

Koichi Kise (Osaka Prefecture University, JP)

Today more and more people monitor their daily physical activities, from jogging, over step counting to sleeping to improve the quality of their physical life. We want to achieve the same for reading activities. As reading is fundamental to knowledge acquisition, we would be able to improve the quality of our knowledge life. We use reading, defined as the process of assigning meaning to characters, words and sentences, as a primary information source, recognizing and monitoring reading activities in unconstrained, natural settings is still largely unexplored. We propose reading activity recognition to analyze the whole process of human reading activity and knowledge acquisition [1]. So far, the activity recognition community focused mainly on detecting an activity, yet especially for reading activities it is important to look at the quality of the activity (how instead of what).

References

1 Tomohiro Nakai, Koichi Kise, and Masakazu Iwamura. Use of affine invariants in locally likely arrangement hashing for camera-based document image retrieval. *Document Analysis Systems VII*, pp. 541—552, 2006.

6 Overview of Talks

In addition to the keynote talks, a number of additional presentations were given to stimulate discussions and broaden the understanding about interesting topics at hand.

6.1 Bringing Social Computing to Smart Environments: Synergies and Challenges

Elisabeth André (Universität Augsburg, DE)

Societal challenges create a high demand for technology able to emulate human-style interaction modes. Currently, most human-machine interfaces focus on input that is explicitly issued by human users. However, often it is the myriad of unconsciously conveyed social and psychological signals that will determine whether an interaction with a machine is successful or not. The talk demonstrated how progress made in the areas of social computing and ambient environments can contribute to a deeper symbiosis in human-machine interaction by collecting subtle behavioral cues under naturalistic conditions and linking them to higher-level intentional states. However, on the way to this goal, a number of challenges need to be solved: Users show a great deal of individuality in their behaviors, and there is no clear mapping between behavioral cues and intentional states. This is in particular true for real-life settings where users are exposed to a more diverse set of stimuli than under laboratory conditions. Furthermore, it isn't obvious how to acquire ground truth data against which to evaluate the performance of system components that map unconsciously conveyed behavioral cues onto intentional states. Finally, we need to cope with limited resources when recording social and psychological cues in a mobile context and responding to them in real-time. Apart from technological challenges, psychological, societal and privacy issues need to be taken into account. Based on an analysis of recent activities in the areas of social computing and ambient environments, the talk outlined a road map for future research [1].

References

1 Johannes Wagner, Jonghwa Kim, and Elisabeth André. From physiological signals to emotions: Implementing and comparing selected methods for feature extraction and classification. In *Multimedia and Expo, 2005. ICME 2005. IEEE Int'l Conf. on*, pp. 940–943. IEEE, 2005.

6.2 Eye-based activity and context recognition

Andreas Bulling (University of Cambridge, GB)

The talk discusses the feasibility and potential of eye-based activity recognition, including an assessment of eye tracking techniques, specifically of optical methods that track features on the eye, and of electrooculography (EOG) measuring eye movement with skin electrodes placed near the eyes. Insights into what eye movement reveals about everyday activity a methods for extraction of features from eye movement patterns that may be useful for activity and context recognition are presented [1].

References

1 Andreas Bulling, Jamie A. Ward, Hans Gellersen, and Gerhard Tröster. Robust recognition of reading activity in transit using wearable electrooculography. In *Proc. of Pervasive '08*, pp. 19–37, 2008.

6.3 Where Activity Recognition supports User Interface Plasticity

Joelle Coutaz (Université de Grenoble, FR)

This short talk presents my background and my interest for Activity Recognition. In short, I do not do "Activity Recognition", but I need solutions from Activity Recognition to build User Interfaces (UI) that are truly plastic.

UI plasticity is the capacity of a UI to adapt to the context of use while preserving human values. In this "context", context of use denotes an information space that describes the humans using the system, the physical and social environment where the interaction takes place, and the physical platform (interaction resources) available in this space.

6.4 Interacting with Machine Learning

James Fogarty (University of Washington – Seattle, US)

This talk provided an overview of the sensing On Shared Datasets and Validation. An concrete example given related to what sensors might predict interruptibility and verifying that implemented sensors are effective, looking more closely at programmer task engagement, demonstrating signal in audio-based water sensing, demonstrating signal in pressure-based water sensing and relating WiFi activity to coffee shop seat availability.

6.5 Social Situations and Co-Activity Detection

Georg Groh, Alexander Lehmann, and Daniel Bader (TU München, DE)

We introduce our current projects and research directions, our interest in activity recognition, our point of view towards datasets and evaluation and the main research questions in are that are relevent in connection with our research.

The main conceptual axis is social contexts on all temporal scales: From Social Situation models as short term social context to long term social context such as topics of social relations, employing machine learning techniques to detect these contexts. An interesting aspect are agent models for combining social contexts on various levels of abstraction, ranging from exchanging raw data to collaborative symbolic reasoning on these contexts. Furthermore, we also investigate various applications of social context, such as social recommender systems, privacy for social networking, social information retrieval etc. One specific project is the detection of co-activities, which neglects the actual activity classification and concentrates on identifying common activities among people [1].

References

1 Wolfgang Woerndl and Georg Groh. Utilizing physical and social context to improve recommender systems. In *Web Intelligence and Intelligent Agent Technology Workshops, 2007 IEEE/WIC/ACM Int'l Conf. on*, pp. 123–128. IEEE, 2007.

6.6 A generalized computational model for complex systems

Jochen Kerdels (FernUniversität in Hagen, DE)

The use of computer models and simulation is a widely adopted approach to study complex systems. To this end a diverse set of computational models like Cellular Automata, Artificial Neural Networks, or Agent-based simulation is being used. As a common denominator virtually all of these approaches favor different variations of complex systems and are tailored to support the description of systems that fit the corresponding variation well. Although this form of specialization has its benefits like ease of modeling with respect to the particular subset of complex systems, the drawbacks of this specialization are a lack of comparability between structurally different systems and a diminished expressiveness with respect to systems that do not fit any particular subset of complex systems favored by existing, specialized models. In this talk a generalized computational model for complex systems is proposed which allows for the description of most types of systems with a single model. Furthermore, the proposed model provides a high degree of encapsulation and reduces the amount of shared knowledge needed among the constituents of the system. The talk closes with a set of example applications of the proposed model to further illustrate the involved concepts and to provide an intuition on how this model may be used [1].

References

1 Jochen Kerdels and Gabriele Peters. A generalized computational model for modeling and simulation of complex systems. 2012.

6.7 Discovery of Everyday Languages

Daniel Kohlsdorf (Georgia Institute of Technology, US)

This talk introduces the concept of simultaneously learning patterns and grammar for activity recognition tasks guided by evaluation [1]. We have to find the atomic patterns (Discovery), define their interactions (Grammar) and finally evaluate to see how well our model fits to the real world.

References

1 David Minnen, Irfan Essa, and Thad Starner. Expectation grammars: Leveraging high-level expectations for activity recognition. In *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE CS Conf. on*, vol. 2, pp. II–626. IEEE, 2003.

6.8 Real Life Activity Recognition

Kai Kunze (Osaka Prefecture University, JP)

This talk investigates how placement variations of electronic devices influence the possibility of using sensors integrated in those devices for activity recognition [1]. The vast majority of activity recognition research assumes well defined, fixed sensor locations. Although this might be acceptable for some application domains (e.g. in an industrial setting), users, in general, will have a hard time coping with these limitations. If one needs to remember to carry dedicated sensors and to adjust their orientation from time to time, the activity recognition system is more distracting than helpful. How can we deal with device location and orientation changes to make activity sensing mainstream?

References

1 Kai Kunze and Paul Lukowicz. Dealing with sensor displacement in motion-based onbody activity recognition systems. In *Proc. of the 10th Int'l Conf. on Ubiquitous Computing*, pp. 20–29, ACM, 2008.

6.9 Computational Behaviour Analysis

Thomas Plötz (Newcastle University, GB)

This talk centers around computational (that is statistical) models that describe and will help in assessing human behaviour. The basis for this is the analysis of behavioural data that is captured in a pretty opportunistic way utilising a variety of sensing modalities, most notably pervasive/ubiquitous sensors (e.g., accelerometers, RFID, environmental sensors), cameras, or microphones. The modelling itself is agnostic in terms of the actual choice of sensing modalities as long as the relevant information for behaviour analysis is captured. Behaviour data are sequential by definition. Consequently, related modelling techniques are focused on sequential models (for example of Markovian type). I am especially interested in quantitative assessments of human behaviour, which is of value for, for example, skill assessment. My day-to-day work can probably best be summarised as applied machine learning for activity recognition.

6.10 Multimodal Activity Recognition using Wearable Sensors and Computer Vision

Bernt Schiele (MPI für Informatik – Saarbrücken, DE)

Understanding sensor information is a fundamental problem in computer science. Scientific challenges cover the entire pipeline from single-sensor processing, over spatial and temporal fusion of multiple and divergent sensor modalities to the complete description of large-scale multimodal sensor streams. At the same time we observe a tremendous increase in both the quantity as well as the diversity of sensor information due to the increasing number of sensors (such as cameras, GPS, or inertial sensors) embedded in a wide variety of digital devices and environments as well as due to the increasing storage of multimodal sensor data (such as surveillance data, personal storage of digital information, multimedia databases, or simply the Internet). While storing and indexing large amounts of sensor data has made tremendous progress, understanding of this multimodal sensor data still lacks far behind. Therefore the long-term goal of our research is to make progress on how to process, structure, access and truly understand multi-sensory data both for online use as well as for large-scale databases.

6.11 Sensitivity for environment-enduced RF-channel fluctuation

Stepahn Sigg(NII – Tokyo, JP)

We consider the detection of activities from non-cooperating individuals with features obtained on the Radio Frequency channel [1]. Since environmental changes impact the transmission channel between devices, the detection of this alteration can be used to classify environmental situations. We identify relevant features to detect activities of non-actively transmitting subjects. In particular, we distinguish with high accuracy an empty environment or a walking, lying, crawling or standing person, in case-studies of an active, device-free activity recognition system with software defined radios. We distinguish between two cases in which the transmitter is either under the control of the system or ambient. For activity detection the application of one-stage and two-stage classifiers is considered. Apart from the discrimination of the above activities, we can show that a detected activity can also be localised simultaneously within an area of less than 1 meter radius.

References

1 Markus Reschke, Sebastian Schwarzl, Johannes Starosta, Stephan Sigg, and Michael Beigl. Context awareness through the rf-channel. *ARCS 2011*, 2011.

6.12 Short Introduction Talks

Additionally, all participants gave short introduction talks about their background split into the sessions depicted in Fig. 7.

Topic	Participants
Interaction I	Elisabeth André (Universität Augsburg, DE)
	Nora Broy (BMW Group Forschung und Technik GmbH - München, DE)
	Elizabeth F. Churchill (ACM SIGCHI, US)
	Joelle Coutaz (Université de Grenoble, FR)
Sensing	Andreas Bulling (University of Cambridge, GB)
	Jingyuan Cheng (Universität Passau, DE)
	Gerald Pirkl (DFKI - Kaiserslautern, DE)
	Kristof Van Laerhoven (TU Darmstadt, DE)
Interaction II	James Fogarty (University of Washington - Seattle, US)
	Niels Henze (Universität Stuttgart, DE)
	Jochen Kerdels (FernUniversität in Hagen, DE)
	Alireza Sahami (Universität Stuttgart, DE)
Real Life AR	Gernot Bahle (DFKI - Kaiserslautern, DE)
	Ulf Blanke (ETH Zürich, CH)
	Daniel Kohlsdorf (Georgia Institute of Technology, US)
	Kai Kunze (Osaka Prefecture University, JP)
Modelling	Mehul Bhatt (Universität Bremen, DE)
	Diana Nowacka (Newcastle University, GB)
	Thomas Plötz (Newcastle University, GB)
	Bernhard Sick (Universität Kassel, DE)
Mobile + Social	Oliver Brdiczka (PARC - Palo Alto, US)
	Thomas Phan (Samsung Research, US)
	Georg Groh (TU München, DE)
	Alexander Lehmann (TU München, DE)
Systems I	Hedda R. Schmidtke (Carnegie Mellon University - Moffet Field, US)
	Stephan Sigg (NII - Tokyo, JP)
Systems II	Michael Beigl (KIT - Karlsruhe Institute of Technology, DE)
	Dawud Gordon (KIT - Karlsruhe Institute of Technology, DE)
	Bernt Schiele (MPI für Informatik - Saarbrücken, DE)
	Hideyuki Tokuda (Keio University, JP)

Figure 7 Introduction talks of the participants divided into topic groups.

Participants

- Elisabeth André
Universität Augsburg, DE
- Gernot Bahle
DFKI – Kaiserslautern, DE
- Michael Beigl
KIT – Karlsruhe Institute of Technology, DE
- Mehul Bhatt
Universität Bremen, DE
- Ulf Blanke
ETH Zürich, CH
- Oliver Brdiczka
PARC – Palo Alto, US
- Nora Broy
BMW Group Forschung und Technik GmbH – München, DE
- Andreas Bulling
University of Cambridge, GB
- Jingyuan Cheng
DFKI – Kaiserslautern, DE
- Joelle Coutaz
Universite de Grenoble, FR
- James L. Crowley
INRIA, FR
- Alois Ferscha
Universität Linz, AT
- James Fogarty
University of Washington – Seattle, US
- Dawud Gordon
KIT – Karlsruhe Institute of Technology, DE
- Georg Groh
TU München, DE
- Niels Henze
Universität Stuttgart, DE
- Jochen Kerdels
FernUniversität in Hagen, DE
- Koichi Kise
Osaka Prefecture University, JP
- Daniel Kohlsdorf
Georgia Inst. of Technology, US
- Kai Kunze
Osaka Prefecture University, JP
- Alexander Lehmann
TU München, DE
- Paul Lukowicz
DFKI – Kaiserslautern, DE
- Diana Nowacka
Newcastle University, GB
- Thomas Phan
Samsung Research, US
- Gerald Pirkl
DFKI – Kaiserslautern, DE
- Thomas Plötz
Newcastle University, GB
- Alireza Sahami
Universität Stuttgart, DE
- Bernt Schiele
MPI für Informatik – Saarbrücken, DE
- Albrecht Schmidt
Universität Stuttgart, DE
- Hedda R. Schmidtke
Carnegie Mellon University – Moffet Field, US
- Bernhard Sick
Universit|"at Kassel, DE
- Stephan Sigg
NII – Tokyo, JP
- Hideyuki Tokuda
Keio University, JP
- Kristof Van Laerhoven
TU Darmstadt, DE
- Katharina A. Zweig
TU Kaiserslautern, DE

Report from Dagstuhl Seminar 12501

Organizational Processes for Supporting Sustainable Security

Edited by

Lizzie Coles-Kemp[1], Carrie Gates[2], Dieter Gollmann[3], Sean Peisert[4], and Christian Probst[5]

1 **Royal Holloway University – London, GB**
2 **CA Labs – New York, US,** carrie.gates@ca.com
3 **TU Hamburg-Harburg, DE,** diego@tu-harburg.de
4 **University of California – Davis, US** peisert@cs.ucdavis.edu
5 **Technical University of Denmark, DK,** probst@imm.dtu.dk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12501 "Organizational Processes for Supporting Sustainable Security" which ran from December 9 to 12, 2012 and was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. We also ran a number of collaborative sessions designed to promote the development of design principles for sustainably secure organizational processes. The first section describes the seminar topics and goals in general. The following section contains abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper.

Seminar 9.–12. December, 2012 – www.dagstuhl.de/12501
1998 ACM Subject Classification K.6.5 Management of Computing and Information Systems: Security and Protection, D.4.6 Operating Systems: Security and Protection
Keywords and phrases Insider threat, Organizational Process, Resilience, Security Policy
Digital Object Identifier 10.4230/DagRep.2.12.37

1 Executive Summary

Lizzie Coles-Kemp
Carrie Gates
Dieter Gollmann
Sean Peisert
Christian Probst

The Dagstuhl seminar "Designing for process resilience to insider threats" was held on December 10–12th December, 2012 (Seminar #12501) to advance our understanding of ways of reducing insider threats through the design of resilient organizational processes.

The 2012 seminar built on the results of its predecessor from 2010 (Insider Threats: Strategies for Prevention, Mitigation, and Response, #10341, Seminar Homepage, Seminar Report). In this seminar we developed a shared, inter-disciplinary definition of the insider and a good formulation for a taxonomy or framework that characterizes insider threats. The seminar also began to explore how organizational considerations might better be incorporated into addressing insider threats.

The purpose of the 2012 seminar was to build on the understanding of the classification of the insider threat as a type of informed threat and the design requirements for tools and policies to respond to this category of threat that we had gained from the 2008 and 2010 Dagstuhl seminars on insider threats (Countering Insider Threats, #08302, and Insider Threats: Strategies for Prevention, Mitigation, and Response, #10341). Our goal was to explore what makes organizational processes resilient to insider threats. The exploration of organizational processes required us to consider the fluid set of informed actors against organizations whose processes and boundaries can be dynamic. It also required us to conceptualise threats and vulnerabilities as "emergent". The conclusions from the previous seminars had resulted in the insight that resilient organizational processes are more resilient with respect to insider threats and more capable of limiting the damage from insider attacks. We also had the insight that resiliency appears to stem from usable, effective, and efficient security having been built into the organizational processes.

The seminar participants contained a carefully balanced mix of social and computer scientists and practitioners in order to explore the technological, organizational and social dimensions of the organizational process and its implementation. In order to productively combine the skills of the different disciplines and perspectives represented, the seminar started with a series of provocations. Debi Ashenden presented a provocation about the competing and sometimes conflicting uses of gamefication in the UK military setting. Kai-Uwe Loser presented a grounded example of personal data management practices and the conflicting perceptions of policy compliance that emerged within the example. Trish Williams presented a provocation about the value of big data in the case of electronic health data.

These design principles reflect a start point for future work on the design of organizational processes that are sustainably secure. Seminar organizers intend to produce a book that extends and explores these principles.

