UNIVERSITÄT
BERN

Institut für Informatik
Universität Bern

www.inf.unibe.ch

INF Annual Report 2024/25

INF Annual Report

Academic Year 2024/2025

September 30, 2025

Contents

1 Institute of Computer Science (INF)

1.1 Address

Neubrückstrasse 10, 3012 Bern, Switzerland
☎ +41 31 684 86 81
✉ info.inf@unibe.ch
https://www.inf.unibe.ch

1.2 Personnel

Members

Florence Aellen, Jesutofunmi Ajayi, Dr. Jayamine Alupotha, Prakash Aryan, Sabyasachi Banik, Vivien Bammert, Mariarosaria Barbaraci, Jan-Andrea Bard, Michal Bechny, Christian Birchler, Roman Bögli, Alexander Boll, Prof. Dr. David Bommes, Dr. Peppo Brambilla, Prof. Dr. Torsten Braun, Prof. Dr. Christian Cachin, Llukman Cerkezi, Jinxuan Chen, Hamadi Chihaoui, Bettina Choffat, Alexandros Christopoulos, Annalisa Cimatti, Riccardo Cusinato, Dr. Reza Darooei, Aram Davtyan, Yihan Deng, Dr. Antonio Di Maio, Kalvin Dobler, Djanira dos Santos Gomes, Chenrui Fan, Benjamin Fankhauser, Prof. Dr. Paolo Favaro, Armand Singh Satjeevan Feuilleaubois, Chuyang Gao, Anthony Gillioz, Pinar Göktepe, Priska Grunder, Elham Hashemi Nezhad, Dragana Heinzen, Martin Heistermann, Sandro Hernández Goicochea, Dr. Benedikt Hitz-Gamper, Fereshteh Jafari, Ali Javadi, Dr. Linlin Jia, Polina Jordan, Denis Kalmykov, Liubov Kamaldinova, Prof. Dr. Timo Kehrer, David Lehnherr, Francesco Leonardi, Roman Macháček, Corina Masanti, Sajad Mazraehkhatiri, Dr. Camille Mignardot, Marcio Moraes Lopes, Franziska Müller, Valentin Müller, Valentin Nigolian, Joël Niklaus, Sajedeh Norouzi, Ivonne del Carmen Núñez Carrillo, Dr. Manuel Ohrndorf, Dr. Sebastiano Panichella, PD Dr. Kaspar Riesen, Atefeh Rohani, Sepehr Sameni, Eric Samikwa, Bruno Santos Martins, Alp Eren Sari, Dr. Jakob Schärer, Milian Luca Scharr, Dr. Philippe Schneider, Alexander Schultheiss, Michael Senn, Viktor Shipitsin, Borja Sierra Miranda, Jonas Spieler, Dr. Thomas Studer, PD Dr. Matthias Stürmer, Hongyan Sun, Mingjing Sun, Thomas Sutter, Radoslava Svihrova, Dr. Rojo Rendrianomentsoa, Dr. Atefeh Rohani, Luca Rolshoven, Dr. Christos Tsigkanos, Prof. Dr. Athina Tzovara, Ramazan Erdem Uysal, Pablo Valenzuela Toledo, Juan Fernando

Villacis Llobet, Anh Duc Vu, Solomon Fikadie Wassie, François-Xavier Wicht, Hexu Xing, Lukas Zenger

Administration

Bettina Choffat, Priska Grunder, Dragana Heinzen, Polina Jordan, Franziska Müller

Technical staff

Dr. Peppo Brambilla, Martin Heistermann, Michael Senn

1.3 Activities

Contribution to the “Studies Orientation Day”, Bern, September 13, 2024

Contribution to the “National Future Day”, Bern, November 13, 2024

Contribution to the “Bachelor’s open days“, December 2+3, 2024

Contribution to the “MINT-Day” for Middle school students, Bern, March 4, 2025

Contribution to the “Master’s open days”, March 18 + 20, 2025

2 Teaching Activities

2.1 Courses for Major and Minor in Computer Science

Autumn Semester 2024

Bachelor 1st Semester (each 5 ECTS)

- Diskrete Mathematik (C. Cachin)
- Grundlagen der Technischen Informatik (J. Schärer)
- Programmierung 1 (K. Riesen)

Bachelor 3rd Semester (each 5 ECTS)

- Computernetze (T. Braun)
- Einführung in Software Engineering (T. Kehrer)
- Digitale Nachhaltigkeit (M. Stürmer)

Bachelor 5th Semester (each 5 ECTS)

- Computergrafik (D. Bommes)
- Mensch-Maschine-Schnittstelle (K. Riesen)
- Machine Learning (P. Favaro)
- Anleitung zu wissenschaftlichen Arbeiten (Die Dozenten der Informatik)

Master Courses (each 5 ECTS)

- Software Product Lines (T. Kehrer)
- Distributed Algorithms (C. Cachin)
- Computer Vision (P. Favaro)
- Applied Optimization (D. Bommes)
- Network Security (T. Braun)

Seminars (each 5 ECTS)

- Seminar: Software Engineering (T. Kehrer)
- Seminar: Pattern Recognition (K. Riesen)
- Seminar: Machine Learning and Artificial Intelligence (P. Favaro)
- Seminar: Computer Graphics & Geometry Processing (D. Bommes)
- Seminar: Cryptography and Data Science (C. Cachin)
- Seminar: Communication and Distributed Systems (T. Braun)
- Seminar: Natural Language Processing (NLP) (M. Stürmer)

Service Courses (each 3 ECTS)

