17th Symposium of the HPI Research Schools

Hasso Plattner Institute
Potsdam | April 25 – 29, 2022

Design IT. Create Knowledge.

HPI Hasso Plattner Institut
Digital Engineering · Universität Potsdam

https://hpi.de/symposium2022

Agenda

Monday, April 25, 2022 Master's Day

Room: D-Space (Main Building)

9:45 - 10:00	**Check-In**
10:00 - 12:00	**Why do your Ph.D.?**
	Hendrik Raetz, Dr. Katharina Baum, Prof. Dr. Felix Naumann, Prof. Dr. Holger Karl Panel Discussion
12:00 - 13:00	**Lunch Break (no catering)**
13:00 - 15:30	**Master's Information Day**
15:30 - 16:00	**Coffee Break**
16:00 - 17:30	**Showcase**
	Wenzel Pünter, Master Thesis Cybersecurity Towards Mining Actionable Cyber Threat Intelligence from Process Behavior Graphs
	Konstantin Dobler, Master Project IT Systems Engineering Generating Art with Multi-Conditional StyleGANs
	Christopher Aust & Ulrike Bath & Florence Böttger & Jonas Krah, Master Project IT Systems Engineering Robust Multi-Agent Reinforcement Learning for Scalable Failure Root-Cause Analysis
17:30 - 19:00	**Get-Together with Finger Food**

Tuesday, April 26, 2022 Future SOC Lab Day

Room: D-Space (Main Building)

9:00 – 9:20	**Opening Words**
	Prof. Dr. Andreas Polze, Operating Systems and Middleware Group, HPI
9:20 – 11:30	**Industry Talks**
	Andreas Rohnfelder, Fujitsu Quantum-Inspired Computing
	Andrey Hoursanov, SAP SE Quantum Experiments @ SAP
	Jan-David Brustik, Hewlett Packard Enterprise Bussiness Model Innovation in the IT-Industry - Why Circular Economy and Ecosystems are the Key to Success
	Jochen Glaser, SUSE Future Open Source Innovations for SAP Systems
12:00 – 13:30	**Lunch Break**
	Steering Committee Meeting
	Poster Session
13:30 – 14:15	**Tour of the New Data Ceter**
	Bernhard Rabe & Dr. Tobias Pape, HPI
14:15 – 15:15	**Research Talks I**
	Nadja Veigel, TU Berlin Mining Flood Insurance Big Data to Reveal the Determinants of Humans' Flood Resilience
	Iliyan Nenov & Prof. Dr. George Mengov, Sofia University St. Kliment Ohridski Neurocomputational Economic Forecasting in Turbulent Times
15:15 – 15:30	**Coffee Break**
15:30 – 16:30	**Research Talks II**
	Stefan Halfpap, HPI Query-Driven Partial Database Replication
	Hendrik Rätz, Wildenstein Plattner Institute Inc. Document Analysis for Cultural Heritage Data
16:30 – 17:00	**Closing Remarks**
	Prof. Dr. Andreas Polze, Operating Systems and Middleware Group, HPI

Wednesday, April 27, 2022 Doctoral Symposium

Room: D-Space (Main Building)

9:00 - 9:10	**Opening Words**
	Prof. Dr. Andreas Polze, Speaker of the SSE Research School, HPI
9:10 - 9:55	**Keynote**
	Dr. Melissa Densmore, Speaker of the Research School Branch at University of Cape Town Co-Designing Digital Resources for Bandwidth-Constrained Communities
10:00 - 11:00	**Session I**
	Huancheng Hu, Cybersecurity - Enterprise Security Group, HPI Dealing with Bad Apples: Organizational Awareness and Protection for Bit- and Typosquatting Attacks
	Shohei Katakura, Human Computer Interaction Group, HPI Kerfmeter: Automatic Kerf Calibration for Laser Cutting
	ZHI Yuan, Nanjing University
	DU Yuntao, Nanjing University Semi-Supervised Learning with Multi-Head Co-Training
11:00 - 11:30	**Coffee Break**
11:30 - 13:00	**Session II**
	Pauline Hiort, Data Analytics and Computational Statistics Group, HPI Drug Response Predictions with Networks from Molecular Data
	Mazhar Hameed, Information Systems Group, HPI An Overview of Data Preparation and Cleaning
	Marco Cipriano, Artificial Intelligence and Intelligent Systems Group, HPI An Alpha-Shape-Based Approach to Refine Manual Annotations of 3D Medical Targets
	WANG Yue, Data Engineering Systems, HPI Optimizing Decentralized Window Aggregation
	Keven Richly, Enterprise Platform and Integration Concepts Group, HPI Budget-Conscious Fine-Grained Configuration Optimization for Spatio-Temporal Applications
	Tzofnat Greenberg-Toledo, Faculty of Electrical and Computer Engineering, Technion Accelerating Deep Neural Network Applications
13:00 - 14:00	**Lunch Break (no catering)**

Wednesday, April 27, 2022 Doctoral Symposium

Room: D-Space (Main Building)

14:00 – 15:30	**Session III**
	Dorina Bano, Business Process Technology Group, HPI Mining Contextual Information from Event Logs
	Mary-Jane Antia, University of Cape Town Automating the Authoring of Competency Questions using Machine Learning
	Felix Grzelka, Operating Systems and Middleware Group, HPI Integration of a Clinical Decision Support System for Remote Patient Monitoring
	Sumit Shekhar, Computer Graphics Systems Group, HPI Low-light Image and Video Enhancement via Adaptive Chromaticity
	Johan Bontes, University of Cape Town Optimizing NP-Hard Problems Using Local Search
	David Sanka Laar, University of Cape Town Requirements for a Business Process Redesign Approach for SMEs in Developing Country Contexts
15:30 – 16:45	**Poster & Ice Cream Session**
16:45 – 17:45	**Session IV**
	Dylan Slack, University of California Irvine Examining Shortcomings and Improving the Reliability of Machine Learning Explanations
	Prof. Erik Sudderth, Speaker of the Research School Branch at University of California Irvine Learning to Make Consistent and Fair Predictions from Sparsely Labeled Data
18:00 – open	**Ph.D. Student Get-Together**

Thursday, April 28, 2022 Symposium Day 1

Room: HS 1 (Lecture Building)

12:30 – 13:00	**Check-In**
13:00 – 13:15	**Opening Words**
	Prof. Dr. Andreas Polze, Speaker of the SSE Research School, HPI
	Prof. Dr. Felix Naumann, Speaker of the DSE Research School, HPI
13:15 – 13:30	**Elevator Pitches – New SSE Members**
	Jorge Francisco Ciprian Sanchez (Computer Graphics Systems Group)
	Mustafa Ghani (System Analysis and Modeling Group)
	Eva Krebs (Software Architecture Group)
	Jobin Idiculla Wattasseril (Computer Graphics Systems Group)
	Afshin Zivi (Cybersecurity – Enterprise Security Group)
	Adrian Jobst (Computer Graphics Systems Group)
13:30 – 13:40	**Elevator Pitches – New DSE Members**
	Tahir Miri (Digital Health – Machine Learning Group)
	Margarita Bugueno (Artificial Intelligence and Intelligent Systems Group)
	Eshant English (Digital Health – Machine Learning Group)
	Fabian Stolp (Digital Health – Connected Healthcare Group)
	Khawla Rjeb (Data Engineering Systems Group)
13:45 – 14:30	**Keynote**
	Dr. Wolfgang Maier, Director of Hardware Development, IBM Research Böblingen How to Cope with Future Compute Demands
	Prof. Dr. Tobias Friedrich, Dr. Marcus Kölling Awarding of the HPI Fellowship
14:30 – 15:00	**Coffee Break**
15:00 – 16:00	**Session I**
	Prof. Dr. Maria Keet, University of Cape Town Evidence-based and Ontology-driven Logics for Conceptual Data Modelling – the Journey
	Prof. Dr. Heiko Hamann, Universität Lübeck Scalability in Robotics and Computing