2 Table of Contents

3 Overview of Talks

3.1 Gemini – New Approach to Data Leakage Prevention

Julie Boxwell Ard (Glenn Dale, US)

Joint work of Ard, Julie Boxwell; Gates, Carrie; Bishop, Matt

In this talk we present a new approach to potentially detecting information leakage or malicious activity, whether it be intentional or unintentional activity. Traditionally, security research has approached insider threat detection from the perspective of modeling human behavior, often at the individual level, and then applying anomaly detection techniques to those models to detect potentially malicious activity. We hypothesize that the workflow and handling of documents has characteristics that can be modeled and analyzed in a manner similar to human activity. That is, we hypothesize that "similar" documents, where similarity might be based on meta-data and/or content, will follow a similar path (or behavior) through an organization. For example, documents created by Alice with similar content might always be reviewed and modified by Bob within a week of creation, followed by editing and moving them to a particular server, etc. We further hypothesize that files that deviate from established patterns merit further investigation and could indicate malicious activity including data exfiltration or malware spreading.

3.2 Attack Time Analysis for Insiders

Florian Arnold (University of Twente, NL)

Attack trees are a widely popular graphical model to represent attack scenarios by defining a system's vulnerabilities and their interdependences. Ray and Poolsapassit extended the basic model by considering attacks from authorized insiders. On the basis of the assumption that attack steps of an intruder might be perfect legitimate operations for an insider, they developed a trimming algorithm to efficiently evaluate an attack tree from the perspective of an insider. We aim to extend their static approach by including a notion of time ands sequencing. The goal of this extended model is to derive a probability distribution for the time until the attacker succeeds. Each attack step is assumed to require a certain time which can be expressed by a acyclic phase-type distribution. The whole attack is then composed by using the maximum, minimum and convolution operation. Based on the work of Pulungan we present a method which can evaluate huge scenarios by compressing the phase-type representation of an attack and introduce an efficient tool chain.

3.3 Bernie the Sheep as a Role Model for Changing Security Behaviour

Debi Ashenden (Cranfield Univ. – Swindon, GB)

How can we persuade end users to behave securely? Accepted approaches have traditionally relied on security awareness programmes, often delivered through computer-based training. Surely if we just give end users the right information in these programmes they'll make good security decisions? As Health Experts have realised this approach to changing behaviour is unlikely to work—otherwise "none of us would be obese, none of us would smoke and none of us would drive like lunatics" (Ian Potter, New Zealand Herald, 2007).

Social marketing, however, is a framework that is increasingly being used to deliver behavioural change for social good (such as healthcare initiatives). It could also offer a promising approach for changing security behaviours in organizations. Social marketing programmes use established marketing concepts such as "exchange" and "competition" to develop an in-depth understanding of the needs and motivations of end users. From this understanding interventions are designed that persuade end users to change their behaviour. "Bernie the Sheep" will be used to illustrate the key concepts of this approach.

3.4 Defining the Cloud Battlefield: Insider Threats in Cloud Computing

Soeren Bleikertz (IBM Research – Zürich, CH)

Joint work of Bleikertz, Sören; Masteli, Toni; Pape, Sebastian; Pieters, Wolter; Dimkov, Trajce
Main reference S. Bleikertz,T. Mastelic, S. Pape, W. Pieters, T. Dimkov, "Defining the Cloud Battlefield: Supporting Security Assessments by Cloud Customers," in Proc. of the IEEE Int'l Conf. on Cloud Engineering (IC2E'13), 2013, to appear.

Cloud computing is becoming more and more popular, but security concerns over shadow its technical and economic benefits. In particular, insider attacks and malicious insiders are considered as one of the major threats and risks in cloud computing. As physical boundaries disappear and a variety of parties are involved in cloud services, it is becoming harder to define a security perimeter that divides insiders from outsiders, therefore making security assessments by cloud customers more difficult.

3.5 Measuring Access, Knowledge, and Trust: A Discussion

Sophie Engle (University of San Francisco, US)

An insider is often defined as a user with access, knowledge, and trust. We propose a semi-automatable approach for a metric of "insiderness" for a user that takes into account all three of these attributes. Specifically, we approximate these attributes by examining access control configurations, access logs, and differences in security policies.

We will use the access control configuration to determine a ratio of configured versus possible access for a particular user, while using access logs to calculate a lower-bound for the amount of knowledge a user has of particular systems and resources. Finally, we will estimate trust by determining the amount of excess access a user has when comparing the configured and feasible levels of policy. This gives an intuitive, concise, semi-automated metric of insiderness that can be used for targeted auditing.

3.6 "Insider Threats" and "Supporting Sustainable Security" – Adding the Dimension of Privacy and Data Protection

Marit Hansen (ULD SH – Kiel, DE)

Privacy and data protection have many flavors: They may be considered as fundamental rights of data subjects, as legally demanded compliance issues, as valuable assets for individuals, for organizations or for the society. Or as an obstacle to security. So what does it mean to add the dimension of privacy and data protection to the discussion on insider threats and support of sustainable security?

First the bad news: The picture gets more complex if further demands and interests have to be considered. Adding the privacy dimension requires a thorough view on the needs, wishes and rights of all parties involved. The good news is that similar procedures and approaches that have proven to be useful in the information security context can be applied when adding, or better: weaving in, privacy and data protection. We will show similarities and differences of the concepts by complementing established security protection goals (confidentiality, integrity, availability) with specific privacy protection goals (unlinkability, transparency, intervenability). These protection goals have to be balanced against each other. Note that the traditional bias towards the organization and its security needs is overcome if the privacy protection goals are taken seriously.

Both technical and organizational measures have to reflect the found balance between the various protection goals. Concerning insider threats, fairness among all parties involved may be achieved by introducing escalating processes that minimize privacy infringements as far (and as long) as possible. Within an organization, checks and balances can be implemented via a staff council and/or an internal data protection officer whose roles in the escalation processes are clearly defined. Sustainable security as well as sustainable privacy and data protection require a regular monitoring of an information security management system (ISMS) and a corresponding, but separate data protection management system. Both systems have to interact to achieve acceptable results.

3.7 Research interests and potential research questions to studying and understanding socio-technical attacks

Jean-Lous Huynen (University of Luxembourg, LU)

My PhD project is about investigating socio-technical aspects of computer security to enrich existing model of security protocols. Kahneman and Tversky's Dual process theory could be a good candidate to better define which triggers are used in socio-technical attacks. Using this theory as a starting point, the purpose of the talk was to ask questions about ways to build experiments to identify these triggers.

3.8 Participatory Designing Work Processes and Security Processes

Kai-Uwe Loser (Ruhr-Universität Bochum, DE)

Joint work of Loser, Kai-Uwe; Nolte, Alexander

Information Security processes are designed processes and they are sociotechnical. As well they are processes of work like other organizational processes. Trust, power, practicability, motivational aspects—like the basic question of who does the work and who has the benefit—or the transparency of rationale behind design decisions are highly relevant here, when it comes to human decision in such a context. Participatory design is an approach that is successful in other contexts, but which is not widely adopted in security processes. One method for the development and reflection of this kind of processes can be supported with the field-tested method of the sociotechnical walkthrough (STWT). Within a project of raising security standards for a university administration infrastructure, STWT was combined with common ISMS methodology. During this project we found indicators for improvement by employing the STWT: technical and organizational measures can be specified in a single effort; contingent relationships can be taken into account as well as vulnerability resulting from characteristics of social structures. Participatory design respects the grounded knowledge of all workers in processes and creates more realistic processes. Descriptions of work processes are supposed to be more realistic instead of abstract ideas of process designers. The earned respect, the joint decisions, the transparent rationale and the more useable processes should contribute to the motivation to obey to the rules of the processes.

3.9 Outsourcing Democracy: Losing Control of e-Voting in the Netherlands

Anne-Marie Oostveen (University of Oxford, GB)

Main reference A.-M. Oostveen, "Outsourcing Democracy: Losing Control of e-Voting in the Netherlands," Policy and Internet , 2(4), pp. 201–220, 2010.
URL http://dx.doi.org/10.2202/1944-2866.1065

Contracting out IT services is a common practice for many governments. This case-study shows that outsourcing is not without risk, especially where elections are concerned. Studying electronic voting in the Netherlands through documents obtained with Freedom of Information requests, we see that government agencies at both local and national level lacked the necessary knowledge and capability to identify appropriate voting technology, to develop and enforce proper security requirements and to monitor performance. Furthermore, over the twenty years that e-voting was used in the Netherlands, the public sector became so dependent on the private sector that a situation evolved where Dutch government lost ownership and control over both the e-voting system and the election process.

3.10 TREsPASS: the socio-technical attack navigator

Wolter Pieters (TU Delft, NL)

Joint work of The TREsPASS consortium
URL http://www.trespass-project.eu/

Information security threats to organizations have changed completely over the last decade, due to the complexity and dynamic nature of infrastructures and attacks. Successful attacks cost society billions a year, impacting vital services and the economy. Examples include Stuxnet, using infected USB sticks to sabotage nuclear plants, and the DigiNotar attack, using fake digital certificates to spy on website traffic. New attacks cleverly exploit multiple organizational vulnerabilities, involving physical security and human behaviour. Defenders need to make rapid decisions regarding which attacks to block, as both infrastructure and attacker knowledge change rapidly.

Current risk management methods provide descriptive tools for assessing threats by systematic brainstorming. Attack opportunities will be identified and prevented only if people can conceive them. In today's dynamic attack landscape, this process is too slow and exceeds the limits of human imaginative capability. Emerging security risks demand tool support to predict, prioritise, and prevent complex attacks systematically.

The TREsPASS project will make this possible, by building an "attack navigator". This navigator makes it possible to say which attack opportunities are possible, which of them are the most urgent, and which countermeasures are most effective. To this end, the project combines knowledge from technical sciences (how vulnerable are protocols and software), social sciences (how likely are people to succumb to social engineering), and state-of-the-art industry processes and tools.

By integrating European expertise on socio-technical security into a widely applicable and standardised framework, TREsPASS will reduce security incidents in Europe, and allow organizations and their customers to make informed decisions about security investments.

This increased resilience of European businesses both large and small is vital to safeguarding the social and economic prospects of Europe.

3.11 From black and white to shades of grey – Static analysis of human behavior

Christian W. Probst (Technical University of Denmark, DK)

Static analysis is a useful tool. It can, for example, compute properties of a system that hold for every possible execution of the system. These characteristics make static analysis very useable for the analysis of systems, and completely unusable for human behavior. This is caused by the fact that static analysis must assume that a system will perform every possible action in a given state; clearly analysing and predicting human behavior requires much more subtle approaches.

We start from modeling systems and analysing them for insider threats. After this we discuss the problem of applying static analysis to human behaviour, and present two solutions, one based on probabilities, the other on simulation.

3.12 Interconnecting Insider-Threat-Defense Responsibilities in Socio-Technical Ecosystems – A managerial perspective

Ingrid Schirmer (Universität Hamburg, DE)

The management of insider threat defense requires integration with IT-Governance processes of an organization based on knowledge of the enterprise architecture. Whereas this integration is a challenge in itself, it is not sufficient.

Today's organizations are acting within socio-technical/business ecosystems with evolving interconnections of organizations, systems, cross-organizational processes, business models with a rapidly growing number of digitally empowered individuals / customers.

Transforming socio-technical ecosystems as a whole by applying IT-innovation (e.g., introducing private clouds) requires in parallel security concepts and their stepwise realization and cultivation within the whole socio-technical ecosystems blurring the insider-focused threat perspective, yet using its achievements. We propose the identification and interconnection of insider threat defense responsibilities in individual organizations as part of a decentralized and long-term security governance in socio-technical ecosystems based on concepts of business ecosystems architectures.

3.13 Where to start with security in healthcare? Dealing with security, insider threat, and increasing mobility in healthcare

Trish Williams (Edith Cowan University – Joondalup, AU)

Drawing on the progression of the Dagstuhl seminars, the definition of insider threat needs to evolve and keep pace with the changing environment. This is presented as a necessity with the context of healthcare as its overarching use-case. The challenges are four fold: integrating security with workflow; influence of trust; a broadening field; and blurring of healthcare devices and software. In addition, there are four assertions to address in relation to mobile health (mHealth) and BYOD, and its impact on the insider threat. In essence, mHealth alters the definition of what constitutes an insider threat. Further, social media, geo-location data and big data are all unaddressed considerations. Before we can do any design it is essential to understand the context and agree (or disagree) on what the impact and effects of the evolution of mobile, social media and the increasing lack of control are. Organizational processes, at this juncture for healthcare, are mature however they do not address the basic security issues regardless of the desire for resilient and sustainable practices. This is compounded by a more fundamental problem of the design-reality gap in incorporating security into health information systems development and the healthcare environment. Whilst those in security appreciate what should be done, we do not "sell" security very well, and thus it is not as effective as we would like or anticipate. The design-reality gap stems from a lack of understanding and engagement with the target environment and a lack of appreciation of its changing nature in light of mobile health. Thus, the issues have to be addressed from the two opposing ends the security continuum: the initial design of security in healthcare systems and the end-user experience, use and consumption of information in health. What problems should be addressed first and how do we prioritise these? How can, or in fact are we able to, address both problems at the same time?

3.14 Problem Statement: System Design beyond Duress Passwords

Alf Zugenmaier (Hochschule München, DE)

Joint work of Zugenmaier, Alf; Coles-Kemp, Lizzie

Currently, most security system designs assume that, after successful authentication, all actions in the system represent the will of the authenticated principal. However, this view does not take into account coercion. Coercion is the threat of a coercer to bring a certain unpleasant consequence upon a coercee, unless the coercee performs (or doesn't perform) a specified action. A coercion attack would be very powerful, as it could be used to turn any attacker into an insider. Thus, it is necessary to define system design principles such that the system becomes resilient to coercion. Coercion resistance mechanisms available in the literature today are mainly limited to e-voting and to duress passwords. We will study the design space framed within the universal credit system for benefit in the UK, which will be digital by default. In this system, citizens will need to obtain an online identity from one of the different online identity assurance providers to participate. Coercion in this system would lead to increased cost in order to provide the required support to the citizens in need.

4 Working Groups

In order to promote a synthesis of the different perspectives represented, the seminar was structured as a combination of collaborative working groups focused on the development of design principles for organizational processes that are sustainably secure and talks. The abstracts from the talks are presented in the next section and reflect the breadth of the disciplines represented at the seminar. Each collaborative working group selected a theme and a scenario from the provocations and explored the organizational processes at work within the scenario. Each group produced their outputs in poster format and presented these outputs to the seminar. In the concluding session on working group work, we used a panel discussion format and the following design principles for organizational processes that are sustainably secure were co-produced:

- Transparency. In the working groups' scenarios transparency in organizational processes and an honest presentation of the values underpinning those proceses was deemed necessary to encourage the desired data handling practices,
- Minimisation. In a number of the working groups, the need for minimisation of data handling practices was an emergent theme. Minimisation as a theme emerged in a number of ways: the need to reduce the number of data handling practices, the need to design simple data handling practices and the need to restrict the number of variants of a particular data handling practice.
- Consent. The theme of consent was reflected in a number of the working groups' conclusions. A spectrum of consent was considered ranging from consent to collect data about staff through to buy in from staff to support a particular organizational process.
- Lawfulness. It was agreed that there needs to be a clear alignment between the legal and regulatory framework in which the organization operates and the organizational processes that are implemented.

5 Open Problems

The design principles require investigation and as such the list of principles produced from the conclusions of the working groups present a set of open problems that require further exploration.

Participants

- Julie Boxwell Ard
Glenn Dale, US
- Florian Arnold
University of Twente, NL
- Debi Ashenden
Cranfield Univ. – Swindon, GB
- Arshid Bashir
RHUL – London, GB
- Sören Bleikertz
IBM Research – Zürich, CH
- Rainer Böhme
Universität Münster, DE
- Lizzie Coles-Kemp
RHUL – London, GB
- Sophie Engle
University of San Francisco, US
- Vaibhav Garg
Indiana University – Bloomington, US
- Carrie Gates
CA Labs -New York, US
- Dieter Gollmann
TU Hamburg-Harburg, DE
- Marit Hansen
ULD SH – Kiel, DE
- Cormac Herley
Microsoft Res. – Redmond, US
- Michael Huth
Imperial College London, GB
- Jean-Lous Huynen
University of Luxembourg, LU
- Dan Ionita
University of Twente, NL
- Florian Kammüller
Middlesex University, GB
- Ana Margarida Leite de Almeida Ferreira
University of Luxembourg, LU
- Makayla Miranda Lewis
RHUL – London, GB
- Kai-Uwe Loser
Ruhr-Universität Bochum, DE
- Anne-Marie Oostveen
University of Oxford, GB
- Wolter Pieters
TU Delft, NL
- Joachim Posegga
Universität Passau, DE
- Marco Prandini
University of Bologna, IT
- Christian W. Probst
Technical Univ. of Denmark, DK
- Ingrid Schirmer
Universität Hamburg, DE
- Sven Übelacker
TU Hamburg-Harburg, DE
- Sam Weber
NSF – Arlington, US
- Sean Whalen
Columbia University, US
- Trish Williams
Edith Cowan University – Joondalup, AU
- Alf Zugenmaier
Hochschule München, DE

Report from Dagstuhl Seminar 12502

Securing Critical Infrastructures from Targeted Attacks

Edited by
Marc Dacier[1], Frank Kargl[2], and Alfonso Valdes[3]

1 **Symantec Research Labs – Sophia Antipolis, FR,** Marc_Dacier@symantec.com
2 **Universität Ulm, DE,** frank.kargl@uni-ulm.de
3 **University of Illinois – Urbana, US,** avaldes@illinois.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12502 "Securing Critical Infrastructures from Targeted Attacks". Through a series of presentations, discussions, and working group meetings, the seminar achieved to shape a clearer picture of what actually constitutes a targeted attack on a critical infrastructure and defined the terms PEST (persistent, sophisticated and targeted) attacks and Critical Cyber Infrastructure in this context. This clearer view will hopefully help the research community and industry to address such threats in a more consistent and holistic way.