- Programmieren für Naturwissenschaften (K. Riesen)
- Programmieren für Naturwissenschaften (Biologie) (P. Liniger)
- Grundkurs Programmieren (B. Hitz-Gamper, M. Stürmer)

Spring Semester 2025

Bachelor 2nd Semester (each 5 ECTS)

- Datenbanken (T. Studer)
- Datenstrukturen und Algorithmen (D. Bommes)
- Computer Architecture (P. Favaro)
- Programmierung 2 (T. Kehrer)

Bachelor 4th Semester (each 5 ECTS)

- Praktikum in Software Engineering (T. Studer)
- Betriebssysteme (T. Braun)
- Berechenbarkeit und Komplexität (T. Studer)
- Algorithmen, Wahrscheinlichkeit und Information (C. Cachin)
- Programmieren für Datenwissenschaften (M. Anwander)

Bachelor 6th Semester (each 5 ECTS)

- Anleitung zu wissenschaftlichen Arbeiten (Die Dozenten der Informatik)

Master Courses (each 5 ECTS)

- Advanced Networking and Future Internet (T. Braun)
- Deep Learning (P. Favaro)
- Cryptography(C. Cachin)
- Graph Based Pattern Recognition (K. Riesen)
- 3D Geometry Processing (D. Bommes)
- Compiler Construction (T. Kehrer)

Seminars (each 5 ECTS)

- Seminar: Machine Learning and Artificial Intelligence (P. Favaro)
- Seminar: Communication and Distributed Systems (T. Braun)
- Seminar: Logic and Theoretical Computer Science (T. Studer)
- Seminar: Applied Optimization (D. Bommes)
- Seminar: Cryptology and Data Security (C. Cachin)
- Seminar: Software Engineering (T. Kehrer)
- Seminar: Pattern Recognition (K. Riesen)
- Seminar: Die Psychologie von IT-Sicherheit und Datenschutz (C. Cachin, M.Elson, L.A. Jansen, M.F. Vogt)

Service Courses (each 3 ECTS)

- Programmieren für Naturwissenschaften (K. Riesen)
- Grundkurs Programmieren (J. Schärer)

2.2 Students

Major Subject Students:	AS 2024	343	SS 2025	322
Minor Subject Students:	AS 2024	165	SS 2025	128
Ph.D. Candidates:	AS 2024	56	SS 2025	63

2.3 Degrees and Examination

PhD: 6

Master: 30

Bachelor: 32

Completition of Minor Studies: 39 (BA Minors: 90E: 0, 60E: 10, 30E: 5, 15E: 7; MA Minor 30E: 1; 885 ECTS)

Semester Examinations AS 2024: 1355 (4689 ECTS)

Bachelor's/Master's Thesis AS 2024: 24 (440 ECTS)

Semester Examinations SS 2025: 861 (3217 ECTS)

Bachelor's/Master's Thesis SS 2025: 28 (500 ECTS)

3 Cognitive Computational Neuroscience Group

3.1 Personnel

Head
Prof. Dr. A. Tzovara ✉ athina.tzovara@unibe.ch
☎ +41 31 684 3306

Office Manager
D. Heinzen ✉ dragana.heinzen@unibe.ch
☎ +41 31 684 4914

Scientific Staff
Dr. F.M. Aellen ✉ florence.aellen@unibe.ch
Dr. C. Mignardot ✉ camille.mignardot.inselspital@gmail.com
R. Cusinato ✉ riccardo.cusinato@unibe.ch
A. Christopoulos ✉ alexandros.christopoulos@unibe.ch

External PhD students
G. Monachino ✉ giuliana.monachino@unibe.ch
M. Bechny ✉ bechnymichal@gmail.com
R. Svihrova ✉ rsvihrova@gmail.com
Y. Deng ✉ dengyihan@gmail.com

3.2 Overview

The Cognitive Computational Neuroscience group conducts research in the areas of neuroscience, machine learning and computational modeling. We use invasive and non-invasive electrophysiological techniques to record neural activity in humans, in combination with machine learning techniques, to study neural functions of the human brain in health and disease. The main areas of focus include: (a) machine learning algorithms for analysing neurological data, and (b) studying the neural correlates of sensory processing and predictions.

3.3 Research Projects

Intrinsic timescales and spatio-temporal integration in wakefulness and sleep

The brain exhibits neural activity that unfolds over multiple timescales and spatial ranges of integration. These patterns, that are observed in the absence of external stimulation, are thought to be a fundamental property of the brain's organization. Intrinsic timescales throughout the cortex are organized in a hierarchical manner, so that areas that are higher in a cortical hierarchy (e.g. associative areas) have longer timescales than those who are lower (e.g. sensory cortices). This property has been evaluated in several studies in wakefulness, but it remains unknown whether and how sleep reconfigures brain hierarchies. In our work, we analysed intracranial electroencephalogrphy (iEEG) recordings, directly from the human cortex, during wakefulness and sleep in a cohort of 106 patients with epilepsy. We computed the intrinsic timescales and spatial integration of iEEG signals via their temporal and spatial autocorrelation in wakfulness and sleep. We analysed broadband signals (across a range of frequencies) and also gamma signals (higher frequencies). We found that broadband and gamma timescales become longer in sleep compared to wakefulness. However, with different patterns: broadband timescales increased along the sensorimotor-association axis, while gamma ones decreased. During sleep we found that slow waves can explain the increase in both broadband and gamma timescales, with broadband ones showing a strong correlation with slow wave density. In terms of spatial integration, we found a high integration at long distances in wakfulness for broadband and at short distances for gamma timescales in sleep. These findings show the co-existence of multiple timescales across the cortex that follow different anatomical hierarchies and are reconfigured by sleep and wake dynamics.