Thursday, April 28, 2022 Symposium Day 1

Room: HS 1 (Lecture Building)

16:00 - 16:15	**Coffee Break**
16:15 - 17:15	**Session II**
	Kordian Gontarska, Operating Systems and Middleware Group, HPI Can we predict death? - Trade-offs in Clinical Decision Support Systems
	Prof. Dr. Peter Tröger, Berliner Hochschule für Technik
17:15 - 17:30	**Coffee Break**
17:30 - 18:00	**Session III**
	Dr. Gregor Seiler, IBM Europe Quantum Safe Cryptography
18:30 - open	**Boat Tour and Visit to Museum Barberini**

Friday, April 29, 2022 Symposium Day 2

Room: HS 1 (Lecture Building)

8:30 – 9:00	**Check-In**
9:00 – 9:30	**Keynote**
	Prof. Dr. Timo Hönig, Ruhr Universität Bochum System Software for the Next Generation: Energy Efficiency as Core Design and Operating Principle
9:30 – 11:00	**Session IV**
	Prof. Dr. Theo D'Hondt, Vrije Universiteit Brussel Memory-Critical Algorithms
	Max Plauth, Operating Systems and Middleware Group, HPI Improving Accessibility of Heterogeneous System Resources for Application Developers using Programming Abstractions
	Rana Shahout, Department of Computer Science, Technion Real Time Flow Monitoring
11:00 – 11:30	**Coffee Break**
11:30 – 13:00	**Session V**
	Dr. Michael Perscheid, Enterprise Platform and Integration Concepts Group, HPI What's Next after In-Memory Databases or How Does the Future of Business Applications Look Like?
	Justin Albert, Digital Health – Connected Healthcare Group, HPI Using Machine Learning to Predict Perceived Exertion During Resistance Training with Wearable Heart Rate and Movement Sensors
	Sarah Dsane, University of Cape Town Using familiar tools in an unfamiliar context: Sharing Co-designing experiences in a cross-cultural low-resource context
13:00 – 13:10	**Closing Remarks**
	Prof. Dr. Andreas Polze, Speaker of the SSE Research School, HPI

Research School Chairs

Prof. Dr. Felix Naumann

Speaker of the Research School on Data Science Engineering, HPI

Prof. Felix Naumann studied mathematics, economy, and computer sciences at the University of Technology in Berlin. After receiving his diploma (MA) in 1997 he completed his PhD thesis in the area of data quality at Humboldt University of Berlin in 2000. In 2001 and 2002 he worked at the IBM Almaden Research Center on data integration topics. From 2003 – 2006 he was assistant professor for information integration, again at the Humboldt-University of Berlin. Since 2006 he holds the chair for information systems at the Hasso Plattner Institute (HPI) at the University of Potsdam in Germany. He has been visiting researcher at QCRI, AT&T Research, IBM Research, and SAP. His research interests include data profiling, data cleansing, and data integration with over 200 scientific publications. Next to numerous PC memberships for international conferences, he has organized several conferences in various roles, including VLDB 2021 as PC co-chair, and he is trustee of the VLDB Endowment. More details are at `https://hpi.de/naumann/people/felix-naumann.html`.

Prof. Dr. Andreas Polze

Speaker of the Research School on Service-Oriented Systems Engineering, HPI

Andreas Polze is the Operating Systems and Middleware Professor at the Hasso Plattner Institute at University Potsdam, Germany. He is also the speaker of the HPI Research School and member of the steering committee of HPI's Future SOC Lab. Andreas received a doctoral degree from Freie University Berlin, Germany, in 1994 and a habilitation degree from Humboldt University Berlin in 2001, both in Computer Science. At HPI, his research focuses on architectures of operating systems, on component- based middleware, as well as on predictable distributed and cloud computing. Andreas Polze was visiting scientist with the Dynamic Systems Unit at Software Engineering Institute, at Carnegie Mellon University, Pittsburgh, USA, where he worked on real-time computing on standard middleware (CORBA) and with the Real-Time Systems Laboratory at University of Illinois, Urbana-Champaign. Current research interests include Predictable Service Computing, Adaptive System Configuration, and End-to-End Service Availability for standard middleware platforms. Together with Charité, GETEMED, and Deutsche Telekom, he has run the Fontane telemedicine project. Joint research with SAP has investigated porting HANA to new processor architectures.

Co-Designing Digital Resources for Bandwidth-Constrained Communities

Dr. Melissa Densmore

Speaker of the Research School Branch at University of Cape Town

Rhetoric around information and communications technology for development (ICT4D) often centers around giving 'have-nots' equal access to the internet. Access to information and social networking resources, it is posited, will lead to 'digital inclusion' and therefore positive social development. However, we argue that framing the problem around internet access limits the potential of communities for leveraging communications technology. Instead, we take a community-centered approach to innovation; from prioritizing issues to address, to active involvement in co-design of digital resources. I discuss what we have learned from our community engagements throughout South Africa, providing a motivation for Inethi, a platform to support community-based digital content and services. Our flagship deployment in Ocean View, Cape Town supports a community-owned-and-operated wireless network. We also have deployed iNethi to support parents of premature babies, and are exploring its usefulness in helping low-skilled workers to find employment. We make an argument for the need for iNethi and establish a research agenda for the iNethi around co-design of content and services, as well as innovative approaches to communications infrastructure.

Melissa Densmore runs the HCI Lab at the University of Cape Town (UCT), where she is an Associate Professor in the Department of Computer Science and Speaker of the Hasso-Plattner Institute Research School at UCT. Prior to joining UCT she did her postdoc at Microsoft Research India, and has also worked for Intel Research Berkeley, Xerox PARC, and a number of Silicon Valley startups. She has been in the ICT4D space since 2004, doing ethnographic field work, systems design, and deployments in Uganda, South Africa, Lesotho, Ghana, India, DR Congo, Rwanda, and Mexico. Melissa completed her PhD at University of California, Berkeley in Information Management and Systems, a 3 year ethnographic study of the use of Internet and mobile technologies by health practitioners and NGO staff in a health financing program in Uganda, has an MSc in Data Communications, Networks and Distributed Systems from University College London, and a BA in Computer Science from Cornell University. She is a member of the ACM SIGCHI Research Ethics committee, the ACM COMPASS steering committee, and the UCT TIA Seed Fund steering committee. Her research seeks to identify and foster community-based innovation around digital content and services for community wireless networks and for maternal and child well-being. The award-winning iNethi project provides a platform for innovative software to allow community-centered digital services, with a goal of using local ICTs to strengthen communities, understanding internet usage in bandwidth-constrained populations, and exploring ethical issues around community-based research.