Seminar 09.–12. December, 2012 – www.dagstuhl.de/12502
1998 ACM Subject Classification K.6.5 Security and Protection
Keywords and phrases Critical Infrastructures, Targeted Attacks, Security
Digital Object Identifier 10.4230/DagRep.2.12.49

1 Executive Summary

Frank Kargl
Mark Dacier
Alfonso Valdes

The last years have highlighted the fact that our ICT security precautions in many critical infrastructure (CI) systems are clearly insufficient, especially if considering targeted attacks carried out by resourceful and motivated individuals or organizations. Critical infrastructures, like energy or water provisioning, transportation, telecommunication, or health support are relying to an ever-larger extent on ICT, often being monitored or controlled in a semi or fully automated way. Disruption of these control processes could turn out to be disastrous, especially as many of these systems are cyber-physical systems that interact with the real world through sensors and actuators and can thus have a direct influence on the physical world not mediated by the common sense of a human being.

Rendering ICT systems in such critical infrastructure unusable or malfunctioning can cause huge economical damages or even endanger human lives. Some examples: it is reported by the Institute for Science and International Security (ISIS) in December 2010[1] that the Stuxnet malware actually damaged around 1000 Uranium enrichment centrifuges in the

[1] http://www.isis-online.org/isis-reports/detail/did-stuxnet-take-out-1000-centrifuges-at-the-natanz-enrichment-plant/

Iranian enrichment facility in Natanz (which was possibly its goal). If the same would happen in a European Uranium enrichment facility, the economical damage would be significant and danger to population due to failure of systems could not be ruled out completely. In 2000, an insider attack on a sewage treatment facility in Queensland, Australia caused millions of liters of raw sewage to spill out into local parks and rivers[2]. The CIP Vigilance Blog collects a long list of such issues[3].

There are many similar examples where Industrial Control Systems (ICS) have been affected due to insufficient security precautions. Moreover, the apparent success in infiltrating Critical Infrastructure environments is calling attention on the ineffectiveness of standard security mechanisms in detecting similar attacks. Stuxnet is believed to have been operating undetected for almost one year leveraging multiple vulnerabilities that were previously unknown, and has been discovered only as a consequence to an operational anomaly that triggered the attention of the field operators. This fact clearly shows that not only our security mechanisms in ICS are insufficient, but that even our methods to find vulnerabilities and detect ongoing or successful attacks in Critical Infrastructure environments are not up to their task. It is very likely that Stuxnet could be the "first of a kind", as demonstrated by the recent apparition of the so-called Duqu threat, apparently based on the same code (see the Symantec thorough analysis for more on this topic[4]).

Similar argumentation can be applied to other forms of control systems like Intelligent Transport Systems, modern health systems, Smart electric grids, and many more. The advanced metering infrastructures (AMI) now being deployed in some form on the electric grids of many countries offers potential benefits in terms of reduction of peak load, which in turn enables green house gas reduction and various economic benefits. However, it introduces potentially hundreds of millions of computationally limited networked endpoints outside of a defensible physical or electronic perimeter. Moreover, smart grids may be subject to attacks that do not require an adversary to compromise a device, whether a smart meter on a residence or a phasor measurement unit (PMU) that contributes to wide area measurement or state estimation. Real-time price signals communicated to smart meters may induce volatility, and if spoofed may lead to destabilizing load fluctuation (see [1]). Spoofing of GPS signals can cause PMUs to lose synchronization, resulting in threats to real-time control and corrupt grid state estimation.

There are many challenges involved in this, especially the heterogeneity of the systems that often involve legacy and proprietary system where not even all specification might be available to security engineers. High dependability and availability requirements of such systems often do not allow fast update cycles in case of security vulnerabilities are disclosed. The trend to use more COTS hardware and software in such systems creates problems and opportunities at the same time. A problem is that all malware that is available in such systems suddenly also becomes available to attackers on Critical Infrastructure ICT and that a lot of known vulnerabilities become exploitable. On the pro side, many established security mechanisms like firewalls, Intrusion Detection Systems, or OS security mechanisms like malware scanners can be applied. However, you often need to specifically adjust them for the new domain (e.g., by having SCADA specific signatures for an IDS). At the same time, the different (dependability) requirements and different applications in Critical Infrastructure

---

[2] http://www.theregister.co.uk/2001/10/31/hacker_jailed_for_revenge_sewage/
[3] http://ciip.wordpress.com/2009/06/21/a-list-of-reported-scada-incidents/
[4] http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_duqu_the_precursor_to_the_next_stuxnet.pdf

Systems often require new or updated approaches, e.g., regarding security updating or security testing methodologies.

The research community has taken up this challenge, as can be seen by the emergence of specific research projects (e.g., EU projects like ReSIST, IRIIS, VIKING, SERSCIS, INSPIRE, CRUTIAL, CRISALIS), and regular contributions on the topic at conferences and workshop (RAID, DIMVA, CCS, LEET, IEEE SSP, NDSS, Usenix Security, etc.). The US Department of Homeland Security and Department of Energy fund numerous projects under programs such as the National SCADA Test Bed (NSTB) and Cyber Security for Energy Delivery Systems (CSEDS). However, we identified that the research community would benefit from being better connected, having identified a clear list of major research challenges, and knowing to what extent they have been addressed so far. Stemming from this motivation, we proposed this Dagstuhl research seminar with the goal to bring together leading researchers both from academia and industry to discuss and evaluate the state of the art and to highlight where sufficient solutions exist today, where better alternatives need to be found, and also to give directions where to look for such alternatives.

One of the most important aspects was to identify whether security challenges and solutions apply to all different areas of CI, be it water, electricity, gas, transport, health-support, public safety infrastructures, or tele-communication. Our initial expectation was that there would be clusters of domains with very similar profiles on the one hand, but also large differences between clusters. This, however, was not clear previously, as many security researchers focused on specific areas or specific aspects of security.

Beyond, during the seminar we also focused on the question how targeted attacks on CI differ from ubiquitous unspecific attacks by malware or occasional hackers. As the later do not focus specifically on CI, they will typically not create large-scale damages — if damages occur, this is typically the consequence of computer systems being down. In contrast, the Stuxnet example illustrates how targeted malware can be injected into target systems in a very stealthy way and can cause subtle damage that can go unnoticed for a long time. Consequently, security countermeasures, reactions, and forensic methods have to differ as well. However, the research community has just started to address the area of targeted attacks.

The seminar started from a set of questions related to this:

- What are the specific security challenges and requirements that are ubiquitous throughout different Critical Infrastructure domains? Where do those domains differ in terms of security?
- What is the status with respect to protection from, detection, and analysis of targeted attacks on Critical Infrastructures? What solutions can be transferred from general ICT? Where have new solutions already been found? Where is further research needed?
- Do these solutions apply to Critical Infrastructures in general, or do we need to work on domain-specific solutions?
- How can the negative effects of successful attacks be contained?
- How can CI be made resilient to attack, and able to maintain critical (possibly degraded) function in the presence of attack?
- How can we bridge the gap between low-level research on the granularity of individual ICT devices or single networks, e.g., to conduct forensic analysis or deploy IDS, on the one hand and on the other the research that assesses the system-wide effects of targeted attacks, e.g., on effect propagation?
- How can technical solutions and organizational policies be aligned and enhanced in a consistent way?

- How can we bridge the gap of knowledge between security experts rarely aware of the specific characteristics of CI systems and CI experts not necessarily up to date with the latest security research outcomes.
- How can we shed some light on CI insecurity without running the risks of opening a Pandora box? What are the consequences of such risks? Are there legal implications to consider?
- How do the approaches of academia and industry in addressing targeted attacks on CI differ?

Many of these questions were addressed during our two and a half day Dagstuhl seminar from December 9 to 12, 2012. We had the envisioned nice mix of participants with an industry participation of over 30 % and experts from various domains of critical infrastructures.

The agenda featured two main plenary talks, nine short presentations, and regular working group breakout sessions. The plenary talks were given by Alvaro Cárdenas Mora from Fujitsu Laboratories of America / UT Dallas who spoke on “Short-term and Long-term research Challenges for Securing Cyber-Physical Systems” and Levente Buttyán from the BME CrySyS Lab who gave us a first hand insight into analysing targeted attack malware in his presentation on “The cousins of Stuxnet: Duqu, Flame, and Gauss”.

The short presentations focussed on a broad variety of topics, some giving broader updates on research agendas and activities like the European CRISALIS project, some others addressing specific areas like train control systems or smart grids. One short talk by Felix Freiling asked the challenging question whether detecting targeted attacks has to be considered impossible by their very nature, a discussion that working group 2 later continued in depth. Other topics addressed in the short presentations included intrusion detection mechanisms for industrial control systems, a report on the CERT run by Siemens, and on societal consequences of cyber attacks on electrical supply systems.

These talks provided perfect input for our working groups. We initially envisioned four working groups with the titles (1) Business Aspects of Security for CI in Different Domains, (2) Attacker Models and Risk Analysis for Targeted Attacks on CI, (3) CI Security in different CI Domains, and (4) Technical Security Approaches for Intrusion Detection and Network Monitoring. However, during initial discussions and working group assignments, groups (1) and (3) found their topics to be closely related and decided to merge.

The merged working groups (1) and (3) first identified the challenge that definitions of what a *Critical Infrastructure* actually is are quite diverse and fuzzy which led to a narrowed down definition of *Critical Cyber Infrastructure* which provided a working definition to then reason about the nature of security incidents and solutions in such systems. One conclusion from their work was that there is a gap between generic IT security and the large number of different CI domains that may be bridged by providing clearer reference scenarios that security researchers can focus on. That way, one could then identify whether generic security solutions can be applied to such scenarios or even cross- scenario or whether specific solutions need to be found.

Working group (2) mostly investigated attacker models for targeted attacks. Again, the term targeted attack was not clear and the first result of the working group was a attack classification scheme to be able to narrow down on this term and distinguish various types of attacks. The group even went beyond the targeted attack term and suggested PEST (persistent, sophisticated and targeted) as a categorization of the most critical types of attacks. In a second meeting, the working group discussed attacker motivations and identified a clear lack of intelligence regarding such motivations. Therefore, a lot today is more guesswork than based on clear facts and more investigations into the nature or PEST attacks seems to be required.

Finally, working group (4) focussed on the technical topic of intrusion detection and network monitoring in ICS, coming up with a list of attack scenarios, technical challenges and ideas for enhancement of countermeasures.

In a final plenary wrap-up discussion, all participants agreed that the seminar's topic was definitely a very challenging one. As both the definition of Critical Infrastructure and Targeted Attack are not even clearly agreed upon and as CIs are so diverse, we were not even able to cover all possible instantiations of CIs by dedicated experts. Especially the work in the working groups provided important first steps towards clearer definition and a common understanding of these issues and as such the seminar has to be considered a success that should be followed up by future activities.

The question whether joint security approaches and solutions for targeted attacks on critical infrastructures can be found can therefore not finally be answered. However, the research community and industry would definitely benefit from a closer cooperation of researchers and practitioners that work on *PEST attacks on Critical Cyber Infrastructures*.

References

1 Roozbehani, M.; Dahleh, M.A.; Mitter, S.K.,"Volatility of Power Grids Under Real-Time Pricing", Power Systems, IEEE Transactions on , vol.27, no.4, pp.1926,1940, Nov. 2012

2 Table of Contents

3 Overview of Talks

3.1 Assessment of Social Impact Costs and Social Impact Magnitude from Breakdowns in Critical Infrastructures

Gunnar Björkman (ABB AG – Mannheim, DE)

This presentation described a method and a tool to calculate societal consequences from outages in the electrical supply. The method and tool have been developed in the EU/Framework 7 project VIKING that was successfully ended in November 2011, see http://www.vikingproject.eu/.

Two types of societal consequences from power blackouts were presented; one that calculates the economic losses to society as lost Gross Domestic Product (GDP), i.e. Social Impact Cost (SIC), and the other as a new type of measure for non- economic damages, Social Impact Magnitude (SIM). For the economic loss calculation, SIC, the national GDP is broken down on individual object level, e.g. public services, small and big companies, hospitals, etc., with a high time resolution. The breakdown of GDP makes it possible to calculate the economic activities for smaller parts of the country and for defined times. The economic loss for society is then calculated as the difference in economic activity for a certain geographic part and for a defined time with and without electrical supply considering the stepwise electrical restoration procedures.

The Social Impact Magnitude (SIM) is a new logarithmic measure considering the number of people impacted by the outage and the outage length. Using the 10- logarithm of the outage length in seconds and thousands of people a measure is reached that closely resembles the well-known Richter scale for earthquakes and is very easy to calculate. Applying the SIM measure on previous, well-known power blackouts gives an intuitively very reasonable value, e.g. the 2003 Northeast American blackout gets a value of 9,67, i.e. a very serious disturbance. The intention of the SIM is to be able to classify power outages in an-easy-to understandable way and to use such classification to plan society responses, e.g. a value below 6 has only regional consequences but values above require a national response.

3.2 The cousins of Stuxnet: Duqu, Flame, Gauss

Levente Buttyán (Budapest Univ. of Technology & Economics, HU)

Joint work of Bencsáth, B.; Pék, G.; Buttyán, L.; Félegyházi, M.
Main reference B. Bencsáth, G. Pék, L. Buttyán, M. Félegyházi, "The Cousins of Stuxnet: Duqu, Flame, and Gauss," in Future Internet, Vol. 4, Issue 4, pp. 971–1003, 2012.
URL http://dx.doi.org/10.3390/fi4040971

Stuxnet was the first targeted malware that received worldwide attention for causing physical damage in an industrial infrastructure seemingly isolated from the online world.

Stuxnet was a powerful targeted cyber-attack, and soon other malware samples were discovered that belong to this family. This presentation first presented our analysis of Duqu, an information-collecting malware sharing striking similarities with Stuxnet. It described our contributions in the investigation ranging from the original detection of Duqu via finding the dropper file to the design of a Duqu detector toolkit and then continued with the analysis of the Flame advanced information-gathering malware.

Flame is unique in the sense that it used advanced cryptographic techniques to masquerade as a legitimate proxy for the Windows Update service. The talk also presented the newest member of the family, called Gauss, whose unique feature is that one of its modules is encrypted such that it can only be decrypted on its target system; hence, the research community has not yet been able to analyze this module. For this particular malware, the authors designed a Gauss detector service and we are currently collecting intelligence information to be able to break its very special encryption mechanism. Besides explaining the operation of these pieces of malware, the presentation also examined if and how they could have been detected by vigilant system administrators manually or in a semi-automated manner using available tools. Finally, the talk discussed lessons that the community can learn from these incidents.

3.3 Short and Long-Term Research Challenges for Protecting Critical Infrastructure Systems

Alvaro A. Cárdenas (The University of Texas – Dallas, US)

Our critical infrastructure systems are being modernized with information and communication technologies to face the operational requirements and efficiency challenges of the 21st century. The smart grid in particular, will introduce millions of new intelligent components to the electric grid, buildings, and homes within the next decade. While this modernization will bring many operational benefits to infrastructure systems, it will also introduce new vulnerabilities, a larger attack surface, and raise privacy concerns.

This presentation focused on the short, medium, and long-term research challenges for protecting cyber-physical systems.

As a short-term goal, it discussed some of the incentives (economic or regulation) to develop, deploy, and maintain control systems following security best practices.

In the medium-term discussion, it focused on the large-scale instrumentation being deployed in critical infrastructure, and the advantages of analyzing this data for better security intelligence. The talk exemplified some of these notions with smart grid data being used for electricity theft and anomaly detection.

The final part of the talk focused on long-term research projects and included the fact that we can combine physical dynamical models of the critical infrastructure with information security models to obtain more resilient and survivable systems against targeted attacks. This part of the talk discussed examples for survivable control of a chemical reactor and other generic control systems resilient to false-data injection and DoS attacks.

3.4 Signature-Less Network Intrusion Detection for Industrial Control Systems

Sandro Etalle (TU Eindhoven, NL)

This talk presented a new technology to do network monitoring and intrusion detection on industrial control system networks, and touched on the limits of the classic intrusion detection technology when applied to industrial control systems.

3.5 Detecting Targeted Attacks Considered Impossible

Felix C. Freiling (Universität Erlangen-Nürnberg, DE)

As a followup to a presentation at the Dagstuhl Seminar 12061, this talk tried to stimulate discussion on what a targeted attack is as opposed to "mass-malware" attacks like Conficker and Storm.

3.6 On the (In)Security of Train Control Systems

Stefan Katzenbeisser (TU Darmstadt, DE)

In the talk, the author described the state-of-the art technologies that are used to safely control trains in the network of German rail. The presentation detailed the future architecture of train control systems, discussed related security aspects and briefly reported on an ongoing project in collaboration with DB Netz that attempts to define the security architecture of next-generation signal boxes.

3.7 The future of Smart Grids

Erwin Kooi (Alliander – Duiven, NL)

The change from a production-follows-demand to a demand-follows-production model introduces challenges for grid operation. These challenges can be addressed by installing more and heavier powerlines or by using data to manage the grid, production and demand more intelligently.

3.8 CRISALIS – Preventing Targeted Attacks on Critical Infrastructures

Corrado Leita (Symantec Research Labs – Sophia Antipolis, FR)

The presentation introduced the CRISALIS FP7 project that aims at providing new means to secure critical infrastructure environments from targeted attacks, carried out by resourceful and motivated individuals.

The discovery of highly sophisticated and targeted attacks such as Stuxnet showed that these threats are a reality. Their success in infiltrating Critical Infrastructure environments is calling attention on the ineffectiveness of standard security mechanisms at detecting them.

Stuxnet, for instance, is believed to have been operating undetected for almost one year leveraging multiple vulnerabilities that were previously unknown, and has been discovered only as a consequence to an operational anomaly that triggered the attention of the field operators. This fact clearly shows that our methods to find vulnerabilities and detect ongoing or successful attacks in critical infrastructure environments are not sufficient. The talk gave an overview over the CRISALIS project which aims at filling this gap with practical, short-term solutions.

See http://www.crisalis-project.eu/ for details.