Research staff: Riccardo Cusinato, Camille Mignardot, Alexandros Christopoulos, Athina Tzovara

Financial support: Swiss National Science Foundation (320030-227728)

Evolution of neural synchrony and complexity in post-hypoxic ischemic coma

To maintain consciousness, the human brain integrates information across regions, while it needs to exhibit a diverse repertoire of neural activity. Patients in a coma after cardiac arrest are often in an unconscious state. To assess their chances of recovery, electroencephalography (EEG) is used at the patients' bedside. In recent work (Alnes et al., 2021), we had shown that, in the first day after coma onset, patients who eventually recover have higher integration of information than non-survivors. It remains unknown how the temporal dynamics of neural synchrony (a measure of information integration) and complexity (a measure of diversity of neural activity) evolve over time of coma. In our work, we assessed the temporal evolution of neural synchrony and complexity of auditory EEG responses in a cohort of comatose patients following a cardiac arrest. We showed that the phase locking value (PLV), a measure of neural synchrony, was informative of the chances to regain consciousness the first 40 hours after a cardiac arrest. After that, its predictive value diminished, as patients with a poor prognosis showed a progressive increase in PLV. By contrast, patients who eventually regained consciousness had a stable and relatively high PLV. We also analysed the complexity of EEG responses, which did not differ between survivors and non-survivors. Neural complexity had a progressive reduction over time of coma. This work suggests that the acute coma phase, and in particular the first 24 hours after coma onset carry high predictive information about the patients' chances to regain consciousness. This finding has implications for guiding future guidelines for quantitative techniques for coma outcome prognostication.

Research staff: Sigurd Alnes, Florence Aellen, Athina Tzovara

Financial support: Swiss National Science Foundation (320030_188737), Olga Mayenfiseh Stiftung

Neural representation of prospective decision outcomes

Humans make numerous decisions every day which are affected by the context around them. For instance, the choice of whether to carry an umbrella often depends on the season when this decision is made. Despite its high relevance, the majority of existing work studies decision-making

irrespective of context. In our work, we aimed at studying how the context into which a decision is made affects decision deliberation, and more specifically the retrieval of prospective outcomes. We performed novel electroencephalography (EEG) experiments, in which participants were presented with pairs of an object and its context and were asked to make decisions, leading to a given outcome. We used machine learning classifiers to model the neural representation ofpossible outcomes, in terms of EEG topographies. These classifiers were then tested during the deliberation period, before participants observed the outcome of their decisions. Our results showed a strong and above chance classification performance in decoding the identity of possible outcomes, based on EEG responses. Moreover, we found above chance decoding of these outcomes during the deliberation period, before the outcomes were presented. Our findings suggest that the most likely outcomes of a given decision were retrieved already 3 seconds before participants made that decision. Moreover, the strength of decoding of prospective outcomes strongly correlated with the behavioral performance of participants, so that participants with stronger neural outcome retrieval made more accurate decisions. This result was only observed in the case where context could alter the outcome of a given decision. Our work suggests that the context of a decision alters retrieval of prospective outcomes and paves the way for future studies in the field of decision making and outcome deliberation.

Research staff: Pinar Göktepe-Kavis, Florence Aellen, Athina Tzovara

Financial support: Swiss National Science Foundation (320030_188737)

Machine learning and statistical methods for healthcare

In recent years, advances in deep learning and computational modeling have transformed healthcare by enabling the discovery of intricate patterns within large and complex datasets. In clinical settings where interpretability and reliability are essential, probabilistic and statistical modeling approaches remain highly valuable. In our work, we combine the strengths of deep learning and advances statistical methods to develop personalized, data-driven systems for patient monitoring. One first focus is on extracting meaningful electrocardiogram (ECG) representation via deep learning. A key objective is to enable effective learning even when only limited labeled data are available. From a technical perspective, this

study emphasizes the exploration, application, and optimization of self-supervised learning approaches and foundation models. These methods make it possible to harness large volumes of unlabeled data and transfer the acquired knowledge to downstream tasks that often suffer from data scarcity. The techniques have been applied to several tasks, including the detection of life-threatening arrhythmias (LTAs), classification of atrial fibrillation (AF), and broader cardiac diagnosis. In addition, new approaches have been developed to improve adaptability to variable ECG lead configurations and to enhance performance on small downstream datasets. A second focus is the development of digital health interventions for behavioral change in lifestyle medicine. By incorporating a causal reinforcement learning framework, adaptation of interventions is enabled for improved timing (just-in-time adaptive interventions) and contextual information, with micro-randomization for allowing subsequent inference. To enhance this sequential decision making, we employ multi-armed bandits, namely Thompson Sampling thanks to its exploration-exploitation trade-off property. For improved personalization we aim to understand causal relationships between lifestyle factors based on the observational data from wearable devices and digital diaries, combined with expert knowledge. Population-level graphs are estimated by causal discovery algorithms, and inference is done with (Bayesian) generalized mixed-effects models to account for intra-individual differences. The framework is a first step towards building personalized behavioral digital twins, a model that adapts based on the real-time data, capable of computing counterfactuals and thus simulating interventions before they are executed. A third focus is on sleep and its characteristics for quantifying new digital biomarkers of present and future diseases. To this end, we investigate primarily sleep-stage dynamics to derive metrics that may serve as sensitive biomarkers. The dynamics are quantified using explainable machine learning tools (causal meta-learners, dynamic Bayesian networks, and forest-based approaches), both to illustrate that sleep-dynamics parameters can be used for diagnostics and to understand the effect of individual predictors. The approaches have been successfully demonstrated in several use-cases, including obstructive sleep apnea, chronic fatigue and pain syndromes, and long-term risk of cardiovascular disease.