Prof. Erik Sudderth

Speaker of the Research School Branch at University of California Irvine

Learning to Make Consistent and Fair Predictions from Sparsely Labeled Data

Real-world datasets are often only sparsely annotated by experts, and in domains like natural language understanding and computer vision, these labels may be implicitly corrupted by biases. We address these issues by learning probabilistic models that simultaneously respect two goals: discriminative classification and interpretable generative modeling. Our new approach scalably enforces prediction and fairness constraints when learning data models, leading to improved topic models of text and Markov models of time-series, as well as "very deep" generative models of images. We also develop new relaxations of fairness constraints that scale to flexible classifiers like deep neural networks for images and text.

Erik B. Sudderth is Professor of Computer Science and Statistics, and Chancellor's Fellow, at the University of California, Irvine. He directs the HPI Research Center in Machine Learning and Data Science at UC Irvine. Erik was previously an Associate Professor of Computer Science at Brown University, and a postdoctoral scholar at the University of California, Berkeley. He received the Bachelor's degree (summa cum laude, 1999) in Electrical Engineering from the University of California, San Diego, and the Master's degree (2002) and Ph.D. degree (2006) in EECS from the Massachusetts Institute of Technology. He received an NSF CAREER award and the ISBA Mitchell Prize.

Dr. Prof. Chongjun Wang was graduated with a Ph.D degree, from Nanjing University in Jun. 2004. Currently he is a full-time Professor at the Dept. of Computer Science and Technology, Nanjing University. He is also the coordinator for the HPI Research School at Nanjing University. His research interests include agent and multi-agent systems, complex network analysis, big data analysis and intelligent systems. He published 2 books and over 50 papers in journals and proceedings these five years. He has won several research funding from serval competitive sources such as the National Key Research and Development Program of China, the National Natural Science Foundation of China, the National 973 Program of China.

Prof. Chongjun Wang

Speaker of the Research School Branch at Nanjing University

Idit Keidar is the Lord Leonard Wolfson Professor at the Technion's Vitrerbi Faculty of Electrical Engineering. She received her BSc (summa cum laude), MSc (summa cum laude), and PhD from the Hebrew University of Jerusalem in 1992, 1994, and 1998, respectively. Subsequently, she was a Rothschild Postdoctoral Fellow at MIT's Laboratory for Computer Science. Prof. Keidar has served as the program chair for a number of leading conferences (PODC, DISC, PPoPP, and SYSTOR). She currently heads the Technion Rothschild Scholars Program for Excellence and consults for Yahoo Labs and Orbs.

Prof. Idit Keidar

Speaker of the Research School Branch at Technion

Prof. Dr. Tobias Friedrich

Algorithm Engineering Group, HPI

Prof. Dr. Tobias Friedrich heads the Algorithm Engineering Group at the Hasso-Plattner-Institut since 2015.

He received his M.Sc. in Computer Science in 2003 from the University of Sheffield and his Diplom in Mathematics in 2005 from the Friedrich-Schiller-Universität Jena. After completing his doctoral studies at the Max-Planck-Institut für Informatik and the Saarland University Saarbrücken in 2007, he became a postdoctoral researcher at the International Computer Science Institute in Berkeley. In 2009, he returned to the MPI in Saarbrücken as a research associate and later senior researcher. Between 2012 and 2015, Prof. Dr. Tobias Friedrich headed the Chair for Theoretical Computer Science I at the Friedrich-Schiller-Universität Jena.

Dr. Marcus Kölling

Chief Financial Officer, HPI

Dr. Marcus Kölling was born in Minden in 1973. He studied business administration and received his MBA from the HHL Leipzig Graduate School of Management. In 2004, Marcus Kölling received his Ph.D. at the Jean Monnet Centre of Excellence at the Wirtschaftsuniversität Wien and continued to work as a lecturer and research associate. In 2006, he returned to Leipzig and led the new Center for Leading Innovation and Cooperation at the HHL (CLIC) as director. The CLIC is an academic think tank focused on innovation. In March 2011, Dr. Marcus Kölling assumed control of the family-owned Kölling Group, which specializes in manufacturing and distribution of insulated glass products as well as associated services. Founded in 1947, the business group employs 220 people. In 2018, Dr. Marcus Kölling handed over control of the Kölling Group and served as chancellor and chief financial officer at the HHL Leipzig Graduate School of Management.

Since May 1st 2021, Dr. Marcus Kölling is Chief Financial Officer of the Hasso-Plattner-Institut in Potsdam. His sphere of responsibility spans the commercial field as well as strategic development, hand in hand with the HPI team.

Speakers

Dealing with Bad Apples: Organizational Awareness and Protection for Bit- and Typosquatting Attacks

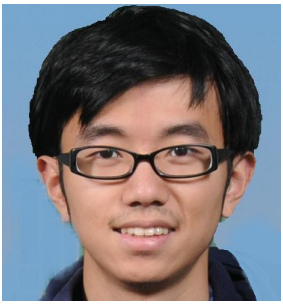

Huancheng Hu

Cybersecurity – Enterprise Security Group, HPI

The domain name system (DNS) maps humanreadable service names to IP addresses used by the network. As it exerts control over where users are directed to, domain names have been targets of abuse ever since the Internet become a success. Over the past twenty years, adversaries have repeated invented new strategies to trick users.

Aside from educating users, it is foremost the responsibility of organizations to monitor for and/or proactively register domain names with abuse potential. This however requires organizations to be aware and translate this into concrete action. While the typorelated attacks of the early 2000s are self-explanatory, other types of domain attacks are not. In this paper, we investigate the level of organizational awareness and preparedness towards two types of DNS abuse, the 20 year old, widely discussed typosquatting and those that result from internal computational errors, the so-called bit flipping attacks.We track the reaction and protection response of 300 large organizations over the course of 6 years, and find that by and large companies take little action towards this threat, with the exception of few well-prepared organizations that even deployed mitigations before the attacked were reported in the academic literature.

I am a PhD student interested in cybersecurity and under Enterprise Security team. My previous academic experience was in Purdue University and San Jose State University in the USA. My research interest is IoT network security.

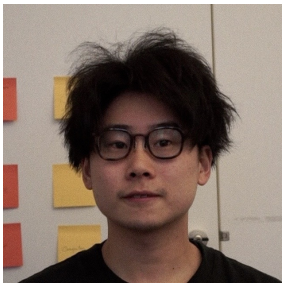

Shohei Katakura

Human Computer Interaction Group, HPI

Kerfmeter: Automatic Kerf Calibration for Laser Cutting

Laser cutting is a fast fabrication technique. 3D modeling software like flatfab enables users to make 3D laser cut models quickly. A software system like Roadkill makes assembling fabricated models more straightforward and faster. However, there is another time-consuming process that has not been discussed yet. It's called calibration. To achieve the desired result, users need to adjust laser energy and the laser-cut model based on kerf, the width of material removed by a laser cutter. My research challenge is automating this calibration process.

In this talk, we focus on kerf. We are developing a mechanical device to measure kerf for a given material and laser energy. The device performs the measurement by: (1) using the laser cutter to create a "gauge", (2) actuating rotating the gauge with the installed motor and (3) assessing the result rotation of gauge through the motor's encoder. The devices then pass the measurement to the software system, compensating the model to fit the measured kerf.