3.9 A CERT for products' perspective on security

Tobias Limmer (Siemens – München, DE)

Multiple trends in cyber security affect product vendors of industrial IT. On the one hand, the amount of targeted attacks that involve critical infrastructures is rising. On the other hand, attention on non-office IT environments by security researchers is rising. This effect is caused by the adaptation of standards, protocols and paradigms that are prevalent in the office IT world into the industrial world. Attack methodologies that have long been used in standard IT environments can increasingly be reused in industrial environments, easing the work of security researchers, both on the black hat and white hat side.

The heterogeneity of devices is an additional problem, as well as static environments that were implemented with focus on safety. One example regarding static environments is the problem of continuous patch deployment of security updates when acceptance tests are required after each system change. Authentication is another problem, as strong authentication methods contradict with quick information transfer that is needed for safety functions.

Product vendors need to adapt to this changed environment by implementing preventative (e.g., secure development) and reactive (e.g., vulnerability handling) measures and focus on bridging the gap between product security and ease-of-use for customers.

3.10 Smart Grid Security Research

Alfonso Valdes (University of Illinois – Urbana, US)

Cyber assets in infrastructure systems such as smart electrical grids potentially enable advances in efficiency and resiliency, but are potentially attractive targets for cyber attack. Securing such systems presents particular challenges, since security solutions and practices from conventional enterprise systems are not always applicable. On the other hand, the constrained function and regularity of communication in infrastructure systems allows security solutions based on such approaches as specification-based intrusion detection. This talk summarized work from SRI International and the University of Illinois towards securing cyber assets in infrastructure systems, with an emphasis on smart electric grids.

4 Working Groups

4.1 Report of Joint Working Groups 1 & 3

Nils Aschenbruck and Working Group 1 & 3 Participants

The working groups 1 on "Business Aspects of Security for Critical Infrastructure in Different Domains" and 3 on "CI Security in different Critical Infrastructure Domains" decided to merge at the beginning of the seminar. So this is their joint report.

4.1.1 Definition and Classification

For both working groups the idea was to classify the different domains of critical infrastructures first. Then, the plan was to examine business aspects (group 1) and security measures, respectively.o

In the literature, there are different classifications. In the US [1], there are 18 critical infrastructures and key resources (CIKR). In the EU [2], there are 9 critical infrastructures. Further examples, e.g., [3], may be found. All the classifications were found to be quite fuzzy, mixing domains and physical infrastructure/facilities, e.g., "banking sector" and "national monuments". Thus, the working groups decided to focus on selected critical cyber infrastructures, while also considering future development such as intelligent cities and smart power grids.

We reached a rough consensus that a critical cyber infrastructure is characterized by the following attributes: (1) the highest impact attack is a cyber attack; (2) a cyber attack potentially results injury or loss of life (although attacks can have economic consequences, such as energy theft or blackmail) (3) there is specific physical infrastructure involved; (4) there is currently rapid adoption of cyber technology due to powerful business cases for the adoption of automation. Concerning security, the specific differences lie in the targets to attack such as: (1) (distributed) control systems; (2) distributed algorithms; (3) physical entities; (3) supply chain; (4) social engineering.

Compared to a standard IT attack, where the machine itself is the goal, targeted attacks on critical cyber infrastructure often have specific attributes. The attacker has a plan for a higher goal. By doing so, he may targets specific groups, or classes of assets. This means that the propagation in contrast to a massive attack is (up to infinitely) focused; sometimes just a single machine is attacked. Usually, the attacker has background knowledge and/or part of the attack is getting more detailed knowledge. The attacker is well prepared and has large resources (knowledge, expertise, time, hardware, etc.).

For the targeted critical infrastructures often the following attributes apply: Standard security solutions, e.g., virus updates, can not be used. The systems originally started as isolated systems, but are now connected to enterprise systems due to business incentives (e.g., remote access to SCADA). Furthermore, classical independent safety loops within the industrial control system are now connected to intelligent electronic devices (IED). Thus, the former hardware independence of the safety loops is lost. Moreover, different critical infrastructures are connected; with consolidation of assets in the electric sector, for example, system operators may operate several distinct energy resources as a virtual power plant. By doing so, access control is a big challenge.

4.1.2 Challenges and Future Research Directions

A core challenge is the lack of domain knowledge. The security research is mainly IT driven, nowadays. On a high level, all challenges seem to be trivial as solutions do already exist in the IT domain. On a low level, everything seems to be Ethernet, which is well known in the IT domain as well. However, there are specific challenges for cyber critical infrastructures. But the specific domains have to be understood first. It would be great to have a set of reference scenarios. Such a set could help to understand different domains and to develop, parameterize, and evaluate security solutions for the reference scenarios. By doing so, it could be examined which solutions are applicable across domains. To figure this out is probably the core challenge regarding critical cyber infrastructures. Besides, the supply chain should be addressed as well. Verifying that hardware has no undocumented functionality is a difficult problem.

References

1 ICS-CERT FAQ, http://www.us-cert.gov/control_systems/csfaq.html.
2 Communication from the Commission to the Council and the European Parliament of 20 October 2004 – Critical Infrastructure Protection in the fight against terrorism [COM(2004) 702 final – Not published in the Official Journal], http://europa.eu/legislation_summaries/justice_freedom_security/fight_against_terrorism/l33259_en.htm.
3 Bundesamt für Bevölkerungsschutz und Katastrophenhilfe, Critical Infrastructures divided by sectors and subsectors, http://www.bbk.bund.de/SharedDocs/Downloads/BBK/DE/Downloads/Kritis/CI_Sectors_Subsectors.pdf.

4.2 Report of Working Group 2 “Attacker Models for Targeted Attacks”

Stefan Katzenbeisser and Working Group 2 Participants

The main goal of working group 2 was to define an attacker model for attacks on critical infrastructures. The participants agreed that the attacker model should include at least the following dimensions:

Targeted vs. opportunistic attacks: for targeted attacks the set of victims is known and rather small, while opportunistic attackers typically do not know the identity of the victims beforehand and do not target a specific group of victims.

Persistence vs. one-time/one-off: A persistent attacker tries to achieve his goal during a long period of time; the attack evolves over time. Even if the victim cleans her infrastructure, a persistent attacker will come back and try to infect through different paths. On the contrary, mass infections can be classified as "one-time/one-off"; once the system is cleaned and the vulnerability is fixed, the attacker will not come back to a specific victim.

Advanced skills vs. script kiddies: Advanced attacks are resourceful and are willing to devote manpower and time. However, the computational resources of an attacker are not necessarily a good measure for judging the level of sophistication, since botnets can provide virtually unlimited compute power. On the contrary, script kiddies have limited skills and resources.

Well-funded organization vs. hobbyist: Another important dimension is the funding available for carrying out the attack.

High impact vs. low impact: The impact of a targeted attack needs to be taken into account, where impact is measured by the damage potential per victim.

Lots of intelligence vs. no background information: Attacks against critical infrastructures are often carried out by attackers with a broad knowledge of the system. There are many similarities between targeted attacks and classical intelligence methods (such as placing a spy in the government of a foreign country).

Human, flexible, adaptive attack vs. fully automated attack: We can observe a big difference in sophistication between attacks that are directly controlled by humans and fully automated attacks. The use of human labor can be a bottleneck, but also allows more flexibility. Flexible attacks can also involve human factor issues, such as social engineering.

The working group concluded that attacks against critical infrastructures are typically characterized as being persistent, sophisticated and targeted (PEST).

The second meeting of the working group focussed on the goals of an attacker who is performing a targeted attack. In general, the goals can be very diverse, such as espionage, information gathering, sabotage, blackmailing, destruction, gaining financial advantage and infiltration.

Unfortunately the precise motivation of an attacker is often unclear; we have not much information about the origins of attacks. The often-cited goal of destructing a critical infrastructure is very close to a cyberwar scenario; nevertheless, there can also be situations that are closely related to cybercrime (such as attacks by insiders, terrorist groups, hacker groups and political activists such as anti-nuclear activists).

Regarding the attack strategy, one should distinguish between the intelligence gathering phase and the actual attack. We must assume that an attacker has sufficient expertise to launch multi-stage attacks that evolve over time and that starts with an intelligence gathering phase, which is followed by an execution phase. We must also distinguish between attacks against the processes and the infrastructure. This distinction is reinforced by the fact that an attacker needs to follow different attack paths when he wants to inflict physical destruction or to gain information (for example by stealing classified documents); for the former, the attacker needs to gain access to the production infrastructure. During the attack execution phase, different goals can be achieved such as denial of service, taking control or sabotage; in particular the latter one may be a goal that is not particularly relevant in classical IT-based infrastructures.

4.3 Report of Working Group 4 "Technical Security Approaches for Intrusion Detection and Network Monitoring"

Damiano Bolzoni, Marco Caselli, Emmanuele Zambon, and Working Group 4 Participants

Last but not least, working group (4) focussed on the technical topic of intrusion detection and network monitoring in ICS, coming up with a list of attack scenarios, technical challenges and ideas for enhancement of countermeasures. The participants discussed such diverse approaches as network- and protocol-based analysis or PLC honeypots, including network-based packet-header periodicity analysis that could help operators build better firewall rules. Practical challenges include issues like rogue hardware being introduced into the system by attackers. One conclusion was that some of the more advanced detection countermeasures may not be ripe yet to be implemented by vendors or would require a very high effort to be deployed in production systems. This is definitely a field where research needs to interact even more closely with vendors to come up with solutions to work on joint definitions, reference models and attacks, testbeds, etc., to start a more focussed research with more practical outcome.

Participants

- Magnus Almgren
Chalmers UT – Göteborg, SE
- Nils Aschenbruck
Universität Osnabrück, DE
- Davide Balzarotti
Institut Eurecom – Sophia Antipolis, FR
- Rafael Barbosa
University of Twente, NL
- Gunnar Bjoerkman
ABB AG – Mannheim, DE
- Damiano Bolzoni
University of Twente, NL
- Levente Buttyan
Budapest Univ. of Technology & Economics, HU
- Alvaro Cárdenas Mora
The Univ. of Texas – Dallas, US
- Marco Caselli
University of Twente, NL
- Marc Dacier
Symantec Research Labs, US
- Sandro Etalle
TU Eindhoven, NL
- Felix C. Freiling
Univ. Erlangen-Nürnberg, DE
- Jakob Fritz
EURECOM – Biot, FR
- Elmar Gerhards-Padilla
Fraunhofer FKIE – Wachtberg, DE
- Dina Hadziosmanovic
University of Twente, NL
- Frank Kargl
Uni Twente, NL & Uni Ulm, DE
- Stefan Katzenbeisser
TU Darmstadt, DE
- Erwin Kooi
Alliander – Duiven, NL
- Maryna Krotofil
TU Hamburg-Harburg, DE
- Klaus Kursawe
ENCS – The Hague, NL
- Corrado Leita
Symantec Research Labs – Sophia Antipolis, FR
- Tobias Limmer
Siemens – München, DE
- Michael Munzert
Siemens – München, DE
- Heiko Patzlaff
Siemens – München, DE
- Andreas Paul
BTU Cottbus, DE
- Franka Schuster
BTU Cottbus, DE
- Valentin Tudor
Chalmers UT – Göteborg, SE
- Alfonso Valdes
Univ. of Illinois – Urbana, US
- Stephen Wolthusen
RHUL – London, GB
- Emmanuele Zambon
SecurityMatters B.V. – Enschede, NL

Report from Dagstuhl Seminar 12511

Divide and Conquer: the Quest for Compositional Design and Analysis

Edited by

Marieke Huisman[1], Barbara Jobstmann[2], Ina Schaefer[3], and Marielle Stoelinga[4]

1 **University of Twente, NL,** Marieke.Huisman@ewi.utwente.nl
2 **VERIMAG – Gières, FR,** Barbara.Jobstmann@imag.fr
3 **TU Braunschweig, DE,** i.schaefer@tu-braunschweig.de
4 **University of Twente, NL,** Marielle.Stoelinga@ewi.utwente.nl

Abstract

On December 16 to 21, the Dagstuhl seminar *Divide and Conquer: the Quest for Compositional Design and Analysis* was organized. Topic was the *compositionality*, a central theme in computer science, but its applications, methods, techniques are scattered around many different disciplines. Therefore, this workshop brought together scientists from different disciplines, including deductive verification, model checking, software product lines, component interfaces.

Seminar 16.–21. December, 2012 – www.dagstuhl.de/12511

1998 ACM Subject Classification D.2.4 Software/Program Verification, D.2.5 Testing and Debugging, F.3 Logics and Meanings of Programs, F.3.1 Specifying and Verifying and Reasoning about Programs

Keywords and phrases Algorithms, Design, Languages, Theory, Verification

Digital Object Identifier 10.4230/DagRep.2.12.64

1 Executive Summary

Marieke Huisman
Barbara Jobstmann
Ina Schaefer
Marielle Stoelinga

Compositionality is a key concept in computer science: only by breaking down a large system into smaller pieces, we can build today's complex software and hardware systems. The same holds true for verification and analysis: realistic systems can only be analyzed by chopping them up into smaller parts. Thus, compositionality has been widely studied in various different settings, and by different communities: people in programming languages, software verification, and model checking have all come up with their own techniques and solutions.

Thus, the goal of this workshop has been to bring together these fields and communities, so that they can learn from and cross-fertilize each other. We have succeeded in doing so: through three extensive tutorials, longer and shorter presentations, and working sessions, researchers from different areas have learned about each others problems, techniques, and approaches.

The scientific programme was built around four corners stones

1. Personal introductions.
2. Three well-received tutorials:
   - Compositional programming by Oscar Nierstrasz
   - Compositional verification by Arnd Poetzsch-Heffter
   - Compositional modelling by Arend Rensink
3. Regular presentations, presenting in-depth technical knowlegde on:
   - Verification of programming languages
   - Automatic synthesis
   - Interface theories
   - Model checking
   - Contract-based design
   - Software product lines
4. Working group sessions:
   - Working group on software product lines
   - Working group on Benchmark for Industrial Verification/Synthesis Problems
   - Working group on Modular Full Functional Specification and Verification of C and Java programs that Perform I/O
   - Model checking vd deducutive verification
   - Compositional Synthesis of Reactive Systems

2 Table of Contents

Working Groups

3 Detailed programme

3.1 Personal introductions

We started out by a three rounds of personal introductions of the participants. Each participant was asked to introduced him or herself in two minutes, focussing on scientific interests. Also, we asked participants to classify themselves according to three basic scientific disciplines, being modeling, verification and/or programming languages. In our opinion, these introductions really helped to get to know each other, to easy communication and discussion, and "break the ice."

3.2 Tutorials

In order to set a common ground between the various fields, we asked three renown experts to give a tutorial, explaining the basic concepts, methods, and techniques in their field. In our opinion, these tutorials were a success, since their quality was excellent, and they were very well received by the participants.

3.2.1 Tutorial on compositional programming by Oscar Nierstrasz

Oscar Nierstrasz started out by a historical overview of compositional techniques in programming languages — that we all know, but were not aware of their historic origins. Then moved to current techniques in compositional programming, and challenges for the future.

3.2.2 Tutorial on compositional verification by Arnd Poetzsch-Heffter

Poetzsch-Heffter started with the presentation of basic principles of compositional verification and suggested to categorize techniques for compositional verification according to four dimensions: (A) What are the components? How are they represented? (B) What are the composition mechanisms? (C) What properties are addressed? How are they specified? (D) What kinds of verification techniques are used? In the main part, the tutorial took a closer look at four different settings to discuss important aspects in the huge space of compositional verification: (1.) Maximal models for model checking step-synchronous Moore machines (2.) Model checking control-flow properties of sequential procedural programs (3.) Modular state-based verification of object-oriented programs (4.) Assume-guarantee reasoning for communicating processes

3.2.3 Tutorial on compositional modelling by Arend Rensink

This presentation formulated a simple formal framework in which some of the essential aspects of compositionality come together. The most important lesson is that the composition operator should be compatible with a notion of abstraction, or semantics, encapsulating the conceptual understanding of the objects under construction. If this compatibility fails, the common solution is to augment the abstraction by putting more information into it. This kind of augmentation is, in fact, the same kind of step as strengthening the induction hypothesis in an inductive proof.

The framework is tested, with varying degrees of success, against a number of cases of composition from different domains, ranging from bisimilarity minimisation to testing and subclassing (seen as a composition operator). Some of these cases nicely illustrate the principle of augmentation; others provide examples where compositionality simply fails to hold.

4 Overview of Talks

4.1 Compositional verification and semantics of concurrent/distributed objects

Wolfgang Ahrendt (Chalmers UT – Göteborg, SE)

We present a semantics, calculus, and system for compositional verification of an object-oriented modelling language for concurrent distributed applications. The system is an instance of KeY, a framework for object-oriented software verification, which has so far been applied foremost to sequential Java. The presented system addresses functional correctness of models featuring local cooperative thread parallelism and global communication via asynchronous method calls. The calculus heavily operates on communication histories specified by interfaces. We also present adenotational semantics and an assumption-commitment style semantics of the logic.

4.2 A Coinductive Big-step Semantics for Distributed Concurrent Objects

Wolfgang Ahrendt (Chalmers UT – Göteborg, SE)

We present a fully compositional big-step operational semantics for globally distributed and locally concurrent objects. The semantics captures both, terminating and non-terminating behaviour. We construct thread histories independently, non-deterministically guessing effects of other threads/objects. Then thread histories are merged to object history, which then are merged to system histories.

4.3 Glue synthesis in BIP

Simon Bliudze (EPFL – Lausanne, CH)

BIP (Behaviour, Interaction, Priority) is a framework for the component-based design and analysis of real-time embedded systems. It has been successfully used for modelling and analysis of a variety of case studies and applications, such as performance evaluation, modelling and analysis of Tiny OS-based wireless sensor network applications, construction and verification of a robotic system. The main characteristic feature of BIP is the clear separation of component behaviour (sequential computation) and coordination. The latter is realised through memory less "glue" consisting of an interaction model defining synchronisations among components and a priority model used for conflict resolution and defining scheduling policies. In this talk, I will briefly discuss automatic synthesis of glue from boolean constraints representing a certain type of safety properties made possible due to the above-mentioned separation of concerns principle.