Research staff: Giuliana Monachino, Radoslava Svihrova, Michal Bechny

Text-based prediction for a patient cohort with sleep disorders

The current clinical documentation of sleep disorders primarily relies on free-text physician reports derived from polysomnography, actimetry, and vigilance tests. In our earlier work, we developed a prototype text-processing pipeline based on embedding techniques and neural network classifiers (Multiclass CNN, Hierarchical Attentive Network), to analyse this free text. We are currently analyzing discrepancies between patient self-perception and clinical investigations using sentiment analysis methods. Specifically, we combined BERT-based models with a GPT-2 generative framework to contrast subjective expressions against standardized objective benchmarks in a daytime sleepiness cohort, demonstrating that generative models can also perform effectively in clinical sentiment analysis tasks. Moreover, we are aiming at extracting comorbidities from clinical reports in patients with sleep disorders. To this aim, we are applying an LLM-based multi-level classifier (Mixtral-24B) with a multi-agent architecture. Future work can generalize this architecture to other multi-level classification tasks.

Research staff: Yihan Deng

3.4 Theses

3.4.1 PhD Theses

- Pinar Göktepe-Kavis, University of Bern, PhD in Computer Science, PhD Thesis: ”Machine learning for decoding brain states and decisions from electroencephalography signals”, September 2024

3.4.2 Master’s Theses

- Dominik Fischli, University of Bern, Master of Computer Science, Master’s Thesis: ”Nonnegative matrix factorization for detecting interictal epileptiform discharges”, June 2025
- Mena Lerf, University of Bern, Master of Artificial Intelligence in Medicine, Master’s Thesis: ”AI to study music perception via intracranial EEG recordings in patients with epilepsy”, November 2024
- Sophie Caroni, University of Fribourg, Specialized Master of Science in Digital Neuroscience, Master’s Thesis: ”Investigation of circadian

patterns in coma with signal processing and machine learning techniques", September 2024

3.5 Further Activities

Presentations

Athina Tzovara

- Intrinsic neural timescales in the temporal lobe - lessons from intracranial recordings in epilepsy patients, Epilepsy Research Symposium, Swiss Epilepsy Center, Zurich, November 2024
- Comment étudier les fonctions cérébrales dans le coma avec l'intelligence artificielle?, Université des aînés de langue française de Berne (UNAB), October 2024
- Computational assessments of brain functions in sleep and coma, Farewell symposium Prof. Bassetti, Bern, September 2024

Florence Aellen

- Slow wave and spindles characteristics in acute and chronic use of BZ receptor agonists, Poster presentation / Multimedia Talks, Clinical Neuroscience Bern, September 2025
- Computational techniques for studying neural functions in coma, Pint of science, Bern, May 2025
- Temporal Dynamics of Heartbeat Evoked Potentials in Comatose Patients in Long Continuous Recordings, Poster presentation, Swiss Society for Neuroscience, Lausanne, February 2025
- NerAi: An analysis tool for predicting survival of comatose patients with AI, Talk, InnoForum: Venture Fellows 2024, November 2024
- Temporal dynamics of heartbeat evoked potentials in comatose patients in long continuous recordings, Poster presentation, Clinical Neuroscience Bern, September 2024

Sigurd Alnes

- Intracranial EEG insights into inter-individual communication, Poster presentation, Society for Neuroscience, Chicago, October 2024

Camille Mignardot

- Corteo: A novel Software for Integrated Brain Network Analysis, Poster presentation, International Conference for Technology and Analysis of Seizures (ICTALS), Montreal, Canada, June 2025

Riccardo Cusinato

- Sleep modulates neural timescales and spatiotemporal integration in the human cortex, Swiss Society for Neuroscience, Lausanne, February 2025
- A journey of sound in the brain: A neuroscience perspective on auditory processing, CAS in AI for creative practices, November 2024

Pinar Göktepe-Kavis

- What does my network learn? Assessing interpretability of deep learning for EEG, Oral presentation, the 48th Annual Meeting of the Japan Neuroscience Society, July 2025

Michal Bechny

- Sleep-Stage Dynamics Predict Current Sleep-Disordered Breathing and Future Cardiovascular Risk, Talk and poster presentation, World Sleep Congress, Singapore, September 2025
- Benchmarking Reveals Shared Biases in Sleep-Staging Algorithms and Points to Age-Aware Solutions, Poster presentation, World Sleep Congress, Singapore, September 2025

Radoslava Svihrova

- Towards Digital Twins with Wearables: A Fully Data-driven Causal Analysis of Lifestyle Effects on Sleep Biomarkers. Poster presentation, World Sleep Congress, Singapore, September 2025
- Wearable-derived Sleep Digital Biomarkers and Well-being: Preliminary Insights from Longitudinal Observational Study. Poster presentation, World Sleep Congress, Singapore, September 2025
- Data-driven causal discovery: Insights from a longitudinal study with wearable data, Poster presentation, 23rd International Conference on Artificial Intelligence in Medicine (AIME 2025), June 2025
- Wearables and sleep: Effect of circadian rhythm disruption on sleep quality, recovery, and well-being, Poster presentation, Congress of European Sleep Research Society in Seville, September 2024
- Wearables and Sleep: Effect of circadian rhythm disruption on sleep quality, recovery and well-being, Giornata della Ricerca e dell'Innovazione in Medicina Umana della Svizzera Italiana, Lugano, September 2024
- AI and Statistics in Medicine AI and Medical Devices Conference, Prague, Czech Republic, September 2024