Shohei Katakura is a Ph.D. student at the Human-Computer Interaction group with Prof. Patrick Baudisch. His research interests are personal fabrication, especially turning digital fabrication machines into appliances by embedding intelligence. He received his Masters in Engineering from Meiji University, Japan.

Drug Response Predictions with Networks from Molecular Data

Pauline Hiort

Data Analytics and Computational Statistics Group, HPI

Networks are important to make interactions amenable to analysis. This has been shown to be particularly useful to investigate the molecular make-up of human cells where molecules regulate and interact with each other. Particularly in disease states, these regulations can change, and in-depth analysis on altered structures in comparison to the healthy state can help find appropriate medication or treatment options. Methods for single types of molecules have been proposed. We established a pipeline called DrDimont – **D**rug **r**esponse prediction from **Di**fferential analysis of **m**ulti-**o**mics **n**et**w**orks. It allows for differential analysis of multiple types of molecules from two conditions, e.g., two patient groups. Based on correlations of abundance measurements of multiple molecular layers, condition-specific networks are generated. Differential predictions, e.g., drug responses, are derived from comparing the condition-specific, multi-layer integrated networks. In a case study, we predict differential drug response in two breast cancer patient groups.

Pauline Hiort is a PhD student in the Data Analytics and Computational Statistics group of Prof. Renard at HPI. She has a bachelor's and master's degree in Bioinformatics. Pauline is currently focusing on employing networks from multiple layers of molecular data to predict drug response in cancer patients.

Mazhar Hameed

Information Systems Group, HPI

An Overview of Data Preparation and Cleaning

A significant obstacle when developing and deploying data science solutions is the poor state of data: files will not load, schemata are outdated, data are ill-formatted, incorrect, or simply missing. Data stewards, data scientists, and developers spend too much time finding, wrangling, and cleaning their training and test data. Only recently has our community begun to recognize such shortcomings as a research (and tooling) opportunity. In the talk, we will examine data quality problems at structural and semantical levels – from the mundane, such as unexpected field delimiters, to the complex, such as automatic structure correction. We will explore the research efforts of our data preparation team to discover and address such issues and highlight outstanding research challenges in the field of data quality and data preparation.

Mazhar Hameed is a 3rd year Ph.D. student specializing in data preparation and cleaning with Prof. Felix Naumann in the Information Systems group at Hasso Plattner Institute. His research interests include character-separated value (CSV) file's dialect detection, field format normalization, and record structure detection and correction.

An Alpha-Shape-Based Approach to Refine Manual Annotations of 3D Medical Targets

Marco Cipriano

Artificial Intelligence and Intelligent Systems Group, HPI

Deep neural networks have provided amazing results for a multitude of computer vision tasks, including medical imaging applications. In this field, technologies like Cone Beam Computed Tomography have become an enabling factor for applying deep learning to 3D images. While the supervised learning paradigm of deep neural networks requires large amounts of annotated data to work, the annotation of three-dimensional images is still a cumbersome procedure and represents a constraint in the quality of the available datasets.

In this talk, I show the limitations of current annotation techniques and describe alternative methods to obtain smooth labels of 3D targets in medical imagery.

He received his B.Sc. and M.Sc. degrees in Computer Engineering from Università degli Studi di Modena e Reggio Emilia, Italy. He spent two years as a Research Fellow at the AImageLab Laboratory at Dipartimento di Ingegneria Enzo Ferrari from the same University. He is currently a Ph.D. student at the Hasso Plattner Institute in the Deep Data Lab. His research interests include deep learning, computer vision, natural language processing, and medical imaging.

WANG Yue

Data Engineering Systems, HPI

Optimizing Decentralized Window Aggregation

Internet-of-Things (IoT) devices are used in many fields, in industry as well as in research. These devices are connected in huge decentralized networks. The decentralized networks produce large data streams that need to be processed in a timely fashion. To evaluate unbounded data streams, events from those data streams are grouped into bounded windows based on queries. The stream processing engines are deployed in huge decentralized networks and perform those windows. These SPEs often run their computation cluster on data centers, which is for processing queries. However, data streams are distributed among decentralized networks. That leads to SPEs having to centralize data streams to their computation clusters. To handle decentralized networks, state-of-the-art processing queries are close to data streams, but they can not perform complex windows. Furthermore, in order to process multiple queries, state-of-arts can slice windows and then share partial results between these queries. Nevertheless, they can not share slices between queries that have different window types, measures, and aggregation functions, yet.

We present Desis, a decentralized stream processing system that is designed for efficiently processing multiple continuous queries in decentralized networks. Instead of collecting events, Desis pushes down processing queries close to data sources which is beneficial in decentralized networks. We propose an optimizer that can share partial results between multiple queries with different window types, measures, and aggregation functions.

I am a Ph.D. student working at the Data Engineering Systems team at the HPI. Currently, I am working on stream processing systems with distributed sensor network. I am also interested in modern hardware and databases.

Budget-Conscious Fine-Grained Configuration Optimization for Spatio-Temporal Applications

Keven Richly

Enterprise Platform and Integration Concepts Group, HPI

Based on the performance requirements of modern spatio-temporal data mining applications, in-memory database systems are often used to store and process the data. To efficiently utilize the scarce DRAM capacities, modern database systems support various tuning possibilities to reduce the memory footprint (e.g., data compression) or increase performance (e.g., additional index structures). However, the selection of cost and performance balancing configurations is challenging due to the vast number of possible setups consisting of mutually dependent individual decisions.

In this talk, we describe a novel approach to jointly optimize the compression, sorting, indexing, and tiering configuration for spatio-temporal workloads. Further, we consider horizontal data partitioning, which enables the independent application of different tuning options on a fine-grained level. We introduce an integer linear programming approach to compute optimized tuning configurations for a given workload and memory budgets. Moreover, we demonstrate on a real-world dataset that our models allow to significantly reduce the memory footprint with equal performance or increase the performance with equal memory size compared to established rule-based heuristics.

Keven Richly is a Ph.D. student at the Hasso Plattner Institute. As a member of the Enterprise Platform and Integration Concepts research group headed by Prof. Hasso Plattner, his research focuses on spatio-temporal data analysis and data management. His current research interests are in the field of workload-driven storage optimizations for columnar in-memory databases to store and process spatio-temporal data in a cost-efficient way.

Tzofnat Greenberg-Toledo

Faculty of Electrical and Computer Engineering, Technion

Accelerating Deep Neural Network Applications

Deep Neural Networks (DNNs) achieve state-of-the-art results in a broad range of applica- tions, such as image and speech recognition and natural language processing. Therefore, these computational models have gained increasing interest from the scientific community. DNNs are usually executed by commodity hardware (mostly FPGA and GPU platforms) and accelerators (such as Google's TPU), as they are compute and memory intensive. However, when executing DNN algorithms, the conventional von Neumann architectures, where the memory and computation are separated, pose significant limitations on performance and energy efficiency. Many studies have thus been devoted to accelerate the execution of DNNs.

In this research, we will investigate different aspects of DNN acceleration and execution, considering both hardware-based and algorithm-based approaches and how they may affect each other.