4.4 Defining a general abstract notion of component

Simon Bliudze (EPFL – Lausanne, CH)

Component-based design generally relies on a clear separation between the mechanisms used to design atomic components and those used to define the "glue" assembling these into higher-level compound components. Different component-based design approaches have different models for components and different – often ad-hoc – glue operators. Defining the appropriate glue for agiven component model and studying glue properties such as composition and interference, requires a formal and generic definition of the notion of glue. The necessity of such definition becomes even clearer when one tries to compare different component frameworks. However, defining a formal generic notion of glue requires first a formal generic notion of component. In a recent paper (http://rvg.web.cse.unsw.edu.au/eptcs/paper.cgi?ICE2012.6), we have proposed one such generic definition. Expected outcome: For this discussion group, I propose to confront this definition with the critique based on other participants' experience and either validate it through examples or give a more appropriate one.

4.5 Compositional Reasoning about Object-Oriented Software Evolution

Einar Broch Johnsen (University of Oslo, NO)

An intrinsic property of real world software is that it needs to evolve. The software is continuously changed during the initial development phase, and existing software may need modifications to meet new requirements. To facilitate the development and maintenance of programs, it is an advantage to have programming environments which allow the developer to alternate between programming and verification tasks in a flexible manner and which ensures correctness of the final program with respect to specified behavioral properties. We propose a formal framework for the flexible development of object-oriented programs, which supports an interleaving of programming and verification steps. The motivation for this framework is to avoid imposing restrictions on the programming steps to facilitate the verification steps, but rather to track unresolved proof obligations and specified properties of a program which evolves. Drawing inspiration from type systems, we use an explicit proof context (or cache) to connect unresolved proof obligations and specified properties, and formulate a soundness in variant for proof contexts which is maintained by both programming and verification steps. The proof context allows a fine-grained analysis of changes to the class hierarchy. Once the set of unresolved obligations in the proof context is empty, the invariant ensures the soundness of the overall program verification.

4.6 Three cases of composition and a question

Ferruccio Damiani (University of Torino, IT)

Three examples of compositional type systems are briefly illustrated. The question is whether people working on programming languages and people workingon formal verification feel the need to identify a suite of code reuse/modularization mechanisms for synergically addressing-fine-grained code reuse- coarse-grained code reuse- spatial/temporal code evolution while being suitable for compositional analysis. Perhaps, being suitable for compositional typing could be a preliminary requisite for such a suite of mechanisms. A reformulation of the question: is it feasible for this research community to agree on a list of recommendations / guidelines / principles to be taken into account when designing a new language (or evolving an existing one) in order to facilitate formal verification?

4.7 Synthesis and Control for Infinite-State Systems with Partial Observability

Rayna Dimitrova (Universität des Saarlandes, DE)

Joint work of Dimitrova, Rayna; Finkbeiner, Bernd

The information available to a component in a distributed system is limited by its interface. Thus, in order to deliver realistic implementations, controller synthesis methods must take into account the fact that the controller has incomplete information about the global state of the system. Incomplete information is a major challenge already for finite-state systems where it makes the synthesis problem exponentially harder. For infinite-state systems the problem is in general undecidable. In particular, for real-time systems the controller synthesis problem becomes undecidable in the presence of incomplete information. We will present a novel approach to timed controller synthesis with safety requirements under incomplete information. We developed the first counterexample-guided abstraction refinement scheme that addresses the two dimensions of complexity – incomplete information and the infinite-statespace. The key innovation of our approach is the automatic synthesis of the observation predicates that are tracked by the controller. Previous methods required these predicates to be given manually. Our procedure relies on abstract counterexamples to guide the search for observations that suffice for controllability. We will outline the techniques for refining the set of observation predicates based on symbolic characterization of spurious counterexamples. We will present experimental results demonstrating better performance than approaches based on brute-force enumeration of observation sets in cases when fine granularity of the observations is necessary.

4.8 Compositional Synthesis of Distributed Systems

Bernd Finkbeiner (Universität des Saarlandes, DE)

In this talk I will illustrate how a logical representation of the synthesis problem for distributed systems facilitates a compositional synthesis approach. I will present a compositional proof rule for Extended Coordination Logic (ECL), a new temporal logic that reasons about the interplay between behavior and informedness in distributed systems. ECL extends linear-time temporal logic with quantification over strategies under incomplete information. ECL subsumes thegame-based temporal logics, including the alternating-time temporal logics, strategy logic, and game logic, and can express the synthesis problem for distributed systems with arbitrary architectures. While ECL is undecidable in general, the compositional proof rule can be used to reduce a general ECL formula to a set of formulas in the decidable fragment of ECL.

4.9 How Decomposition Enhances Security Analysis

Kathi Fisler (Worcester Polytechnic Institute, US)

Joint work of Fisler, Kathi; Krishnamurthi, Shriram

Factoring security policies out into separate program components enables interesting forms of compositional reasoning about security properties. Policies are expressive yet declarative. Both property-based verification and exhaustive semantic differencing (a property-free formal analysis) are tractable on policies. Moreover, running these analyses on policies can leave simpler residues to verify about the programs that use policies. The lesson for this seminar is that decomposition into modules in different languages can open some interesting avenues for making verification feasible and tractable.

4.10 Compositional Verification of Procedural Programs

Dilian Gurov (KTH – Stockholm, SE)

The talk gives a high-level overview of a line of work that Marieke Huisman and I started in early 2001. We develop a verification method for control-flow based temporal safety properties that uses relativization on local component properties as a means of handling variability in code: not yet available component code, evolving components, as well as components that exist in multiple variants as resulting from software product lines. We develop the theory based on flow graphs, simulation and maximal models, and explain the various difficulties in practically implementing the approach. These include a suitable choice of specification formalisms, plus algorithmic support for flow graph extraction, maximal flow graph construction, and model checking. As an example application area we show how hierarchical variability models of software product lines can be used to drive verification in a devide-and-conquer style, allowing for scalable analysis of large numbers of software products.

4.11 Abstract Symbolic Execution

Reiner Haehnle (TU Darmstadt, DE)

Modern software tends to undergo frequent requirement changes and typically is deployed in many different scenarios. This poses significant challenges to formal software verification, because it is not feasible to verify a software product from scratch after each change. It is essential to perform verificationin a modular fashion instead. The goal must be to reuse not merely software artifacts, but also specification and verification effort. In our setting code reuse is realized by delta-oriented programming, an approach where a core program is gradually transformed by code "deltas" each of which corresponds to a product feature. The delta-oriented paradigm is then extended to contract-based formal specifications and to verification proofs. As a next step towards modular verification we transpose Liskov's behavioural subtyping principle to the delta world. Finally, based on the resulting theory, we perform a syntactic analysis of contract deltas that permits to automatically factor out those parts of a verification proof that stays valid after applying a code delta. This is achieved by a novel verification paradigma called "abstract symbolic execution".

4.12 GCM/ProActive: a distributed component model, its implementation, and its formalisation

Ludovic Henrio (INRIA Sophia Antipolis – Méditerranée, FR)

The main claim of this talk is to show that software components and active objects provide an efficient programming model and verification setting. I present a component model that aims at large scale distributed systems, and is implemented as part of the ProActive library. the component model is named GCM, which stands for Grid Component Model. I present several works using formal methods to prove properties on the component model, or on some applications composed with it. I conclude with a few perspectives including verification of adaptive components, and better adaptation to multicore architectures.

4.13 VeriFast: A Powerful, Sound, Predictable, Fast Verifier for C and Java

Bart Jacobs (KU Leuven, BE)

VeriFast is a verifier for single-threaded and multithreaded C and Java programs annotated with preconditions and postconditions written in separation logic. To enable rich specifications, the programmer may define inductive datatypes, primitive recursive pure functions over these datatypes, and abstract separation logic predicates. To enable verification of these rich specifications, the programmer may write lemma functions, i.e., functions that serve only as proofs that their precondition implies their postcondition. The verifier checks that lemma functions terminate and do not have side-effects. Verification proceeds bysymbolic execution, where the heap is represented as a separation logic formula. Since neither VeriFast itself nor the underlying SMT solver do any significant search, verification time is predictable and low. We have used VeriFast to verify fine-grained concurrent data structures, unloadable kernel modules, JavaCard programs, and a network routing program embedded in a home gateway. In this talk, we demonstrate the features of VeriFast for C and Java in a 60-minute tutorial.

4.14 A Variability-Aware Module System

Christian Kaestner (Carnegie Mellon University – Pittsburgh, US)

Module systems enable a divide and conquer strategy to software development. To implement compile-time variability in software product lines, modules can be composed in different combinations. However, this way, variability dictates a dominant decomposition. As an alternative, we introduce a variability-aware module system that supports compile-time variability inside a module and its interface. So, each module can be considered a product line that can be type checked in isolation. Variability can crosscut multiple modules. The module system breaks with the antimodular tradition of a global variability model in product-line development and provides a path toward software ecosystems and product lines of product lines developed in an open fashion. We discuss the design and implementation of such a module system on a core calculus and provide an implementation for C as part of the TypeChef project. Our implementation supports variability inside modules from #ifdef preprocessor directives and variable linking at the composition level. With our implementation, we type check all configurations of all modules of the open source product line Busybox with 811 compile-time options, perform linker check of all configurations, and report found type and linker errors- without resorting to a brute-force strategy.

4.15 Compositionality for Complex Event Processing and Aspects

Shmuel Katz (Technion – Haifa, IL)

Main reference M. Goldman, E. Katz, S. Katz, "MAVEN: modular aspect verification and interference analysis," Formal Methods in System Design, vol. 37, no. 1, pp. 61–92, 2010.
URL http://dx.doi.org/10.1007/s10703-010-0101-1

Aspects can be viewed as system transformers, and then specified by assume guarantee pairs, where the assumption relates to the system to which the aspect is to be woven, and the guarantee to the resultant system after weaving. The correctness of the aspect relative to this assumption can then be shown by creating a model of the assumption's tableau with the aspect state machine model woven in, and checking whether the gurarantee holds for that model. Interference among aspects can also be defined as whether one aspect disturbs the assumption or the guarantee of another. Similar ideas can be used to specify and verify event detectors and responses for complex event processing.

4.16 Composition on the Web

Shriram Krishnamurthi (Brown University – Providence, US)

The modern Web is full of composition. Some of the most interesting (andterrifying) instances are in the browser itself. Broadly, there are two kinds of composition: between the browser and extensions, and within the page itself. My presentation briefly outlines some of these scenarios, to provide challenging examples of composition scenarios on which people can try out their theoretical ideas.

4.17 Compositional Programming

Oscar M. Nierstrasz (Universität Bern, CH)

This talk surveys the evolution of compositional paradigms in programming fromthe 1950s through to the present day. In particular, we explore the innovations introduced by procedural programming, object-oriented langauges, component-based development, and model-driven engineering.

4.18 Synthesis of Control for Component-based systems using Knowledge

Doron A. Peled (Bar-Ilan University – Ramat-Gan, IL)

In distributed systems, local controllers often need to impose global guarantees. A solution that will not impose additional synchronization may not be feasible due to the lack of ability of one process to know the current situation at another. On the other hand, a completely centralized solution will eliminate all concurrency. A good solution is usually a compromise between these extremes, where synchronization is allowed for in principle, but avoided whenever possible. In a quest for practicable solutions to the distributed control problem, one can constrain the executions of a system based on the pre-calculation of knowledge properties and allow for temporary interprocess synchronization in order to combine the knowledge needed to control the system. This type of control, however, may incur a heavy communication overhead. We introduce the use of simple supervisor processes that accumulate information about processes until sufficient knowledge is collected to allow for safe progression. We combine the knowledge approach with a game theoretic search that prevents progressing to states from which there is no way to guarantee the imposed constraints.

4.19 Compositional verification: A tutorial

Arnd Poetzsch-Heffter (TU Kaiserslautern, DE)

The tutorial starts with the presentation of basic principles of compositional verification and suggest to categorize techniques for compositional verification according to four dimensions:

- What are the components? How are they represented?
- What are the composition mechanisms?
- What properties are addressed? How are they specified?
- What kinds of verification techniques are used?

In the main part, the tutorial takes a closer look at four different settings to discuss important aspects in the huge space of compositional verification:

1. Maximal models for model checking step-synchronous Moore machines
2. Model checking control-flow properties of sequential procedural programs
3. Modular state-based verification of object-oriented programs
4. Assume-guarantee reasoning for communicating processes

(The selection of the four settings was partly done to foster the communication of the seminar participants, and partly reflects my restricted knowledge. It does not provide a fair and balanced representation of the space.)

4.20 Typical Worst-Case Analysis of Real-Time Systems

Sophie Quinton (TU Braunschweig, DE)

We present a new compositional approach providing safe quantitative information about real-time systems. Our method is based on a model to describe sporadic overload and bursts at the input of a system. We show how to derive from such a model safe quantitative information about the response time of each task. Experiments demonstrate the efficiency of this approach on a real-life example.

4.21 Component signatures of networking process components require protocol and role declarations

Johannes Reich (SAP AG – Walldorf, DE)

Although from an implementation perspective executable process components are characterized by their computational function, this is not the way they present themselves from their partners' perspective in their network like interactions. From their partners' perspective, only a projection of the interacting process components becomes visible, which is best described as their role in a protocol. Thus, to support component based design of processes, these process components require the declaration of their role and protocol in their component signature.

4.22 A logical perspective on (finite) software systems and their composition

Johannes Reich (SAP AG – Walldorf, DE)

A system notion is proposed that rests on the mathematical function notion, transforming in one time step some input and internal state values onto some output and internal state values. The question whether a certain computational entity represents a system thereby becomes the question to identify the internal and i/o states together with the system function and the corresponding time step. The effect of system interaction on system composition can be classified as:

1. Parallel processing or strict sequential interaction results in strictly hierarchical super system formation
2. Deterministic bidirectional interactions together with certain consistency conditions result in (recursive) super system formation
3. Nondeterministic bidirectional interactions together with certain consistency conditions results in provably no super system formation
4. further, non-classified relations.

4.23 Compositionality, huh?

Arend Rensink (University of Twente, NL)

In this presentation I have formulated a simple formal framework in which some of the essential aspects of compositionality come together. The most important lesson is that the composition operator should be compatible with a notion of abstraction, or semantics, encapsulating the conceptual understanding of the objects under construction. If this compatibility fails, the common solution is to augment the abstraction by putting more information into it. This kind of augmentation is, in fact, the same kind of step as strengthening the induction hypothesis in an inductive proof. The framework is tested, with varying degrees of success, against a number of cases of composition from different domains, ranging from bisimilarity minimisation to testing and subclassing (seen as a composition operator). Some of these cases nicely illustrate the principle of augmentation; others provide examples where compositionality simply fails to hold.

4.24 Comparing Verification Condition Generation with Symbolic Execution

Malte Schwerhoff (ETH Zürich, CH)

Joint work of Schwerhoff, Malte; Kassios, Ioannis T.; Müller, Peter

There are two dominant approaches for the construction of automatic program verifiers, Verification Condition Generation (VCG) and Symbolic Execution (SE). Both techniques have been used to develop powerful program verifiers. However, to the best of our knowledge, no systematic experiment has been conducted to compare them – until now. We have used the specification and programming language Chalice and compared the performance of its standard VCG verifier with a newer SE engine called Syxc, using the Chalice test suite as a benchmark. We have focused on comparing the efficiency of the two approaches, choosing suitable metrics for that purpose. Our metrics also suggest conclusions about the predictability of the performance. Our results show that verification via SE is roughly twice as fast as via VCG. It requires only a small fraction of the quantifier instantiations that are performed in the VCG-based verification.

4.25 Compositional Verification of Actors

Marjan Sirjani (Reykjavik University, IS)

Rebeca is designed as an imperative actor-based language with the goal of providing an easy to use language for modeling concurrent and distributed systems, with formal verification support. I will explain the language Rebeca and the supporting model checking tools. Abstraction and compositional verification, and state-based reduction techniques including symmetry reduction of Rebeca will be discussed. As an example of a domain specific example, I will show how we used Rebeca for model checking SystemC codes.

4.26 Risk Management meets model checking: compositional analysis of DFTs

Marielle Stoelinga (University of Twente, NL)

Dynamic fault trees (DFTs) are a versatile and common formalism to model and analyze the reliability of computer-based systems. This talk presents a formal semantics of DFTs in terms of input/output interactive Markov chains (I/O-IMCs), which extend continuous-time Markov chains with discrete input, output and internal actions. This semantics provides a rigorous basis for the analysis of DFTs. Our semantics is fully compositional, that is, the semantics of a DFT is expressed in terms of the semantics of its elements (i.e. basic events and gates). This enables an efficient analysis of DFTs through compositional aggregation, which helps to alleviate the state-space explosion problem, and yields a very flexible modeling and analysis framework by incrementally building the DFT state space. We have implemented our methodology by developing a tool, and showed, through a number of case studies, the feasibility of our approach and its effectiveness in reducing the state space.

4.27 Modularity and compositionality in embedded system design: interface synthesis and interface theories

Stavros Tripakis (University of California – Berkeley, US)

Compositional methods, that allow to assemble smaller components into larger systems both efficiently and correctly, are not simply a desirable feature in system design: they are a must for designing large and complex systems. In this talk I will present some recent work on this general theme, motivated by embedded system applications. A key notion is that of "interface" which allows to abstract a component, hiding details while exposing relevant information. I will present methods for automatic bottom-up synthesis of interfaces for hierarchical synchronous and dataflow models, motivated by the need for modular code generation from such models. I will also present two interface theories for the same models. Interface theories can be seen as behavioral type theories. They include the key notion of refinement, which captures substitutability: when can a component be replaced by another one without compromising the properties of the entire system. I will present two interface theories: – synchronous relational interfaces, targeted at synchronous systems and functional properties; – actor interfaces, targeted at dataflow models and performance properties such as throughput or latency.