Giuliana Monachino

- Addressing ECG data scarcity with AI: From Transfer Learning to Foundation Models. Oral presentation at AI in Electrophysiology Workshop at 23rd International Conference on Artificial Intelligence in Medicine (AIME 2025), June 2025

Conference and Scientific Boards/Committees

Athina Tzovara

- Organization for Human Brain Mapping (OHBM), member of Diversity and Gender Committee, 2024
- Swiss league against epilepsy, Board member

Journal Committees

Athina Tzovara

- Handling Editor for Imaging Neuroscience

Reviewing Activities

Journal Reviews
Athina Tzovara

- eClinicalMedicine
- Journal of Neuroscience
- Sleep

Florence Aellen

- European Journal of Neuroscience

Riccardo Cusinato

- Imaging Neuroscience

Grant reviews
Athina Tzovara

- European Research Council
- Bern Medtech Collaboration Call

3.6 Publications

Journal Publications

- Alnes SL, Aellen FM, Rusterholz T, Pelentritou A, Hänggi M, Rossetti AO, Zubler F, De Lucia M, Tzovara A, (2025) Temporal dynamics of neural synchrony and complexity of auditory EEG responses in post-hypoxic ischemic coma, Resuscitation, https://doi.org/10.1016/j.resuscitation.2025.110531
- Cusinato R, Seiler A, Schindler K, Tzovara A, (2025) Sleep modulates neural timescales and spatiotemporal integration in the human cortex, Journal of Neuroscience, https://doi.org/10.1523/JNEUROSCI.1845-24.2025
- Boccalaro IL, Aime M., Aellen FA, Rusterholz T, Borsa M, Bozic I, Sattin A, Felin T, Gutierrez Herrera C, Tzovara A, Adamantidis AA, (2025), A role for the thalamus in danger evoked awakening during sleep, Nature Communications, https://doi.org/10.1038/s41467-025-62265-0
- Monachino G, Zanchi B, Wand M, Conte G, Tzovara A, Faraci FD, (2025), Overcoming data scarcity in life-threatening arrhythmia detection through transfer learning Communications Medicine, https://doi.org/10.1038/s43856-025-00982-9
- Benchy M, Fiorillo L, van der Meer J, Schmidt M, Bassetti C, Tzovara A, Faraci FD, 2025, Beyond accuracy: a framework for evaluating algorithmic bias and performance, applied to automated sleep scoring Scientific Reports, https://doi.org/10.1038/s41598-025-06019-4
- Benchy M, Kishi A, Fiorillo L, van der Meer J, Schmidt M, Bassetti C, Tzovara A, Faraci FD, 2025, Novel digital markers of sleep dynamics: causal inference approach revealing age and gender phenotypes in obstructive sleep apnea Scientific Reports, https://doi.org/10.1038/s41598-025-97172-3
- Svihrova R, Dei Rossi A, Marzorati D, Tzovara A, Faraci FD, (2025), Designing digital health interventions with causal inference and multi-armed bandits: a review Frontiers Digital Health, https://doi.org/10.3389/fdgth.2025.1435917
- Zanchi B, Monachino G, Fiorillo L, Conte G, Auricchio A, Tzovara A, Faraci F, Synthetic ECG signals generation: A scoping review (2025) Computers in biology and medicine, https://doi.org/10.1016/j.compbiomed.2024.109453

- Zanchi, B., Monachino, G., Faraci, F. D., Metaldi, M., Brugada, P., Sarquella-Brugada, G., Behr, E. R., Brugada, J., Crotti, L., Belhassen, B., Conte, G. (2025). Synthetic electrocardiograms for Brugada syndrome: from data generation to expert cardiologists evaluation. European Heart Journal - Digital Health. https://doi.org/10.1093/ehjdh/ztaf039
- Deng, Y., Denecke, K. (2025). A Comparison of Potentials and Limitations of Transformer Models for Aspect-Based Medical Sentiment Analysis. In: Cafolla, D., Rittman, T., Ni, H. (eds) Artificial Intelligence in Healthcare. AIiH 2025. Lecture Notes in Computer Science, vol 16038. Springer, Cham. https://doi.org/10.1007/978-3-032-00652-3-25
- Pelentrinou A, Cataldi J, Zubler F, Iten M, Haenggi M, Ben-Hamouda N, Rossetti AO, Tzovara A, De Lucia M (2025), Complex auditory regularity processing across levels of consciousness in coma: Stage 1 Registered Report. Brain Communications, https://doi.org/10.1093/braincomms/fcae466
- Quiroga-Martinez RD, Fernandez-Rubio G, Bonetti L, Achyutuni KG, Tzovara A, Knight RT and Vuust P, (2024), Decoding reveals the neural representation of held and manipulated musical thoughts. PLoS Biology, https://doi.org/10.1371/journal.pbio.3002858
- Göktepe-Kavis P, Aellen FM, Cortese A, Castegnetti G, de Martino B, Tzovara A, (2024) Context changes retrieval of prospective outcomes during decision deliberation, Cerebral Cortex, https://doi.org/10.1093/cercor/bhae483
- Tzovara A, Andrillon T, Preller KH, Sarasso S (2024) New trends in the empirical study of consciousness: Measures and mechanisms European Journal of Neuroscience, https://doi.org/10.1111/ejn.16606