We investigated thus far accelerating DNNs using emerging memristive memory technologies, such as resistive RAM (RRAM) and spin-transfer torque magnetic RAM (STT-MRAM). The emergence of memristors has enabled combining data storage and computation using the same physical entities. Hence, it is possible to compute an energy efficient and highly parallel analog multiply and accumulate (MAC) operation in-place, also known as processing-near-memory (PNM).

I am a Ph.D. candidate, working on architecture and algorithm levels. My research interests are efficient deployment and acceleration for Deep Neural Networks (DNN). In this research, I am using the memristor's properties to implement energy-efficient dedicated architecture for DNN. I am working at the Technion as a teaching assistant and a lab instructor. I have a B.Sc. in Electrical engineering from the Technion. During my undergraduate studies, I worked as a logic design engineer at Intel Corp.

Mining Contextual Information from Event Logs

Dorina Bano

Business Process Technology Group, HPI

Business process mining is becoming an increasingly important field for understanding the behavioral perspective of any given organization. In a process mining project, process experts are tasked with discovering or improving the operational business processes. They do so by analyzing event logs, the starting point of any process mining endeavor. Despite event logs capturing behavioral information, we argue that they are also a rich source of domain-specific information. This information is not represented explicitly in a process model but, nevertheless, it provides valuable contextual information. To this end, we propose a semi-automatic approach to discover a data model that complements traditional process mining techniques with domain-specific information.

Process discovery is an important area in the field of process mining, where most discovery algorithms focus on process control flow, giving little attention to the data-flow perspective. As a result, the discovered process models lack information about data dependencies, and process experts need to manually enrich the discovered process models accordingly. This requires deep domain knowledge, is not scalable, and is error-prone. To overcome this limitation, we propose an approach that aims to discover the data objects and their behavior by investigating how event attributes are manipulated during process execution. The resulting data objects are used to enhance the discovered process model.

My name is Dorina Bano and currently I am working as Research Assistant at Hasso Plattner Institute, University of Potsdam. I have studied Computer Science at the University of Paderborn in Germany, and I have 4 years of experience in the industry as a Software Engineer. During my work experience at Entimo AG, besides acting as a software engineer, I have been closely involved in several Pharmaceutical projects. I was in daily contact with pharma companies to gather their requirements, plan the product roadmap and, finally, presenting the results. Currently, I am writing my PhD thesis related to the Process Mining and Business Process Management area.

Mary-Jane Antia

University of Cape Town

Automating the Authoring of Competency Questions using Machine Learning

Competency Questions (CQs) continue to play very important roles in Ontology Engineering. However, the authoring of CQs continue to be impeded by a number of challenges: bad grammar, quality of questions, answerability, the manual nature of the process which is time consuming, etc. In response to some of these challenges, there has been an increase in research on the translation of CQs to axioms and SPARQL/SPARQL-OWL queries. While studies on the use of controlled natural language have shown progress in rectifying some of the issues, CQs still need to be authored manually first. In this study we leverage the use of machine learning (specifically, transfer learning methodology) as well as templates to create a model that automates the process of authoring CQs. We make use of the CQs used in creating the CLaRO templates which have been certified in terms of grammatical quality, answerability and to a certain level, generalizability, to unseen CQs. We use the CQs underlying the CLaRO templates to develop a model from which knowledge learned can be transferred to other questions generated. We evaluate the model with questions generated from a COVID-19 text corpus. Afterwards, we extract patterns from the questions and map these questions to the CLaRO templates. The next step will be to extract CQs which have characteristics of questions used in the model Then subsequently, extracting the Competency Questions which carry the characteristics of the questions used in the model. Expected results will include the model metrics and accuracy, possible new template patterns, and CQs for the COVID-19 text corpus.

My name is Mary-Jane Antia, I am married and a mum of three incredible boys. I studied Information Technology for my master's degree where my research was on applying the concept of knowledge Graph to multilingual learning of high school Life Science subject. I am currently working on natural language generation and Ontology engineering for my PhD in Computer Science; both masters and PhD at the University of Cape Town. My PhD thesis focus is on Competency Questions authoring automation and metrics for evaluating them. I am keen to understand and apply machine learning concepts to Ontology related problems. I enjoy the outdoors especially nature; so, when I get chance to escape from the non-negotiable tasks, I go hiking!

Integration of a Clinical Decision Support System for Remote Patient Monitoring

Felix Grzelka

Operating Systems and Middleware Group, HPI

Clinical Decision Support Systems (CDSS) can help medical personnel, among other things, by sorting patients based on individual patient risk and need for medical interventions. Such an AI-based system was built for the Telemed5000 Project, which aims to improve the care of heart failure patients (HF) by employing remote patient monitoring. The goal of the project is to scale the number of patients that can be monitored by such a system by sorting the patients according to the predicted risk of cardiovascular interventions. For a CDSS to work, it needs to be integrated and interoperable with different components, such as patient records and data formats produced by medical devices.

In this talk, I will give an overview of the architecture of the Telemed5000 System, run a Live End-to-End Demo, and discuss some of the challenges when deploying deep neural networks to production.

Felix Grzelka is a Ph.D. student at the Operating Systems and Middleware group of Prof. Dr. Polze at HPI. His research interests include the approvability of AI as medical products and reproducibility of data pipelines and data science in general. He finished his B. Sc. in computer science at the Brandenburg University of Technology Cottbus-Senftenberg in 2018 and his M. Sc. in IT-Systems Engineering at HPI in 2021.

Sumit Shekhar

Computer Graphics Systems Group, HPI

Low-light Image and Video Enhancement via Adaptive Chromaticity

Image acquisition in low-light conditions suffers from poor quality and significant degradation in visual aesthetics. This affects the visual perception of the acquired image and the performance of various computer vision and image processing algorithms applied after acquisition. Especially for videos, the additional temporal domain makes it more challenging, wherein we need to preserve quality in a temporally coherent manner. We present a simple yet effective approach for low-light image and video enhancement. To this end, we introduce "Adaptive Chromaticity", which refers to an adaptive computation of image chromaticity. The above adaptivity allows us to avoid the costly step of low-light image decomposition into illumination and reflectance, employed by many existing techniques. All stages in our method consist of only point-based operations and high-pass or low-pass filtering, thereby ensuring that the amount of temporal incoherence is negligible when applied on a per-frame basis for videos. Our results on standard low-light image datasets show the efficacy of our algorithm and its qualitative and quantitative superiority over several state-of-the-art techniques.

I am in the fourth year of my PhD at the Computer Graphics Systems group at HPI. My research focuses on creating efficient image and video processing algorithms using their intrinsic attributes. Previously, I finished my Master's degree in "Visual Computing" at Saarland University.

Optimizing NP-Hard Problems Using Local Search

Johan Bontes

University of Cape Town

Many industrial optimization problems, such as planning, scheduling, routing, and bin packing are part of the NP-Hard class. This means that – even for moderately sized problems – optimal solutions cannot be found in reasonable time. If we relax our ambition and drop the requirement for optimality we find that local search using satisfiability finds better solutions in shorter time than most competing approaches. Commercial mixed linear programming (MLP) approaches still slightly outperform our method, but the latter are problematic, because its internals are not made public to the scientific community.

In this talk we present our work on local search and how it can be applied to optimize hard combinatorial optimization problems.