4.28 Interface Theories: Design Choices

Stavros Tripakis (University of California – Berkeley, US)

Interface theories such as interface automata were introduced by Alfaro and Henzinger in the early 2000s. A key characteristic of these theories is the "asymmetry" of inputs and outputs, and the fact that interfaces are not "input-enabled": they may declare some inputs as illegal. This results in a "demonic" notion of composition and an "alternating" notion of refinement. This talk discusses the design choices that naturally lead to these definitions.

4.29 Compositional Behavioral Modeling with Real Time Games

Andrzej Wasowski (IT University of Copenhagen, DK)

Joint work of Wasowski, Andrzej; Larsen, Kim; Legay, Axel; Nyman, Ulrik; David, Alexandre

The presentation introduces the concept of specification theory (or interface)theory for automata-like models, including the main motivation and ingredient operators. I show how we instantiated this paradigm in the tool ECDAR, which aims at stepwise design and verification of real-time embedded controllers. The tool is based on the semantic model of Timed Games, and its associated symbolic solving algorithms. I will present the main objects of ECDAR's specification theory (implementations, specifications, properties), its transformation operators (conjunction, parallel composition, quotient) and its verification operator (satisfaction, refinement). I focus mostly on examples of structuring models and correctness proofs. Time permitting, I will show patterns that allow Assume/Guarantee style of verification, and combine them to obtain (finite) inductive proofs of correctness using refinements in ECDAR.

5 Working Groups

5.1 Working Group: How can model checking and deductive verification benefit from each other/verification of very large systems

Huisman, Marieke; Sirjani, Marjan

A group of about 15 people with very different backgrounds attended this working group. Initially, the discussion went in many different directions, and the point was raised that work on this had been attempted already 20 years ago. However, it some point it was realised that much of this unclearity was caused by the term deductive verification, which can have many different meanings. Thus it was necessary to clarify the difference between theorem proving and program verification. Theorem proving means that one writes some logical properties in a formal language and uses the theorem prover to verify the properties. Two kinds of theorem provers exist: interactive theorem provers, typically used for proving properties in higher-order logic, where the user has to guide the verification process, and automated theorem provers, typically used for first-order logic, where an algorithm tries to construct a proof. Some well-known examples of interactive theorem provers are: HOL, PVS, Isabelle and Coq. Some well-known examples of automated theorem provers are all STMLib-compliants tools, such as Vampire and Z3. Program verification is really focused on the verification of annotated software. Their formal foundations are for example Hoare logic, weakest precondition calculus or separation logic. Typically, they use verification condition generation or symbolic execution to generate proof obligations in first-order logic. To verify these proof obligations, automated theorem proving is used. Then the discussion continued about how program verification and model checking could benefit from each other. It was observed that model checking is currently moving from concurrent models to concurrent software, whereas program verification is moving from sequential software to concurrent software. Thus, the verification targets seem to meet. However, program verification typically focuses on data-oriented properties, whereas model checking focuses on control-flow-oriented properties. Finally, we identified several interesting examples. Program verification could benefit from model checking in the following examples:

- a distributed computation: a server sends a task to a worker. If the task is too big, the worker spawns a new thread to perform part of the task (and this is repeat until the task is small enough). Eventually all the results from the task are merged. This control pattern is difficult to capture with classical program verification techniques;
- verification of e-voting software: typically security properties such as absence of information leakage depend on control and data;
- counter example generation. Model checking could also benefit from program verification, for example in the following case: – control flow properties of code with complicated data structures: suppose we have a device driver that uses a balanced search tree as internal data structure.

The tree would have a method to compute its size. Program verification can be used to prove that the size does not change when for example rebalancing the tree, and the information that the size of the tree is not changed between two consecutive calls can then be used by the model checker to reduce the state space. A next step to continue the results from this workgroup would be to concretely start working on these examples.

5.2 Working Group: Compositional Synthesis of Reactive Systems

Barbara Jobstmann

The working group consisted of four people, all with a similar background in reactive synthesis and automata-based game theory. Therefore, there was no need to discuss the precise meaning of the problem and we could dive right into the technical details about the subject. We brainstormed on the different approaches we knew and explained to each other their key aspects. More precisely, Marielle mentioned recent work on "Compositional Synthesis of Safety Controllers" by her student Wouter Kuijper. We discuss the benefits of using safety specifications and the difficulties arising with more general specifications. Barbara mentioned that similar issues also arise in a problem of aiming to synthesis missing environment assumptions. Here the idea is the given an unrealizable specification we aim for a minimal assumption that makes the specification realizable. Finally, Ufuk discussed his recent work on semi-compositional version of the generalized reactivity-1 approach. Barbara and Rayna provided feedback and suggested extensions. Over all the discussion was very pleased, focused, and technical and more than worth the short amount of time we spent on it.

5.3 Working Group: Modular Full Functional Specification and Verification of C and Java programs that Perform I/O

Bart Jacobs

We discussed preliminary ideas on how to verify I/O properties of C and Java programs. For concreteness, we used the VeriFast program verification tool to illustrate the ideas through simple example programs. We started with a simple C program that prints "Hi" using two calls of the 'putchar' function. Our first specification for this program used an abstract separation logic predicate 'world' with an single parameter of type 'list of character' that represents the characters that have been printed. The precondition stated 'world([])' and the postcondition stated 'world(['H', 'i'])'. This specification approach works fine if we also check termination. Otherwise, the program can circumvent it by first printing e.g. "Bye" and then going into an infinite loop, so that the postcondition is never checked. The second specification used the same 'world' predicate, but now the parameter denotes the characters that are yet to be printed by the program. The precondition states 'world(['H', 'i'])' and the postcondition states 'world([])'. This specification properly enforces the safety property that the program only ever prints a prefix of "Hi", even if it does not terminate. However, this specification approach cannot accommodate nondeterministic behavior. Therefore, we next came up with a predicate 'world' with two parameters: one of type 'list of character', denoting the characters that have been printed, and one of type 'set of lists of characters', denoting the set of allowed values of the first argument. The precondition of 'putchar' requires that the old value of the list, with the new character appended to the end, is in the set. An alternative approach is to have a single argument of type 'iospec', where 'iospec' is coinductively defined as 'function from I/O actions to I/O results' and 'I/O result' is defined as either 'Not Allowed' or 'Allowed (iospec)' which specifies the new I/O specification that holds after the action is performed. We also discussed if it would be possible (and a good idea) to build on this approach to also check liveness properties, by adding 'tau' actions (a.k.a.stutter steps or silent actions). We remarked that it would not in general be easy to modularly prove conformance of a program against such a specification. We noted that it was important to record all actions in a single stream; otherwise, it would not be possible to specify an ordering between the various types of actions. Conclusions: We have some ideas on how to specify and verify I/O properties of Cand Java programs. However, we need to check their feasibility by trying them out on larger examples. We anticipate that additional mechanisms will be necessary to achieve a truly modular approach.

Participants: Erik Poll, Dilian Gurov (first half), Einar Broch Jonsen, Malte Schwerhoff, Wolfgang Ahrendt, Bart Jacobs .

5.4 Working Group: Benchmark for Industrial Verification/Synthesis Problem

Johannes Reich

The interest in the working group o was pretty high. Johannes Reich gave a short introduction (see slides) about potentially interesting problems from an industry perspective. Then we discussed the matter along the following lines: There was general agreement that a benchmark would be an appropriate tool to gain attraction and attention for the advanced engineering methods commonly termed "formal methods" The main issue is to find a showcase that is immediately relevant for the industry like SAP and which provides provides a scalable and significant problem. One example would be the RSA-challenge of prime number factorization, with its immediate relevance for asymmetric encryption techniques. The examples of the introduction were seen as too unspecific. In general, the area of formal methods does not come with _the_ problem but instead shows quite some heterogeneity. In some sense there seems to be a henn-and-egg problem to bring together the more technical view of the formal methods experts and the business view of the industry. An other approach could be that a company like SAP provides an example application where the power of several formal methods like verification, synthesis, etc. could be demonstrated – possibly to arrive at a show case for a benchmark. Interesting candidates for SAP could be the tax engine or the pricing engine, where correctness is top priority and errors might get a high visibility. Another area could be product security or areas where high testing efforts could be ameliorated by formal verification. Thanks for the feedback and best regards to all participants!

Further sources and competitions

Several other competitions and one article was mentioned by the participants:

References

1 M Huisman, V Klebanov and R Monahan (2011) On the Organisation of Program Verification Competitions, http://ceur-ws.org/Vol-873/papers/paper_2.pdf
2 Transformation Tool Competition: http://planet-sl.org/ttc2011 for the last edition, there is a next one coming up for 2013
3 Knowledge Engineering for Planning and Scheduling: http://icaps12.poli.usp.br/icaps12/ickeps
4 Microsoft Research: 2012 Verified Software Milestone Award, http://dream.inf.ed.ac.uk/vsi/

5.5 Working Group: Compositional Verification of Software Product Lines

Ina Schaefer

A software product line is a set of software systems developed from a common set of core assets. There exist several implementation techniques for product lines, such as annotative, compositional or transformational approach, and different analysis strategies which apply to all possible verification approaches. The product-based analysis strategy generates each possible product and verifies it in isolation. The advantage of this strategy is that single-product analyses can be applied without any change, but for large product lines this approach is inefficient or even infeasible. The family-based approach generates a meta-representation of all products which can be analyzed at once in order to derive properties of all products. While this is more efficient than the product-based approach, it uses a closed-world assumption making it fragile to product line evolution. The feature-based analysis strategy takes acompositional approach by analyzing the building blocks of the productsseparately. However, in most cases, this analysis step is not sufficient because it does not take the dependencies and interactions between the building blocks into account. Hence, in addition to the feature-based analysis step another product-based or family-based analysis step is required. In the working group, we discussed the potential of the feature-based compositional analysis strategy for product lines and its limitations. In general, we came to the conclusion that the applicability of this strategy depends on the considered use case. In particular, to assess the above question one needs to fix the implementation approach for the products, the properties that should be verified, the specification technique for the product line, and the analysis technique. It seems that feature-based analysis is in particular well suited to model-checking where traditional assume-guarantee reasoning is adapted as follows: If a feature satisfies its specification under the assumption that it is added to a product with a particular required property, then the resulting composed product guarantees the specification of the feature if there are no interactions between the already contained features and the newly composed one.

6 Programme day-by-day

Monday December 17

9:00	Welcome from the organisers
9:05	Tutorial on compositional programming Oscar Nierstrasz
10:35	Coffee
11:00	Personal introductions
12:15	Lunch
13:45	Tutorial on compositional verification Arnd Poetzsch-Heffter
15:15	Personal introductions
15:30	Coffee
16:00	Tutorial on compositional modelling Arend Rensink
17:30	Remaining personal introductions

Tuesday December 18

9:00	Einar Broch Johnsen Shriram Krishnamurthi Compositionality for web services
10:30	Coffee
11:00	Doron Peled Christian Kaestner A Variability-Aware Module System Wolfgang Ahrendt Compositional semantics and of concurrent/distributed objects
12:15	Lunch
13:45	Sophie Quinton Typical worst-case analysis of real-time systems
14:45	Working groups
17:00	Shmuel Katz Compositionality for Complex Event Processing and Aspects Bernd Finkbeiner Compositional synthesis of distributed systems

Wednesday December 19

9:00	Stavros Tripakis Compositionality and modularity: interfaces everywhere! Wolfgang Ahrendt Compositional verification of concurrent/distributed objects Johannes Reich A logical perspective on software systems and their composition
10:30	coffee
11:00	Ferruccio Damiani Dynamic delta-oriented software product lines Dilian Gurov Compositional Verification of Programs with Procedures Stavros Tripakis design choices for interface theories
12:15	lunch
14:00 14:30	photo excursion

Thursday December 20

9:00	Bart Jacobs VeriFast: A Powerful, Sound, Predictable, Fast Verifier for C and Java
10:00	Reiner Haehnle Abstract Symbolic Execution
10:30	coffee
11:00	Rayna Dimitrova Synthesis and Control for Infinite-State Systems with Partial Observability Malte Schwerhoff Comparing Verification Condition Generation with Symbolic Execution Johannes Reich Component signatures of networking process components require protocol and role declarations
12:15	Lunch
13:45	Marjan Sirjani Compositional Verification of Actors (30 min) Simon Bliudze Glue Synthesis in BIP (30 min.)
14:45	working groups
17:00	Ludovic Henrio GCM/ProActive: a distributed component model, its implementation, and its formalisation

Friday December 21

9:15	Mariëlle Stoelinga
9:30	Andrzej Wasowski Compositional design with stepwise refinement using real time games
10:00	Christian Kaestner
10:30	Coffee
11:00	Reports from working groups and plenary discussion
12:00	Closing of the seminar
12:15	Lunch

Participants

- Wolfgang Ahrendt
Chalmers UT – Göteborg, SE
- Simon Bliudze
EPFL – Lausanne, CH
- Einar Broch Johnsen
University of Oslo, NO
- Ferruccio Damiani
University of Torino, IT
- Rayna Dimitrova
Universität des Saarlandes, DE
- Christian Eisentraut
Universität des Saarlandes, DE
- Bernd Finkbeiner
Universität des Saarlandes, DE
- Kathi Fisler
Worcester Polytechnic Inst., US
- Susanne Graf
VERIMAG – Gières, FR
- Dilian Gurov
KTH – Stockholm, SE
- Reiner Hähnle
TU Darmstadt, DE
- Ludovic Henrio
INRIA Sophia Antipolis – Méditerranée, FR
- Marieke Huisman
University of Twente, NL
- Bart Jacobs
KU Leuven, BE
- Barbara Jobstmann
VERIMAG – Gières, FR
- Christian Kästner
Carnegie Mellon University – Pittsburgh, US
- Shmuel Katz
Technion – Haifa, IL
- Shriram Krishnamurthi
Brown Univ. – Providence, US
- Malte Lochau
TU Darmstadt, DE
- Oscar M. Nierstrasz
Universität Bern, CH
- Doron A. Peled
Bar-Ilan University – Ramat-Gan, IL
- Arnd Poetzsch-Heffter
TU Kaiserslautern, DE
- Erik Poll
Radboud Univ. Nijmegen, NL
- Sophie Quinton
TU Braunschweig, DE
- Johannes Reich
SAP AG – Walldorf, DE
- Arend Rensink
University of Twente, NL
- Ina Schaefer
TU Braunschweig, DE
- Malte Schwerhoff
ETH Zürich, CH
- Vasiliki Sfyrla
VISEO – Lyon, FR
- Marjan Sirjani
Reykjavik University, IS
- Lei Song
Universität des Saarlandes, DE
- Martin Steffen
University of Oslo, NO
- Marielle Stoelinga
University of Twente, NL
- Ufuk Topcu
University of Pennsylvania, US
- Stavros Tripakis
University of California – Berkeley, US
- Andrzej Wasowski
IT Univ. of Copenhagen, DK

Report from Dagstuhl Seminar 12512

Representation, Analysis and Visualization of Moving Objects

Edited by
Joachim Gudmundsson[1], Patrick Laube[2], and Emiel Van Loon[3]

1 **The University of Sydney, AU,** joachim.gudmundsson@sydney.edu.au
2 **Universität Zürich, CH,** patrick.laube@geo.uzh.ch
3 **University of Amsterdam, NL,** vanloon@uva.nl

Abstract

From December 16 to December 21, 2012, the Dagstuhl Seminar 12512 "Representation, Analysis and Visualization of Moving Objects" was held in Schloss Dagstuhl – Leibniz Center for Informatics. The major goal of this seminar was to bring together the diverse and fast growing, research community that is involved in developing better computational techniques for spatio-temporal object representation, data mining, and visualization of moving object data. The participants included experts from fields such as computational geometry, data mining, visual analytics, GIS science, urban planning and movement ecology. Most of the participants came from academic institutions but some also from government agencies and industry. The seminar has led to a fruitful exchange of ideas between different disciplines, to the creation of new interdisciplinary collaborations and to recommendations for future research directions. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper.

Seminar 16.–21. December, 2012 – www.dagstuhl.de/12512
1998 ACM Subject Classification E.1 Data Structures, F.2 Analysis of Algorithms and Problem Complexity, H.2.8 Database Applications
Keywords and phrases movement analysis, algorithms, visualization, geo-ecology, urban planning
Digital Object Identifier 10.4230/DagRep.2.12.89
Edited in cooperation with Georgina Wilcox

1 Executive Summary

Joachim Gudmundsson
Patrick Laube
Emiel Van Loon

The seminar brought together researchers and domain experts involved in developing and utilizing methods for knowledge extraction from massive amounts of data from moving objects. This knowledge is essential to substantiate decision making in public and private sectors, in application domains such as transportation modelling, urban planning, tourism, wildlife ecology, spatial epidemiology, location-based services, flight safety, and marine safety. Moving object data typically include trajectories of discrete spatial objects (e.g. humans, vehicles, animals, and goods), continuous phenomena (e.g. storms, ocean currents) as well as trajectories of abstract concepts (e.g. information flow, moving data points in attribute space) or even vectors of spreading diseases. Technologies for object tracking are low cost and increasingly reliable in terms of coverage and accuracy, hence movement records are nowadays generated in huge volumes on a routine basis, using diverse technologies such as radio frequency mapping, Global Navigation Satellite Systems, video sequences and Doppler radar.

The computational analysis of movement data has seen a successful first decade with progress made in capturing, preprocessing, storing, indexing and querying movement data, combined with promising results in visualizing movement and detecting movement patterns. However, whereas such basic progress in handling movement data was needed for establishing a new field and attracting funding, attention must now move on towards the extraction of useful information and process knowledge from tracking data.

In many application fields the need for analysing large sets of trajectories is evident and crucial; however, only very rudimentary automated analysis tools are available and anything more advanced is currently analyzed manually. As an example there are several companies tracking the movement of football players through video with a frequency of at least 25Hz and accuracy of approximately 10cm. But most of the analysis and the annotation (passes, throw-ins, goals etc) of the data are still made manually. Thus the analysis part has been neglected and in comparison with the image processing part it is technologically far behind in the development. A reason for this is the obvious lack of theoretical and practical solutions for many crucial fundamental problems.