4 Communication and Distributed Systems Group

4.1 Personnel

Head

Prof. Dr. T. Braun	✉ torsten.braun@unibe.ch ☎ +41 31 684 4994

Office Manager

P. Grunder	✉ priska.grunder@unibe.ch ☎ +41 31 684 8681

Scientific Staff

Dr. E. Samikwa	✉ eric.samikwa@unibe.ch ☎ +41 31 684 6692
Dr. A. Di Maio*	✉ antonio.dimaio@unibe.ch ☎ +41 31 684 6687
J. Ajayi	✉ jesutofunmi.ajayi@unibe.ch ☎ +41 31 684 6686
J. Chen*	✉ jinxuan.chen@unibe.ch ☎ +41 31 684 6685
C. Gao*	✉ chuyang.gao@unibe.ch ☎ +41 31 684 6689
E. Hashemi Nezhad*	✉ elham.hasheminezhad@unibe.ch ☎ +41 31 684 6682
S. Banik	✉ sabyasachi.banik@unibe.ch ☎ +41 31 684 6693 (as of 01.01.2025)
S. Norouzi*	✉ sajedeh.norouzi@unibe.ch ☎ +41 31 684 6684
I. Núñez Carrillo	✉ ivonne.nunez@unibe.ch ☎ +41 31 684 6681
M. Sun*	✉ mingjing.sun@unibe.ch ☎ +41 31 684 6692
S. Wassie*	✉ solomon.wassie@unibe.ch ☎ +41 31 684 6683
F. Marggi	✉ fabrice.marggi@unibe.ch

	☎ +41 31 684 6665
B. Martins	✉ bruno.martins@unibe.ch
	(01.01.2025 to 30.06.2025)
H. Xing*	✉ hexu.xing@unibe.ch
	☎ +41 31 684 6694
Z. Xu*	✉ zimu.xu@unibe.ch
	☎ +41 31 684 6695
D. Xenakis*	✉ dimitrios.xenakis@unibe.ch

External Ph.D. Students

P. Hammler	✉ patric.hammler@unibe.ch
Y. Farhat	✉ yamshid.farhat@bkw.ch

Joint Supervision Degree Ph.D. Students

L. De Sousa Pacheco	✉ lucas.pacheco@unibe.ch
M. Lopes	✉ marcio.moraeslopes@unibe.ch
	(as of 01.01.2025)

*With financial support from a third-party credit

4.2 Overview

The research group “Communication and Distributed Systems” has been investigating how mobile communication systems and networks can support multimedia and mixed reality applications and cloud computing services with high quality, reliability, and energy efficiency demands. Moreover, we are investigating localization mechanisms for wireless devices and new Future Internet paradigms such as Information-Centric Networking (ICN) and the Internet of Things (IoT). Distributed and Federated Machine Learning are emerging approaches for mobility prediction, mixed reality, and IoT. Emerging applications such as immersive communications require both high bandwidth and low delay. Those are supported by mobile edge computing and service function chaining. Moreover, we are investigating how Federated Learning (FL) can be used for massive Multiple Input Multiple Output (MIMO) communications in 6th generation mobile networks (6G). We are also researching the integration of (distributed) ML in 6G network architectures.

4.3 Research Projects

Efficient Distributed and Federated Machine Learning for Internet of Things

The integration of machine learning models within the Internet of Things (IoT) ecosystems presents significant challenges, including the expensive execution of deep neural networks (DNNs) on resource-constrained devices, ensuring privacy in data-sensitive applications, and managing the dynamic and heterogeneous nature of data and resources available in IoT devices. This research tackles these challenges by proposing a set of novel distributed machine-learning strategies, each designed to optimize different aspects of machine-learning workflows in IoT systems. These strategies include Early Exit of Computation (EEoC), Distributed Micro-Split Deep Learning (DISNET), Adaptive Resource-Aware Split-learning (ARES), and Dynamic Federated Split Learning (DFL), offering solutions to manage inference and training processes across heterogeneous IoT networks efficiently. Through these approaches, the research contributes to the emergence of the Artificial Intelligence of Things (AIoT), where intelligent decision-making is integrated with IoT infrastructure.
EEoC introduces an adaptive mechanism for optimizing DNN inference tasks by allowing early exits based on the computation intensity and resource availability, thereby reducing latency and conserving energy on IoT devices. DISNET extends this efficiency to a broader scale by implementing a micro-split deep learning approach that facilitates flexible, distributed, and parallel execution of DNN tasks, ensuring minimal inference latency and reduced energy consumption while maintaining accuracy across heterogeneous IoT devices. Building on the foundation of efficient inference, ARES introduces a dynamic, resource-aware approach for the distributed training of machine learning models specifically tailored for edge IoT environments. It significantly accelerates local training processes in resource-constrained devices and minimizes the effects of slower devices on global model performance through device-targeted split points while adjusting to time-varying training conditions. DFL complements ARES by enhancing the Federated Learning framework to accommodate IoT devices' and training data's heterogeneity and dynamism. It optimizes training efficiency through a resource-aware federated approach with similarity-based clustering, addressing the heterogeneity of training data and resources.

Research staff: E. Samikwa, T. Braun.