Johan Bontes did his Masters in Computer Science at the University of Cape Town (UCT) and is currently in the third year of his PhD in computer science, also at UCT, working on high performance automated reasoning. His work focusses on applying compute power to solve practical problems.

David Sanka Laar

University of Cape Town

Requirements for a Business Process Redesign Approach for SMEs in Developing Country Contexts

As engine of economic development in developing countries, and for competitiveness, small and medium-sized enterprises (SMEs) need to adopt innovative business practices to deal with their challenges such as resource poverty, and limited business skills in the volatile business climate. Business process (BP) redesign provides transformational capabilities that can improve the performance of these enterprises through innovation. However, the extent of BP redesign adoption and process maturity among SMEs in developing country contexts is woefully low as these enterprises lack access to appropriate BP redesign methods. Thus, the overarching research question is How can SMEs in developing country contexts redesign their BPs? To address this question, the Design science research methodology is employed to achieve three main objectives as to: (1) determine the requirements of SMEs for process redesign, (2) develop a BP redesign approach for SMEs, and (3) validate the proposed approach in SMEs. This paper focuses on the first objective, using a qualitative exploratory design to investigate the environment of selected SMEs in Ghana to understand their needs with regards to BP redesign. The findings show that the environment of the SMEs is highly uncertain, and they require an evolutionary, agile and context sensitive BP redesign approach. These findings serve as criteria for developing and evaluating BP redesign methodologies suitable SMEs in developing country contexts.

David Sanka Laar is currently a PhD candidate at the Department of Information systems, University of Cape Town, South Africa. His research interests are in the area of ICT4D, Business process redesign, Enterprise systems, and Software engineering, employing user-centred design, qualitative and design science research methods. His PhD research centres on redesigning business processes in small and medium-sized enterprises in developing country contexts. David is from Ghana, and he is passionate for and open to collaboration in the use of ICTs for emancipating the marginalised in society through innovation. He can be contacted on kontilaar@yahoo.co.uk; https://www.facebook.com/david.laar.5/

Examining Shortcomings and Improving the Reliability of Machine Learning Explanations

Dylan Slack

University of California Irvine

For domain experts to leverage machine learning (ML) models in critical settings, such as healthcare, they must trust model predictions. Consequently, researchers have proposed numerous ways to explain the predictions of complex ML models. However, these approaches suffer from several critical drawbacks such as, instability, inconsistency, and lack of guidance about accuracy and correctness. This talk will contextualize these shortcomings and propose new methods that leverage improved uncertainty modeling to help overcome several critical issues of explanations.

I am a third-year PhD candidate at University of California, Irvine advised by Sameer Singh and Hima Lakkaraju. I work on machine learning as part of UCI NLP, UCI CREATE, and the HPI Research Center. My research is supported by an HPI fellowship. I've previously interned at AWS in Summer 2020 and Google AI in Summer 2021.

Speakers

Dr. Wolfgang Maier

Director of Hardware Development, IBM Research Böblingen

How to Cope with Future Compute Demands

Information processing capabilities play a central role in our private life, business environment as well as scientific issues. No other technology branch has taken a stronger influence on our social habits and structures over the recent decades. Along with this evolution economic power as well as political stability indicate significant dependencies on the future engagement in the advanced development of information technology in all its various subjects. A key field in this respect is the future development of advanced hardware components, which address a variety of challenges given by dynamically changing boundary conditions. These include higher demand of computing power and higher energy efficiency paired with technical feasibility and economical solidity.

The presentaion will discuss how new advanced workloads cause increased demand in computing power and gives deeper insights which possibilities are exhausted to deal with this. Specifically the latest trends in the field of microprocessor development in regard to semiconductor advancements, architectural and topological adjustments as well as completely new compute paradigms like quantum computing will be introduced.

Wolfgang Maier is Director of Hardware Development at the german sited IBM R&D lab close to Stuttgart.

His career with IBM started as design engineer for Mainframe Firmware in 1996. 3 years later he accepted the role of department manager for System z IO Firmware and started to focus on the employment of industry standard IO technologies in IBM high end servers. During his international assignment in Austin (Tx) he directed the first implementation of the Infiniband technology for the IBM POWER line. After his return to germany he extended the use of industry standards within the mainframe and headed the development of central IO hardware as well as mainframe and POWER CMOS processors. His current responsibilities also span the system packaging and system control area with special focus on workload optimization and power efficiency. In addition in recent years his organisation has build a Quantum Computing team, which signs responsible for the implementation and operation of the first installed IBM Quantumsystem in Germany. His current focus in innovation is on Semiconductor Technology and Quantum Computing. Wolfgang Maier obtained his PhD in Laserphysics from the University of Tuebingen in 1996. He enjoys to spend his spare time with his family, loves skiing and cycling.

Evidence-based and Ontology-driven Logics for Conceptual Data Modelling – the Journey

Prof. Dr. Maria Keet

University of Cape Town

Interoperability among UML class diagrams, EER, and ORM diagrams is still important for collaboration in software design and for data integration. A logic-based reconstruction may assist, but the multiple extant formalisations are incompatible due to the plethora of logics used or due to diverse hidden choices that permeate the formalisations, which are consequences of the lack of a systematic design process. We aim to address these problems by, first, structuring the logic design process in a methodological way. Second, we specify minimal logic profiles availing of this process, including considerations on ontological commitments embedded in the languages and availing of evidence of language feature usage, and of computational complexity insights. The profiles characterise the essential logic structure needed to handle the semantics of conceptual models, therewith enabling the development of interoperability tools. The talk will also reflect on the various processes of doing the research over the timespan of the seven years that it took.

Prof. Dr. Maria Keet is an Associate Professor with the Department of Computer Science, University of Cape Town, South Africa. She focuses on ontologies, conceptual data modelling, and their interaction with natural language tools, which has resulted in around 150 publications, including at several top-rated conferences and journals. She is, and was, PI on two NRF-funded projects on natural language generation with ontologies, she was PI on an DST/MINCyT-funded bi-lateral project with Argentina on ontology-driven conceptual modelling, and she was involved in several EU projects. She has served in numerous Program Committees and reviewed for numerous journals; recently she's a PC chair for the ISWC 2022 Resources Track, IJCAI 2021 Demos, and EKAW 2020. She also wrote an award-winning textbook on ontology engineering. Before her employment at UCT, Maria was Senior lecturer at the School of Computer Science, University of KwaZulu-Natal, South Africa and before that, an Assistant Professor at the KRDB Research Centre, Free University of Bozen-Bolzano, Italy. She obtained a PhD in Computer Science at the KRDB Research Centre in 2008.

Prof. Dr.
Heiko Hamann

Universität Lübeck

Scalability in Robotics and Computing

How to find an efficient group size or number of CPUs? How to balance opportunities for collaboration and constraints due to contention? The scalability properties of robot swarms have astonishing similarities to those of distributed computing, communication networks, car traffic, and even natural systems. We have a look on how to model system scalability in a generic way and on how to determine appropriate group sizes. With a better understanding of how system performance scales with system size, we can develop more efficient systems.

Since 2017 Heiko Hamann is professor for service robotics at the University of Luebeck, Germany. His main research interests are swarm robotics, bio-hybrid systems, evolutionary robotics, and modeling of complex systems. He enjoys working in multi-disciplinary teams with ethologists, plant biologists, architects, and psychologists in EU-funded projects, for example flora robotica, WatchPlant, and ChronoPilot.