For that reason, this seminar focussed on formalizing methods for algorithmic analysis, visual analytics, data mining and knowledge discovery, defined by a multidisciplinary team of researchers and practitioners.

Participants and Format

The seminar brought together researchers from several disciplines involved in developing and utilizing computational techniques for spatiotemporal object representation, data mining, and visualization. This community encompasses an interdisciplinary mix of methodologically oriented as well as application oriented researchers. The methods-oriented researchers are from fields such as theoretical computer science, spatial databases, knowledge discovery and data mining, visual analytics, and geographic information science. They were complemented by application scientists from various fields, especially behavioral ecology and urban planning.

Drawing upon positive experiences in previous seminars of this series, oral presentation sessions were complemented by special sessions dedicated to open research questions and project ideas, as well as to discussions in small, concurrent break-out groups focussing on a specific domain or case studies. Since the participants of the seminar came from quite different backgrounds, concise survey talks on the first two days were delivered on movement ecology (Ran Nathan), industry movement research (Ben Loke) and spatial data mining (Győző Gidófalvi).

A data challenge was organized for the participants prior to the seminar. This ensured that the participants were well aware of the application domain which was the focus of the seminar and it gave the domain experts a possibility to see the potential use of various different approaches. The data challenge included a bird migration data set provided by Emiel van Loon.

Many interesting research results were presented, demonstrating the progress in this field. The participants were highly satisfied with the quality of the seminar; especially the involvement of domain specialists from relevant application domains was highly appreciated.

2 Table of Contents

3 Overview of Talks

3.1 Towards an algorithmic framework for movement analysis with movement models

Kevin Buchin (TU Eindhoven, NL)

Joint work of Buchin, Kevin; Sijben, Stef; Arseneau, T. Jean Marie; Willems, Erik P.

In trajectory data a low sampling rate leads to high uncertainty in between sampling points, which needs to be taken into account in the analysis of such data. However, current algorithms for movement analysis ignore this uncertainty and assume linear movement between sample points. In this talk I present a framework for movement analysis using the Brownian bridge movement model (BBMM), that is, a model that assumes random movement between sample points. Many movement patterns are composed from basic building blocks like distance, speed or direction. I will show how to express their distribution over space and time in the BBMM and will demonstrate our framework by computing a "following" pattern. Our motivation to study the BBMM stems from the rapidly expanding research paradigm of movement ecology. To this end, I will present an application of the framework to the simultaneous movement of groups of wild and free-ranging primates.

3.2 Segmenting geese tracks by movement states

Maike Buchin (TU Eindhoven, NL)

Joint work of Buchin, Maike; Kruckenberg, Helmut; Kölzsch, Andrea
Main reference M. Buchin, H. Kruckenberg, A. Kölzsch, "Segmenting Trajectories based on Movement States," in Proc. of the 15th Int'l Symp. on Spatial Data Handling (SDH), 2012.

We report on work that originated during the Lorentz Center 2011 workshop. Dividing movement trajectories according to different movement states of animals has become a challenge in movement ecology, as well as in algorithm development. In this study, we revisit and extend a framework for trajectory segmentation based on spatio-temporal criteria for this purpose. We adapt and extend the framework for segmenting by movement states of an individual. We implement the framework and evaluate it using the example of two spring migration tracks of white-fronted geese. Manual and automatic classifications are compared and extensions discussed.

3.3 Movement in cordon structured spaces

Matt Duckham (The University of Melbourne, AU)

Joint work of Duckham, Matt; Bleisch, Susanne; Both, Alan; Lyon, Jarod; Laube, Patrick; Wark, Tim
Main reference A. Both, M. Duckham, P. Laube, T. Wark, J. Yeoman, "Decentralized monitoring of moving objects in a transportation network augmented with checkpoints," The Computer Journal, in press.

Traditional definitions of a "trajectory" (e.g., "set of n moving points whose locations are known as t consecutive time steps,") conjure up movement in space through time: the position of a moving object recorded at successive times [1]. However, a complementary view of movement is to record the time at which an object moves past known checkpoints, called cordons. These cordon-structured spaces and movements are common in a wide range of applications, from e-tolling and public transport smart cards; through resource tracking for emergency response; to environmental monitoring of tagged animal movements. Trajectory and cordon-structured movements are to some extent inter-convertible: both are representations of the same underlying phenomena. However, two important differences exist. First, the often lower, but more importantly variable spatial and temporal granularity of motion tracking in cordon-structured spaces makes many trajectory analysis tools not applicable to cordon structure movement data (such as path shape and movement velocity derivation). Second, the location of cordons is typically salient in the application. Cordons are placed in locations that provide a meaningful qualitative structure to the space. Unlike trajectories, observations of spatial location in cordon structured movement are not arbitrary.

References

1 Gudmundsson, J., van Kreveld, M., and Speckmann, B. Efficient detection of patterns in 2D trajectories of moving points. *GeoInformatica* 11, 195–215, 2007
2 Both, A., Duckham, M., Laube, P., Wark, T., Yeoman, J. (2013) Decentralized monitoring of moving objects in a transportation network augmented with checkpoints. *The Computer Journal*, in press.

3.4 Brief survey of trajectory prediction

Győző Gidófalvi (KTH – Stockholm, SE)

Joint work of Győző Gidófalvi; Fang Dong
Main reference G. Gidófalvi, F. Dong, "When and Where Next: Individual Mobility Prediction," in Proc. of MobiGIS, 2012.

Over the past decade, several technological and social trends enabled the large scale collection of moving object trajectories. The potential utility of likely regularities in these trajectories has spawned a new sub-field in data mining: trajectory data mining. After a brief review of frequent sequential pattern mining, the talk presents a brief survey and taxonomy of trajectory prediction—an important trajectory mining task—and illustrates how different approaches perform trajectory prediction at different spatio-temporal scales and under different application settings [1, 2, 3].

References

1 G. Gidófalvi, M. Kaul, C. Borgelt and T. Bach Pedersen. *Frequent Route Based Continuous Moving Object Location- and Density Prediction on Road Networks*. In Proc. of ACM-GIS, pp. 381-384, 2011.

2 G. Gidófalvi and F. Dong. *When and Where Next: Individual Mobility Prediction*. In Proc. of MobiGIS, 2012.

3 A. Monreale, F. Pinelli, R. Trasarti and F. Giannotti. *WhereNext: a Location Predictor on Trajectory Pattern Mining*. In Proc. of KDD, pp. 637-645, 2009.

3.5 The classification of passes in football

Joachim Gudmundsson (University of Sydney, AU)

Joint work of Estephan, Joël; Gudmundsson, Joachim

Studies have shown that coaches are less than 45% correct in their post-game assessment of what occurs during the 90 minutes of a football match [1]. As a consequence there has recently been a strong research interest in automatically analyst the performance of a player or a team. The time and money lately being exhausted on sports analysis only begins to exemplify its importance to the sporting community.

In this talk we explore the automatic classification of the quality of passes in team sports. We first generated an extensive data set of football passes and had them manually classified by experts. Models were then built using a variety of different feature selection methods and supervised learning algorithms. Four models were selected, two classifying the risk of a pass and the others the worth of a pass. The results were quite compelling/positive/encouraging with the accuracy rates ranging from 71% to 96%.

The developed models have been added to the ever growing set of tools that provide coaches and players with an in-depth analysis of the game.

References

1 I.M. Franks and G Miller. *Eyewitness testimony in sport*. Journal of Sport Behaviour, 9:39–45, 1986.

3.6 Analyzing and visualizing movement data with head/tail breaks

Bin Jiang (University of Gävle, SE)

Main reference B. Jiang, "Head/tail Breaks: A New Classification Scheme for Data with a Heavy-tailed Distribution," arXiv:1209.2801 [physics.data-an].
URL http://arxiv.org/abs/1209.2801

This paper summarizes a novel method for efficiently and effectively analyzing and visualizing movement data. The novel method, called head/tail breaks, was initially developed as a classification method for data with a heavy tailed distribution. Since most movement data exhibit heavy tailed distributions, head/tail breaks can be an important means for data analysis and visualization.

3.7 Trajectory grouping structure

Marc van Kreveld (Utrecht University, NL)

Joint work of Buchin, Kevin; Buchin, Maike; van Kreveld, Marc; Speckmann, Bettina; Staals, Frank

The collective motion of a set of moving entities such as a set of people, birds, or other animals, is characterized by groups arising, merging, splitting, and ending. Given the trajectories of these entities, we define and model a structure that captures all of such changes using the Reeb graph, a concept from topology. The trajectory grouping structure has three natural parameters that allow more global views of the data in group size, group duration, and entity inter-distance. We prove complexity bounds on the maximum number of maximal groups that can be present, and give algorithms to compute the grouping structure efficiently. We also study how the trajectory grouping structure can be made robust, that is, how brief interruptions of groups can be disregarded in the global structure, adding a notion of persistence to the structure. Furthermore, we showcase the results of experiments using data generated by the NetLogo flocking model and from the Starkey project. The Starkey data describes the movement of elk, deer, and cattle. Although there is no ground truth, the experiments show that the trajectory grouping structure is plausible and has the desired effects when changing the essential parameters. Our research provides the first complete study of trajectory group evolvement, including combinatorial, algorithmic, and experimental results.

3.8 Noldus InnovationWorks moving objects research themes

Ben Loke (Noldus Information Technology BV, NL)

In this talk we present an overview of the current Noldus InnovationWorks research themes.

Movement tracking: the AnyTrack project, a tool that should track any number of objects in any context with any type of sensor. With examples of use for the wild-life tracking project E-Track.

Automatic behavior recognition: from behavior recognition for rodents with a video based tracking tool, EthoVision XT, to people tracking, and gesture, pose and hand tracking techniques.

System integration: trend to analyze multi-sensor data from different systems into one integration tool The Observer XT for detailed synchronized analysis.

3.9 M-Atlas: A platform for mobility knowledge discovery

Mirco Nanni (ISTI-CNR – Pisa, IT)

Joint work of Giannotti, Fosca; Nanni, Mirco; Pedreschi, Dino; Pinelli, Fabio; Renso, Chiara; Rinzivillo, Salvatore; Trasarti Roberto
Main reference F. Giannotti, M. Nanni, D. Pedreschi, F. Pinelli, C. Renso, S. Rinzivillo, R. Trasarti, "Unveiling the complexity of human mobility by querying and mining massive trajectory data," in VLDB Journal (VLDBJ), Special Issue on "Data Management for Mobile Services", Vol. 20, Issue 5, pp. 695–719, 2011.
URL http://dx.doi.org/10.1007/s00778-011-0244-8

Real applications involving the analysis of moving objects require a complex combination of tasks of several different kinds, including data mining (or other analysis) tools. In this talk I presented an analytics platform for supporting such kind of process, thought for — but not limited to — vehicle traffic trajectories, highlighting both the basic analytics tools provided by the system and its ability to build derived analysis functionalities and whole knowledge extraction processes.

3.10 Movement ecology

Ran Nathan (The Hebrew University of Jerusalem, IL)

Main reference R. Nathan, W.M. Getz, E. Revilla, M. Holyoak, R. Kadmon, D. Saltz, P.E. Smouse, "A movement ecology paradigm for unifying organismal movement research," in Proc. of the National Academy of Sciences of the United States of America, Vol. 105, No. 49, pp. 19052–19059, 2008.
URL http://dx.doi.org/10.1073/pnas.0800375105

Understanding and predicting the dynamics of complex ecological and human-dominated systems are best accomplished through the synthesis and integration of information across relevant spatial, temporal and thematic scales. Recent advances in mechanistic modeling, data analysis techniques and tracking technology have enriched our capacity to disentangle the key parameters affecting dispersal, migration, foraging and other movement processes and to accurately quantify movement patterns. In lieu of this favorable background, movement ecology has recently emerged to facilitate the unification of movement research. Movement ecology aims at investigating the explicit links between the internal state, the motion and the navigation capacities of the individual, and the external (biotic and abiotic) environmental factors affecting its movement. Therefore, it provides a natural platform for examining the mechanisms underlying movement processes and patterns, their causes and consequences in changing environments. In this talk I presented the vision and basic principles of the movement ecology approach [1], and illustrated its application to the study of foraging of vultures [2] and navigation of fruit bats [3].

References

1 Nathan, R., W. M. Getz, E. Revilla, M. Holyoak, R. Kadmon, D. Saltz and P. E. Smouse. 2008. *A movement ecology paradigm for unifying organismal movement research*. Proceedings of the National Academy of Sciences of the United States of America, 105: 19052–19059.
2 Nathan, R., O. Spiegel, S. Fortmann-Roe, R. Harel, M. Wikelski and W. M. Getz. 2012. *Using tri-axial acceleration data to identify behavioral modes of free-ranging animals: Gen-*

eral concepts and tools illustrated for griffon vultures. Journal of Experimental Biology, 215: 986–996.

3 Tsoar, A., R. Nathan, Y. Bartan, A. Vyssotski, G. Dell'Omo and N. Ulanovsky. 2011. *Large-scale navigational map in a mammal.* Proceedings of the National Academy of Sciences of the United States of America, 108: E718–E724.

3.11 Workshops, data challenges and some thoughts

Ross Purves (University of Zurich)

In my talk I presented my take on the process of moving from our discussions in 2010 at Dagstuhl, to a data challenge on gull data at the Leiden Centre in 2011 and a workshop on Progress in Analysis of Real Movement Data in autumn 2012. I emphasized the utility of challenges in bringing us together to discuss common themes, even if direct results in terms of publications are less obvious. However, links made have resulted in new publications and it is clear that data challenges have great utility. I concluded by suggesting that data challenges should meet three criteria:

1. Be realistic (i.e. use real data and ask real questions)
2. Go beyond simple primitives
3. Allow us to compare high level methods (for example analysis not query/ derivation)

3.12 Semantic enrichment of trajectory data

Chiara Renso (ISTI-CNR – Pisa, IT)

Joint work of Renso, Chiara; Vania Bogorny
Main reference V. Bogorny, C. Renso, A. Ribeiro de Aquino, F. de Lucca Siqueira, L. Otavio Alvares, "CONSTAnT – A Conceptual Data Model for Semantic Trajectories of Moving Objects," in Transaction in GIS, 2013, to appear.

Mobility of individuals is a complex phenomena. The last few years has seen a growing interest in representing analysing and understanding the mobility of individuals, basically relying on the concept of a trajectory to represent the history of evolving locations. However, the raw trajectory representation as a sequence of georeferenced points and the paired analysis techniques so far do not produce satisfying results in terms of understanding and proper deployment for the final application. The proposal of representing trajectories in a more meaningful way, namely, semantic trajectories, is a step towards this ambitious goal. We propose a conceptual model called CONSTAnT for semantic trajectory representation that enhances the well known "stop and move" including many semantic facets that should be combined with raw trajectories. Analysis methods aimed at (1) inferring semantics from data to annotate trajectories and (2) finding patterns from semantically enriched trajectories have therefore to be investigated.

3.13 Alpha-visibility

Jörg-Rüdiger Sack (Carleton University, CA)

Joint work of Ghodsi, Mohammad; Maheshwari, Anil; Nouri-Baygi, Mostafa; Zarrabi-Zadeh, Hamid

We study a new class of visibility problems based on the notion of α-visibility. A segment t is said to be α-visible from a point p, if p can see t with an angle at least α; that is, if there exists an empty triangle with one vertex at p and side opposite to p on t such that the angle at p equals α. This notion of visibility appears to be natural and has the potential to be of practical value. For example, the smallest angle we can observe directly is determined by the ratio of the wavelength of light to the diameter of the eye's pupil, which is lower bounded by a constant 10^{-4} radians. Also, some optical/digital imaging devices have similar limitations, quantified by their resolutions.

Our α-visibility model is inspired by those limitations, and may provide a more realistic alternative to the classical visibility models studied in the literature. In this talk, we present some of the first results for several variants of the point/segment visibility problems in the α-visibility model. We discuss open problems and anticipate that this model will be extended in further studies.

3.14 Towards an interdisciplinary analysis platform: the benefits of R as an open source analysis framework

Kamran Safi (MPI für Ornithologie, DE)

As ecology is heading towards BigData with rapid speed, computational ecology has emerged as an interdisciplinary discipline. But, besides the problems of storage and versioning that large data sets bring, the issue of efficient data analysis and progress in methodology needs addressing. Here, I present MOVE (http://cran.r-project.org/web/packages/move/index.html) an R library that provides an interface with Movebank, allowing simple data access from R. This library should allow the future to make advances in development of methods in movement analysis more straight forward. Finally, I presented an outlook of ICARUS an initiative to track animals from space beginning in 2014 on the international space station. ICARUS will in the near future provide orders of magnitude more data and with that challenges in data storage and analysis.

3.15 Data challenge: Learning Oyster catcher behavior

Daniel Weiskopf (Universität Stuttgart, DE)

In this talk, the results of our work on the Data Challenge are presented. With the Data Challenge, data on the movement of oyster catchers was provided to the participants before the seminar. We built a processing pipeline consisting of some custom code for data cleaning and preprocessing and subsequent supervised learning with RapidMiner and Weka. First, we reproduced the results previously reported in a paper on the Data Challenge data set and then compared different classification methods provided by RapidMiner. For our choice of data cleaning, W-SOM and Weka-LMT performed best, and slightly better than the decision tree of the original paper. Subsequently, we analyzed the time series of the recorded data to extend the feature vector. We chose to include additional information from the accelerometer. The extended feature vector led to further slight improvement of classification results. Finally, we briefly report on attempts to apply other visual analytics techniques to the Data Challenge: scatter/gather browsing of trajectories [1] and inter-active learning and visualization of classifiers [2].

References

1 B. Höferlin, M. Höferlin, D. Weiskopf, G. Heidemann. *Interactive Schematic Summaries for Exploration of Surveillance Video*. Proceedings of ICMR '11, 2011, Article No. 9. DOI: 10.1145/1991996.1992005.