Mitigating E2E Latency for Future Mobile VR Applications

The next generation of Virtual Reality (VR) applications is expected to provide advanced experiences through Six Degrees of Freedom (6DoF) content, which requires higher data rates and ultra-low latency.
Our work on Flow Latency and Throughput Aware Sensitive Routing (FLATWISE) introduces a novel intra-domain routing algorithm with throughput guarantees for minimizing the overall end-to-end (E2E) latency for all flows deployed in the network [Medeiros et al., 2025]. We investigate the Joint Flow Allocation (JFA) problem to find paths for all flows in a network such that it determines the optimal path for each flow in terms of throughput and latency. The JFA problem is NP-hard. We use a mixed integer linear programming to model the system, along with the FLATWISE heuristic, which is one order of magnitude faster than optimally solving the JFA problem. FLATWISE provides an adaptive routing approach that can squeeze or relax the path calculation based on the E2E latency requirement of 6DoF VR applications. The primary characteristic of FLATWISE is to approximate the E2E latency of the calculated path with the E2E latency required by each 6DoF VR application by analyzing the impact of path assignment on other VR flows. Simulations demonstrate that FLATWISE significantly reduces flow latency, over-provisioned latency, E2E latency, and algorithm execution time.

Research staff: A. Medeiros, T. Braun.

Intelligent Orchestration for Scalable Systems

Our first research field is dedicated to orchestrating the virtualized network infrastructure for 5G and beyond. The primary challenge in this area is managing a shared physical infrastructure that must simultaneously support services with highly heterogeneous requirements, such as high-speed mobile broadband and ultra-reliable low-latency communications. We leverage technologies like Network Function Virtualization (NFV) and Software-Defined Networking (SDN) to address this and create dedicated network slices. However, the real-time orchestration of these slices is a complex optimization problem, traditionally solved with heuristic or static algorithms that struggle with scalability and uncertainty. To overcome these limitations, we developed HELIOS [Ajayi et al., 2025], a hierarchical learning approach for network slice provisioning. HELIOS employs a Hierarchical Multi-Armed Bandit model to learn an optimal placement policy that maximizes the acceptance rate of service requests while minimizing

average resource utilization. This design allows for exploring the vast solution space and exploiting learned patterns efficiently. Our results show that this intelligent approach can admit over 25% more slice requests in certain scenarios compared to baseline methods, demonstrating a significant step forward in automated network management. The trade-off for this performance gain is a short initial learning phase, which we have minimized to ensure practical deployment viability.
Complementary to this, our second research field focuses on designing large-scale computing architectures to support and optimize complex distributed applications, particularly Distributed Machine Learning (DML). Modern computing has evolved into a Cloud Continuum (CC), a vertical integration of layers from IoT and Edge devices to the Cloud. To scale beyond a single provider, the Cloud Federation (CF) concept enables horizontal collaboration between different domains. In this area, we have identified a significant research gap: the lack of a unified architecture that integrates both CC and CF to optimize DML execution. Our research is focused on designing and validating such an integrated architecture. The proposed framework includes i) dedicated modules for adaptive resource management, using a Deep Reinforcement Learning (DRL) strategy; ii) a business model based on a Zero-Trust Architecture (ZTA) to govern secure, federated interactions; iii) a context-awareness framework to translate high-level application and legal constraints into machine-enforceable rules, leveraging fuzzy logic to reason with imprecise information; and iv) a system performance module to optimize the trade-off between energy consumption and application scalability, which is addressed using metaheuristic-based optimization. This comprehensive approach ensures that distributed AI applications can scale effectively, securely, and efficiently across multiple collaborative domains.

Research staff: J. Ajayi, M. Lopes, T. Braun.

Indoor and Outdoor Localization

In the era of ubiquitous connectivity, accurate indoor positioning remains a significant challenge. Our research addresses this by developing innovative cooperative localization techniques that harness the collective capabilities of wireless sensor networks. We aim to create accurate and scalable solutions that can operate effectively in complex environments, paving the way for advanced location-based services across various applications.
We present a series of works based on our initial publication, Anchor-free Ranging-Likelihood-based Cooperative Localization (ARLCL).This founda-

tional work introduced ARLCL as a novel, anchor-free, and technology-agnostic localization framework that utilizes inter-exchanged ranging signals for simultaneous positioning. Building on this foundation, our subsequent research addressed the persistent challenge of flip ambiguity in cooperative localization techniques.
These occur when local estimations are correct, but the overall network structure is misaligned due to relative positioning errors. Our upcoming paper, "Iterative Inlier Selection (IIS) for Flip-Ambiguity Mitigation in Cooperative Localization," introduces a corrective method based on Iterative Inlier Selection. Our evaluation demonstrated IIS's effectiveness, with median performance increases exceeding 50% and error reductions of up to 85% for larger swarms. This work pioneers the application of inlier/outlier analysis to improve cooperative localization systems.
Further expanding on ARLCL, we developed a fully distributed version of the algorithm. This adaptation maintains consistent performance as network size increases, demonstrating linear scalability crucial for real-time operations in extensive networks. The distributed approach enhances privacy by minimizing data exchange between nodes, making it suitable for applications where location confidentiality is paramount.
Recognizing the need for comprehensive evaluation resources, we also created "TWR-CLOUD Bern: TWR-based Cooperative Localization with an Open UWB Dataset from Bern-University [Xenakis et al., 2025]," submitted to *IEEE Data Descriptions*—the leading IEEE journal focused on datasets—where it is currently under review. This dataset, the most extensive for UWB mesh networks to date, involves 40 nodes across a 500 square meter area, complemented by LiDAR-based ground-truths. A dedicated webpage for this open-access dataset and its documentation has also been developed and released at https://twr-cloud.inf.unibe.ch, facilitating transparent evaluation and standardization of method comparisons in the field.
Our latest work introduces sequential position Optimization using the Likelihood-based Inferred Distances (SOLID) framework. SOLID addresses the practical challenges of deploying scalable and cost-effective location-based services in the expanding Internet of Things landscape. It supports technology-agnostic ranging models and offers three positioning modes: anchors (SOLID-A), cooperation (SOLID-C), and their combination (SOLID-AC). Evaluated using 40 UWB nodes across 500m², SOLID demonstrated significant accuracy improvements through cooperation, with enhancements near 500% in common deployments. These interconnected research efforts represent a comprehensive approach to advancing cooperative localization techniques, from fundamental algorithms

to practical implementations and evaluation frameworks. Our work contributes to developing more accurate, efficient, and versatile positioning solutions for diverse wireless sensor network applications.