Can we predict death? - Trade-offs in Clinical Decision Support Systems

Kordian Gontarska

Operating Systems and Middleware Group, HPI

Being able to automatically predict clinically relevant events like death, hospitalizations, strokes, or infarctions may drastically improve personal health care. Constant patient monitoring enables better medical treatment as it allows practitioners to react on time and provide the appropriate treatment. Telemedicine can provide constant remote monitoring so patients can stay in their homes, only requiring medical sensing equipment and network connections. A limiting factor for telemedical centers is the amount of patients that can be monitored simultaneously.

We aim to increase this amount by implementing a clinical decision support system in our project consortium Telemed5000. This talk aims to give you an overview and understanding of the telemedical scenario, the data we worked with, our solution approach, and the insights we gained along the way.

Kordian is currently a PhD candidate in the Operating Systems and Middleware research group of Prof. Dr. Andreas Polze at the Hasso-Plattner Institute. He holds a Bachelor and Master in Computer Science from the Free University of Berlin. His research interests concentrate around developing fast adapting machine learning models in the tele-medicine domain and for the configuration of distributed system parameters.

Prof. Dr. Timo Hönig

Ruhr Universität Bochum

System Software for the Next Generation: Energy Efficiency as Core Design and Operating Principle

Electrical energy is the single most important operating resource for computer systems. Therefore, the conscious use of energy resources is crucial and it depends on the respective systems: from "smart dust" and battery-powered mobile systems to high-performance computing clusters that operate under thermal constraints.

This keynote talk focuses on system software and elaborates how energy-aware programming techniques serve as a core design principle for the construction of energy-efficient systems. Runtime and operational aspects are discussed using an operating system that takes advantage of dynamic electricity prices to ensure the efficient operation of heterogeneous HPC clusters. This talk will further take a look on most recent hardware developments (i.e., non-volatile memory) to identify upcoming challenges that will dramatically reshape the structure of future operating systems for next-generation computer systems.

Prof. Dr.-Ing. Timo Hönig is leading the Bochum Operating Systems and System Software (BOSS) research group at the Ruhr-Universität Bochum (RUB). Dr. Hönig is Professor at RUB where his group is affiliated with the Faculty of Computer Science. At RUB, his research focuses on operating systems, system software, and energy-aware systems. He explores techniques for energy-aware programming, the design of energy-efficient system software, and high-performance computing systems with highly dynamic energy demand. In 2017, he received his doctorate (Dr.-Ing.) at Friedrich-Alexander-Universität Erlangen-Nürnberg. Dr. Hönig gained ten years of industry experience as senior software engineer and research fellow at SUSE Linux GmbH in Nürnberg (Germany) and in the Embedded Hardware Development Group at IBM Research Böblingen (Germany). He is a member of ACM and GI. Timo Hönig serves on international program committees and regularly reviews submissions to renowned journals (e.g., ACM TACO).

Memory-Critical Algorithms

Prof. Dr. Theo D'Hondt

Vrije Universiteit Brussel

We call algorithms memory-critical if they critically use the memory that is actually the domain that they act upon. This complicates the immediate use of memory-related language abstractions. An example of a memory-critical algorithm is a recursion stack. A more challenging one is a garbage collector.

Typically, memory-critical algorithms are designed as state machines. In this way well-known tools become accessible for the verification, documentation and specification leading to an actual implementation.

Memory-critical algorithms typically involve a clever technique to subsume the memory required by the algorithm under the memory that is the domain of the algorithm. The former is typically a greatly resized version of a regular approach to the algorithm. They often end up being very terse and hard to conceive or comprehend, and when expressed as actual code they are equally hard to debug or optimise.

Memory-critical algorithms have been the object of scientific publications, educational material and actual software since the early 60's. In order to present these algorithms, techniques have been proposed ranging from mathematical formalisms, graphical representations and pseudocode. Unfortunately, these all suffer from the pithy nature of these algorithms and the lack of a uniform medium to reason about them. Hence their notoriously bad reputation with students of advanced algorithms.

This presentation reports on an initiative to bind formalism, graphics and (executable) code in a framework to reason about memory-critical algorithms. It started out as an accumulation of convenient representations, and grew into something that has all the hallmarks of a systematic approach. This will be illustrated by applying this approach to several historical algorithms.

Max Plauth

Operating Systems and
Middleware Group, HPI

Improving Accessibility of Heterogeneous System Resources for Application Developers using Programming Abstractions

The ever-growing demand for compute resources has reached a wide range of application domains, and with that has created a larger audience for compute-intensive tasks. At the same time, heterogeneous computer architectures have become vital in order to deliver the corresponding growth in compute capacity. However, in order to make use of the diverse capabilities of heterogeneous compute resources, application developers have to deal with the complexities introduced by each class of compute resources. With the goal of making heterogeneous computing accessible for a wider range of application developers, the goal of my thesis is to investigate programming model abstractions for various types of compute resources. Addressing scale-out GPU deployments, CloudCL attempts to provide a unified programming model that hides complexities of both OpenCL and distributed computing frameworks. Targetting large-scale NUMA systems, the C++-based PGASUS framework simplifies data placement based on polymorphic memory allocators.

In 2014, Max Plauth obtained his M.Sc. in IT Systems Engineering from the Hasso Plattner Institute at the University of Potsdam in Germany. In his master's thesis, he studied the benefits of leveraging graphics processors for computational tasks in audio signal processing scenarios. Subsequent to his graduation, he joined the workforce of the IVU Traffic Technologies AG as a software engineer. He contributed to a software product that supports the daily workflows of transport companies by aggregating large data volumes. After having gained one year of industry experience, he returned to the academic world in February 2015 and started his PhD in the Operating Systems and Middleware Group at the Hasso Plattner Institute. His research interests include programming abstraction mechanisms for heterogeneous hardware resources.

Real Time Flow Monitoring

Rana Shahout

Department of Computer Science, Technion

Keeping data on network flows, referred to as flow monitoring, is a critical part of a wide variety of networking applications, including load balancing, denial of service detection, traffic engineering, and fairness. Along with flow monitoring, network routers need to route messages quickly. This limits the usage of slower memory types; instead, we utilize high-speed but scarce and costly memory.

Our solutions focus on minimizing the number of counters needed, thereby allowing the system to monitor a large number of elements using only a small amount of space. This approach often sacrifices accuracy in favor of space frugality.

Rana Shahout is a Ph.D. candidate at the Technion. She is working under the supervision of Prof. Roy Friedman. Her research focuses on streaming algorithms include network monitoring and interval query capabilities for real-time network problems. She completed her BSc at the Computer Science Department at the Technion and worked at Mellanox and Yahoo.

Dr. Michael Perscheid

Enterprise Platform and Integration Concepts Group, HPI

What's Next after In-Memory Databases or How Does the Future of Business Applications Look Like?

The Enterprise Platform and Integration Concepts Group of Prof. Dr. h.c. mult. Hasso Plattner is well-known for its research and contributions to in-memory databases over the last decades. His group envisions to find the fastest way to get insights out of enterprise data. Therefore, they are researching on technical aspects and design principles of enterprise systems with the goal to maximize performance, cost efficiency, and business value. To be more specific, they focus on optimizing data management on modern in-memory and cloud hardware, rethinking software engineering for enterprises, and enabling machine learning for data-driven decision support.