2 B. Höferlin, R. Netzel, M. Höferlin, D. Weiskopf, G. Heidemann. *InterActive Learning of Ad-Hoc Classifiers for Video Visual Analytics*. Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST), 2012.

3.16 Visual analytics of eye-tracking trajectories

Daniel Weiskopf (Universität Stuttgart, DE)

Joint work of Andrienko, Gennady; Andrienko, Natalia; Burch, Michael; Weiskopf, Daniel
Main reference G. Andrienko, N. Andrienko, M. Burch, D. Weiskopf, "Visual Analytics Methodology for Eye Movement Studies," in IEEE Transactions on Visualization and Computer Graphics, 18(12):2889–2898, 2012.
URL http://dx.doi.org/10.1109/TVCG.2012.276

In this talk, methods for the spatio-temporal visual analytics of eye- tracking data are discussed [1]. The basic approach is to apply visual analytics methods previously developed for geospatial analysis. However, due to the characteristics of eye-tracking data, not all of the geospatial techniques are applicable. Appropriate analysis techniques are illustrated by running through an eye-tracking example data set acquired in the context of evaluating different type of tree visualization techniques [2, 3]. Finally, challenges and open questions of spatiotemporal analysis of eye- tracking data are discussed.

References

1 G. Andrienko, N. Andrienko, M. Burch, D. Weiskopf. *Visual Analytics Methodology for Eye Movement Studies*. IEEE Transactions on Visualization and Computer Graphics 18(12) 2889-2898, 2012. DOI: 10.1109/TVCG.2012.276

2 M. Burch, N. Konevtsova, J. Heinrich, M. Höferlin, D. Weiskopf. *Evaluation of Traditional, Orthogonal, and Radial Tree Diagrams by an Eye Tracking Study*. IEEE Transactions on Visualization and Computer Graphics 17(12):2440-2448 (2011). DOI: 10.1109/TVCG.2011.193

3 M. Burch, G. Andrienko, N. Andrienko, M. Höferlin, M. Raschke, D. Weiskopf. *Visual Task Solution Strategies in Tree Diagrams*. Proceedings of IEEE PacificVis 2013.

4 Working Groups

4.1 Data Challenge

As part of this Dagstuhl seminar, a data challenge was prepared. The challenge was on purpose kept small. It mainly aimed at familiarizing the participants with the topics of the seminar, establishing a harmonized language amongst a very diverse group of participants and ultimately shortening the warm-up phase at the seminar. The data challenge consisted of two distinct parts: one aimed at an analysis question and another aimed at interactive visualization. Both parts used the same data set of bird movement which has been published and described in [1]. The goal of the analysis-part of the data challenge was to find features and algorithms for segmentation, shape and movement analysis which could be successful in classifying foraging behavior. The goal of the visualization-part was to develop concepts and tools to visually explore differences in (model-predicted or observations of) bird movement and behavior. The task description and data were made available six weeks before the Dagstuhl seminar and all seminar participants were invited to join.

4.1.1 Results on the analysis challenge

Kevin Buchin, Maike Buchin, Rob Weibel, and Daniel Weiskopf

Rob Weibel and Daniel Weiskopf considered point-based (not considering segment shape or geographical relationships) classifiers and clustering techniques. A range of classifiers were compared under a variety of cross-validation schemes. The conclusion of this work was that it is worthwhile to cross-compare and integrate the various successful classifiers by the different groups and also to define features at different resolutions (in this way encapsulating different levels of spatio-temporal heterogeneity). The tools used for this part in the data challenge comprised RapidMiner (http://sourceforge.net/projects/rapidminer/), Weka (http://www.cs.waikato.ac.nz/ml/weka/) and some tailor-made C++ and R code.

Kevin and Maike Buchin considered methods to include temporal relations into the classification problem. Here a naive Bayes classifier was compared to a hidden Markov model. It appeared that considering time-dependencies did improve the correct classification, without even involving additional features. Hence considering time dependencies seems to be a fruitful avenue in enhancing the correct classification of movement data. This work has been implemented in R.

4.1.2 Results on the visualization challenge

Susanne Bleisch and Aidan Slingsby

Aidan Slingsby demonstrated a visualization interface. It visualizes the location of individuals (on top of Bing satellite imagery), with highlighting according to behavior and details per gps-point. Next to the geographic display the interface contains a timeline, showing the habitat, modeled behavior and observed behavior in parallel. And finally, the interface also shows barcharts of time spent in each habitat, modeled behavior and observed behavior per individual or jointly for the several individuals. All three panels are linked and the geographical display as well as the time-line can be used for zooming, panning or selecting a subset of the data.

Susanne Bleisch demonstrated an interactive visualization of the challenge data, comprising a set of spatial views with coloring by individual and subsetting by behavior. The data were shown at different scale levels and shown against a background that did show wrongly classified results. With both tools, an overview of the classification errors (e.g. the distribution over space or time, the individuals where errors are most prevalent) were quickly obtained; both tools were implemented in Java.

The contributions to this part of the data challenge demonstrated that interactive data visualization is of great value to effectively explore the properties of both movement data and prediction errors of that data. However, it also became clear that it is difficult to use generic tools to make in-depth analysis.

4.1.3 Future directions and panel discussion

The discussion started with a review and summary of previous activities related to data challenges in the area. Initial ideas about data challenges and benchmarking in movement analysis came out of the predecessor Dagstuhl seminars Nr. 08451 in 2008 and Nr. 10491 in 2010. A similar yet different group of experts at the GIScience 2010 workshop on Movement Pattern Analysis in Zurich further stressed the need for community initiatives facilitating a direct comparison of methods. Finally, at seminar Nr. 10491 in 2010 a group of researchers around Bettina Speckmann, Emiel van Loon and Ross Purves initiated a follow-up event hosted by the Lorentz Center in Leiden, NL [2].

In the Lorentz data challenge, the task focused on calculating a number of simple movement statistics that are generally reported in animal ecology (like trip duration, distance traveled and average speed). It appeared that in spite of the detailed guidelines and the relatively complete data sets, relatively large differences arose in the calculated summary statistics by the different participants. The conclusion of this data challenge was that it is dangerous to conduct meta-analysis on studies that only report summary statistics of movement; at the same time the need for better definitions or generally accepted and robust algorithms to compute the required statistics was emphasized. Furthermore, the importance of storing complete analysis workflows (i.e. including all data pre-processing steps) to enable reproducible research was demonstrated through this data challenge.

A workshop on Progress in Movement Analysis hosted at University of Zurich took a related yet different approach. This workshop included explicitly only reports of methodological contributions when they emerged multidisciplinary collaborations with domain experts (for example bringing together both ecological and algorithmic expertise). Integrating the experiences made at Leiden, at PMA10 and MPA12, and the mini-challenge stages at the current seminar, the discussion then turned to the contributions such data challenges might offer the community.

The group quickly agreed that the challenges led to moderate progress but fell short of achieving the hoped for general comparability of methods. Several reasons were given for this shortcoming. First, comparing to, for instance, the VisMaster challenge in the rather coherent information visualization community, the movement analysis community is rather diverse featuring researchers with very different research cultures. Second, it was stressed that all initiatives targeting methodological progress are organized and hosted by methods experts, repeatedly struggling attracting large numbers of domain experts and hence often falling short of being truly multidisciplinary.

The group then discussed pros and cons of data challenges, potentially aiming at establishing a long-lasting series of challenges co-hosted with a perennial conference.

Pros:
- Challenges are a good way to get students involved
- Challenges bring knowledge together and can promote the exchange of knowledge on how to apply established techniques

Cons:
- Challenges rather encourage practice of established methods, rather than innovation. Authors of innovative methods will rather publish their work somewhere else with more prestige.

The group concluded with the insight that even though the so far staged challenges significantly contributed to more transparency in methodological movement analysis and have propelled interdisciplinary work, the potential is now exploited as the small group has had ample opportunities for exchange. Before organizing new challenges, the community agreed on making more accessible what has already been achieved (Leiden Challenge). This could mean, (a) putting all results and code in the public domain, or (b) putting effort into developing techniques and tools to a high standard such that is can be easily reused (for example as *R*-libraries). Future challenges could receive more attention when co-hosted with an established conference (GIScience, ACM SIG Spatial) and when offering a more formal incentive (access to really interesting data sets). Finally, the group discussed funding opportunities for larger scale challenges. Initial ideas came from the movement ecology community, around possible challenges with multi-sensor, combined location-acceleration data.

References

1 J. Shamoun-Baranes, R. Bom, E.E. van Loon, B.J. Ens, K. Oosterbeek, W. Bouten. *From sensor data to animal behaviour: an oystercatcher example.* PLoS One, 7(5), e37997, 2012. http://dx.doi.org/10.1371/journal.pone.0037997

2 J. Shamoun-Baranes, E.E. van Loon , R.S. Purves, B. Speckmann, D. Weiskopf, C.J. Camphuysen. *Analysis and visualization of animal movement.* Biol Lett 8:6–9, 2012. http://dx.doi.org/10.1098/rsbl.2011.0764

4.2 Defining the theoretical core of movement science

A first series of working group sessions aimed in parallel for a definition of the theoretical core of movement science. The groups addressed the problem from three different perspectives. The first working group assembled an inclusive list of research questions emerging an eclectic set of domains studying movement (bottom up). The second group took a top-down perspective and worked on a table of contents for a planned inclusive textbook on movement analysis. The last working group studied a concrete use case, aiming at identifying the fundamental questions and challenges through a systematic walk through the use case.

Rather doing this than assembling yet another research agenda. List/TOC can serve as a starting point for potential review articles.

4.2.1 Collection of application driven questions

Andrea Kölzsch, Ran Nathan, Ross Purves, Roeland Scheepens, Aidan Slingsby, Bettina Speckmann, Egemen Tanin, Emiel Van Loon, Georgina Wilcox, and Jo Wood

The working group assembled a bank of classified application questions. The list may help in seeing how methods related to moving objects could be transferred from one domain to another, or to see if there are certain types of recurring questions for which there are not sufficient methods to address them. What started at the seminar is now an ongoing initiative where researchers are still encouraged to maintain and further develop the list and classification. The only conditions for questions is that they must be related to application domains and not directly about methods or specific data sets.

4.2.2 Table of contents of a Movement Science book

Susanne Bleisch, Maike Buchin, Györö Gidófalvi, Joachim Gudmundsson, Patrick Laube, Ben Loke, Mirco Nanni, Chiara Renso, Rodrigo I. Silveira, Egemen Tanin, and Rob Weibel

This group aimed at narrowing down and defining the core of movement science by assembling a table of content for a planned text book on the subject. The working group attracted a diverse set of researchers from the major areas advancing movement science, including geographical information science, computational geometry, computer science and database research. For the below listed sections, the group discussed and assembled the relevant methods according the following criteria: (i) Gather methods the group members have directly worked on, (ii) complement this list with relevant methods the group members have not directly worked on but consider relevant, (iii) exclude methods considered outside the focus of the planned book.

The characteristic of every method or methods category to be included were specified according to the following roster. This exercise help distinguishing and categorizing methods and finally positioning these methods in the table of contents.

- **Why** and **what for** would one use that method?
- **What** is processed? Does the method process fixes (observation points), subtrajectories (segments), trajectories, single moving point objects, or groups of moving point objects?
- **How** does the method actually work? What is the mathematical, statistical or algorithmic underpinning of the method?
- Are there unifying **criteria** that help categorizing the methods?

The tentative core of movement science mirrored in an table of contents for a text book are structured as follows.

1. **Conceptual Models of Movement.** Conceptual data models for movement space (Euclidean vs. network vs. Voronoi spaces), models for trajectories, other models for representing movement.

2. **Data sources, preprocessing, cleaning.** Data quality and uncertainty, map matching and filtering. Clearly, many preprocessing methods can be considered proper analysis methods.
3. **Management of movement data**. Storing, databases, moving object databases, indexing, querying, streams.
4. **Analysis.** Deriving movement parameters, outliers, similarity, clustering, segmentation, sequence analysis, movement patterns, trajectory simplification, smoothing, *aggregation*, sampling issues, movement simulation, visualization and mapping, visual analytics, semantic annotation, data integration.
5. **Evaluation.** Sensitivity experiments, Monte Carlo simulations, well-documented case studies and benchmarking.
6. **Tools.** Practical advice, tools that do a fraction, include a matrix techniques x tools, e.g. MoveBank.

The example for the task *aggregation* shall illustrate the chosen procedure. **Why** is aggregation of trajectory data required? Aggregation copes with information overload and allows an overview of large trajectory data sources, it also compensates for location uncertainty of individual fixes. **What** is processed? Typical aggregation methods act on fixes of trajectory or fixes of many trajectories of always the same or a group of moving point objects. **How** do aggregations methods work? Approaches include the selection of a representative trajectory, a median trajectory, flow maps, convex hulls, kernel density estimations, heat maps, and 3D shapes and density maps. What are **criteria** useful for a categorization? Methods can be separated depending on the nature of the results (entities or fields?), interrelations can be found to *clustering.*

Instead of listing a plethora of methods, the group worked out the contribution of the present disciplines (as above geographical information science, computational geometry, computer science and database research). The main contribution of the group can be seen in putting together and structuring the methods with respect to specific ways of addressing a problem and generic properties of the discussed methods. Clearly, there are many different ways of organizing methods into sections representing categories or steps along the analytical process. It is furthermore little surprising that the group identified many interrelations between methods. For instance, *aggregation* can be carried out based on *clustering* methods, or *segmentation* of trajectories can be done using *clustering* of fixes. Nevertheless, the discussion converged towards a conclusive categorized table of contents.

4.2.3 Interactive annotation: interfacing expert opinion and automated annotation tools

Kevin Buchin, Urska Demsar, Andrea Koelzsch, Kamran Safi, Emiel van Loon, and Georgina Wilcox

Movement analysis often requires to annotate trajectories with additional information. In movement ecology, for example, different stages of movement, characterized by different types of behavior, inform ecologists about the amount of time spent on different activities. Such activity budgets, including the precise onset and offset, are an essential part of life history strategies and can provide valuable information necessary to understand selective processes and conservation of populations. However, with the advent of cheap GPS devices and the ability to gather information with better spatial and temporal resolution, the method of manually annotating trajectories faces a logistic problem. It is therefore crucial to reconstruct the decisions experts make to annotate trajectories and automatically assign behavior to trajectories with a high reliability. Such a procedure requires an interface between the expert who is doing the annotation and the computer that evaluates these annotations to eventually provide automated annotation tools.

The group discussed the value of a thorough understanding of the workflow required for an efficient interface. Obtaining the best results both from domain experts and the chosen software is clearly a function of ease of use and information provided. Thus, increasing the willingness of experts to devote time to train a supervised machine learning algorithm requires some interaction with a system that needs to provide the necessary information in an intuitive way. At the same time, such an interface would need to provide the experts with clear definitions of behavior that they should assign and potentially involve the iterative presentation of important layers after previous assessment in addition to information such as satellite maps and topography or landuse.

Investigating the inter- and intra-expert variance would allow researchers to re-evaluate the definitions of behavioral classes and reduce ambiguities in their definitions. It would also allow a better understanding of the processes and decisions that experts take in the process of decision making. Also a comparison of mismatches between experts (or the median expert assessment) with the machine based assessment will allow to better understand the necessary information, that both the machine and the experts would require, for a more reliable assessment of behavioral patterns. Ultimately, the iterative annotation of trajectories is a problem that has implications well beyond annotation of trajectories for the assessment of time budgets. The principles can be applied to a wide variety of data that require annotation and where the amount of the data asks for an automated assessment, such as acceleration data or other types of movement data.

The conclusions of the group were to pursue a two step approach. In a first step a study would be conducted in the field of movement ecology where experts are asked to annotate behavioral stages to a set of water fowl data. These expert opinions would then be modeled using different statistical and machine learning procedures with previously environmentally enriched data to identify the important contextual layers that correlate with the different behavioral stages. This study can be published as such in an ecological frame work, however, along this project the workflow and the processes involved in the course of this project would be closely monitored. The experts who would do the annotations would be questioned as to their experience and suggestions in improving the workflow for them, in terms of visualization and feedbacks from the machines. This consideration of the workflow and reported suggestions for requirements would allow the group in a second step to formulate well defined challenge for computer science groups to provide tools and methods that substantially ease the process and provide efficiency gain in iterative annotation. Such a tool can then provide help bridging the gaps between man and machine where with every new annotated trajectory new and better models could be generated allowing to replicate the decisions that experts take. It would be also possible with such a system to investigate differences between experts and lay audience and the potential for crowd sourcing evaluated.

Participants

- Susanne Bleisch
The University of Melbourne, AU
- Kevin Buchin
TU Eindhoven, NL
- Maike Buchin
TU Eindhoven, NL
- Urska Demsar
School of Geography & Geosciences, GB
- Matt Duckham
The University of Melbourne, AU
- Györö Gidófalvi
KTH – Stockholm, SE
- Joachim Gudmundsson
The University of Sydney, AU
- Bin Jiang
University of Gävle, SE
- Andrea Kölzsch
NIOO-KNAW, NL
- Patrick Laube
Universität Zürich, CH
- Ben Loke
Noldus Information Technology BV, NL
- Mirco Nanni
ISTI-CNR – Pisa, IT
- Ran Nathan
The Hebrew University of Jerusalem, IL
- Ross Purves
Universität Zürich, CH
- Chiara Renso
ISTI-CNR – Pisa, IT
- Jörg-Rüdiger Sack
Carleton Univ. – Ottawa, CA
- Kamran Safi
MPI für Ornithologie, DE
- Roeland Scheepens
TU Eindhoven, NL
- Rodrigo I. Silveira
UPC – Barcelona, ES
- Aidan Slingsby
City University – London, GB
- Bettina Speckmann
TU Eindhoven, NL
- Egemen Tanin
The University of Melbourne, AU
- Marc van Kreveld
Utrecht University, NL
- Emiel Van Loon
University of Amsterdam, NL
- Robert Weibel
Universität Zürich, CH
- Daniel Weiskopf
Universität Stuttgart, DE
- Georgina Wilcox
The University of Sydney, AU
- Jo Wood
City University – London, GB