After finishing his studies at the HPI back in 2008, Michael did his dissertation in Computer Science under the supervision of Prof. Dr. Robert Hirschfeld in the HPI research school on "Service-oriented System Engineering". After that, he joined the SAP Innovation Center in 2014 and coordinated projects such as Hasso Plattner's and Bernd Leukert's published book "The In-Memory Revolution" as well as Hasso Plattner's SAPPHIRE keynotes since 2014. Since 2015, Michael was the head of Chairman and Strategic Projects at SAP. In this role, he led highly skilled, results-oriented, and technology-focused development teams to come up with tangible results for high-priority inquiries from the Executive Board and Supervisory Board of SAP. In May 2020, he received the request for a special strategic project, and he returned to the Hasso Plattner Institute to direct the department of Professor Plattner and to strengthen the collaboration with SAP. In 2021, Michael has been awarded with an ACM Senior Membership.

Using Machine Learning to Predict Perceived Exertion During Resistance Training with Wearable Heart Rate and Movement Sensors

Justin Albert

Digital Health – Connected Healthcare Group, HPI

Intensity quantification is an essential aspect of weight training to adjust training routines and prevent injuries. Usually, training intensity can be assessed objectively by, e.g., measuring the lifted weight. Subjective measurements such as questionnaires or ratings of perceived exertion (RPE) can also deliver important information from athletes. A standard RPE scale is the Borg scale, ranging from six (minimal exertion) to 20 (maximal exertion). In this project, the aim is to predict RPE values from athletes performing physical exercise. We recorded data using inertial measurement unit (IMU) sensors, electrocardiography (ECG), and RGB-D cameras of athletes performing squats for a given number of sets. We calculate heart-rate variability (HRV) parameters from the ECG signals and various statistical features from the IMU and camera data. After the feature extraction process, we utilize machine learning methods such as Support Vector Machines or K-Nearest-Neighbors to estimate the subjects' RPE values. Also, we identify the most meaningful features from the three sensor modalities for explainability.

Justin Albert joined the HPI Data Science and Engineering Research School in October 2019. Before, he obtained a B.Sc. degree in Computer Science in New Media, followed by an M.Sc. degree in Computer Science with an emphasis on Computer Vision and Machine Learning. His Ph.D. research focuses on the simultaneous use of optical sensors (cameras) and wearable sensors in biomechanical applications. He aims to develop machine learning and sensor fusion methods to analyze the performance of physical exercises performed in the recreational or rehabilitation setting.

Sarah Dsane

University of Cape Town

Using familiar tools in an unfamiliar context: Sharing Co-designing experiences in a cross-cultural low-resource context

Sharpies, sticky notes, and large swaths of paper are familiar tools for engagements with our participants. And yet, as we work to be more inclusive with our designs and engagements, these tools are not familiar to all. In this pictorial, we reflect on how households of pregnant women in Ghana responded to the elements of our 'basic' co-design toolkits, including sticky notes, pens, coloured pencils, design cards, and other props for empowering user engagements. Getting them to express themselves using these items to trigger the ideation process was challenging, yet the design cards – some prepared initially as part of another project in South Africa, were effective for elucidating contextualised conversations. In this pictorial we reflect on various aspects of the design cards and other materials, drawing lessons learned on engaging in contexts that may be less familiar with our traditional co-design tools.

Sarah Dsane is a Lecturer in the Computer Science Department at Koforidua Technical University. She is a gender advocate and has embarked on several projects aimed at equipping young women in TVET. She is currently in her third year of PhD in Human-Computer Interaction at the University of Cape Town. Her research explores explicitly an aspect of maternal and child health called the First Thousand Days (FTD), where she seeks to understand how parents are using digital resources in bandwidth-constrained communities in Cape Town and Ghana. She employs the use of co-designing and participatory design in her research.

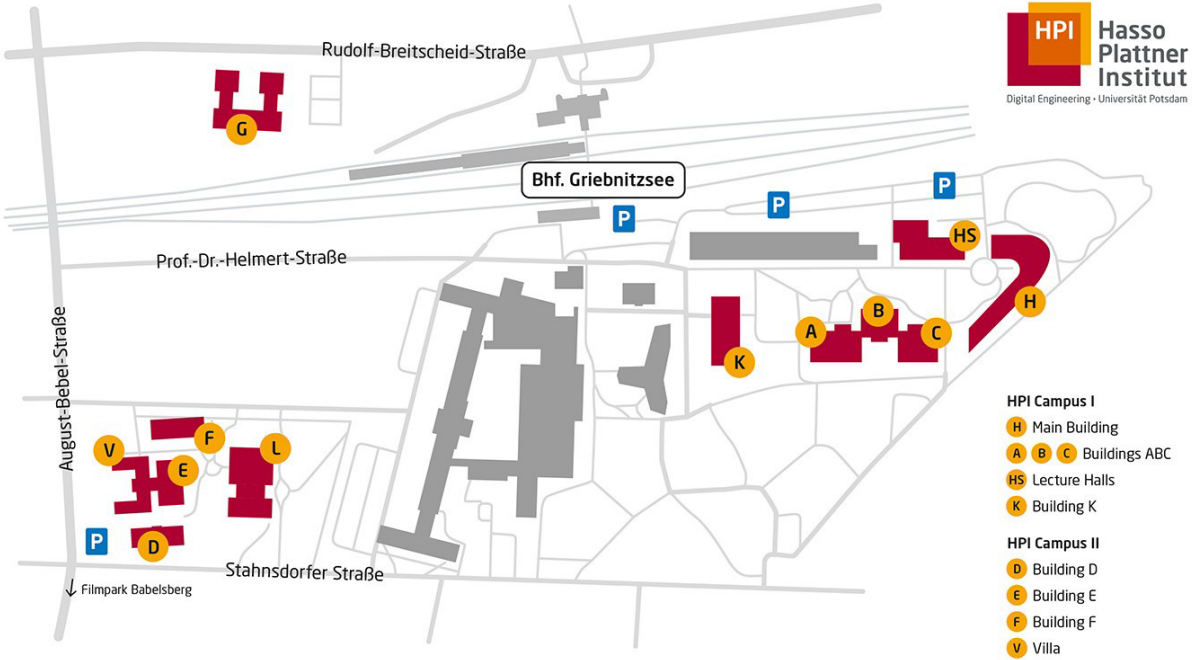

HPI Hasso Plattner Institut
Digital Engineering · Universität Potsdam

HPI Campus I

- H Main Building
- A B C Buildings ABC
- HS Lecture Halls
- K Building K

HPI Campus II

- D Building D
- E Building E
- F Building F
- V Villa
- L Building L

HPI Campus III

- G Building G

Maps

Notes

Contact

Hasso-Plattner-Institut für Digital Engineering gGmbH
Campus Griebnitzsee | University of Potsdam

Prof.-Dr.-Helmert-Straße 2-3
D-14482 Potsdam, Germany

Phone +49 331 5509-0
E-Mail hpi-info@hpi.de
Web https://www.hpi.de

Photo: HPI / Kay Herschelmann, Charles Yunck