Research staff: D. Xenakis, T. Braun.

Federated Learning in Vehicular Networks

Beamforming techniques use massive antenna arrays to formulate narrow Line-of-Sight signal sectors to address the increased signal attenuation in millimeter Wave (mmWave). However, traditional sector selection schemes involve extensive searches for the highest signal strength sector, introducing extra latency and communication overhead. This research introduces a dynamic layer-wise and clustering-based federated learning (FL) algorithm for beam sector selection in autonomous vehicle networks called enhanced Dynamic Adaptive FL (eDAFL) [Pacheco et al., 2025]. The algorithm detects and selects the most important layers of a machine learning model for aggregation in the FL process, significantly reducing network overhead and failure risks. eDAFL also considers an intra-cluster and inter-cluster approach to reduce overfitting and increase the abstraction level. We evaluate eDAFL on a real-world multi-modal dataset, demonstrating improved model accuracy by approximately 6.76% compared to existing methods, while reducing inference time by 84.04% and model size up to 52.20%.

Research staff: L. Pacheco, T. Braun

Massive MIMO

Massive Multiple Input Multiple Output (mMIMO) antenna systems will be an important technology for future generation, e.g., 6G, mobile telecommunication networks to meet emerging throughput, reliability, and latency requirements. Directional transmission becomes increasingly important for exploiting spatial multiplexing in future mobile and wireless networks. In a mMIMO system, signals are transmitted simultaneously across multiple antenna arrays, allowing the system to steer the directional behavior by adjusting the phases and amplitudes of the transmitted signals. However, the large number of antennas in mMIMO systems makes real-time optimization of these parameters highly complex and computationally demanding. Machine Learning (ML) concepts such as Reinforcement Learning (RL) and Deep Learning (DL) can be used to not only estimate but also predict

channel conditions and to determine parameters for rapid beam forming. Federated Learning (FL) allows ML to be performed in a distributed way. This approach is promising since many user devices can be involved in the learning process, but FL for channel estimation and beam forming is not well explored and understood.

Channel estimation is crucial for enabling efficient beamforming, and machine learning–driven solutions are increasingly gaining traction. While DL–based channel estimation algorithms achieve high accuracy, their significant computational complexity limits real-time deployment. To overcome this challenge, we propose a lightweight channel estimation model based on graph neural networks (GNNs). By integrating domain knowledge, our model adapts effectively to dynamic wireless environments, making it particularly suitable for fast-fading channels and real-time applications, while being 250× lighter than competing approaches [Norouzi et al., 2025]. Furthermore, we leverage federated learning frameworks to allow base stations (BSs) to collaboratively train channel estimation models. This not only reduces training time but also enhances robustness in distributed deployments. In particular, our results show that BSs with stronger estimation capabilities (such as those using more pilot signals and achieving higher estimation accuracy) can effectively assist other BSs, leading to system wide performance gains [Norouzi et al., 2025]. Future work will explore scenarios with severe pilot contamination, particularly in cell-free massive MIMO (mMIMO) systems, where a large number of distributed BSs cooperatively serve users without relying on cell boundaries. While this architecture mitigates inter-cell interference, it is highly susceptible to pilot contamination, making channel estimation especially challenging. To address this, we envision distributed and coordinated training of channel estimation models, naturally aligning with federated learning frameworks to reduce overhead while improving robustness under heterogeneous and dynamic conditions.

Traditional RF-based beam forming for mmWave communication requires estimating channel characteristics, generating appropriate precoding weights, and sweeping through MIMO code elements. Multi-modal beamforming using Federated Learning (FL) can leverage resources like GPS, Lidar, and image data, significantly accelerating beam searching while enhancing data privacy. The heterogeneity of vehicles, however, affects the availability of computing resources for training machine learning models. Moreover, the multi-modal fusion network may contain billions of parameters, leading to extended training time for FL. To address these challenges, we propose a novel Deep Reinforced Federated Split Learning framework (DRFSL) tailored for multi-modal beamforming with differ-

ent sub-model architectures. DRFSL efficiently utilizes MEC computing and adapts the collaborative and distributed training to dynamic network conditions and system heterogeneity by incorporating deep reinforcement learning and split learning with FL. Experimental evaluation using real-world datasets demonstrates that DRFSL minimizes average training time by 49.45% and inference time by 24.43% and can achieve higher accuracy within the same timeframe compared to the existing FLASH framework [Chen et al., 2025a]. However, the accuracy of the FLASH model remains a problem that hinders the benefits of using FL. Thus, we then propose a novel federated attention-based fusion learning framework named FedAttention for multi-modal beamforming in the Internet-of-Vehicle (IoV). FedAttention further improves the model's generalization ability by utilizing the CNN-Transformer architecture and making full use of the MEC servers for the potential federated split learning to enhance efficiency. Based on the real-world datasets, FedAttention achieves 98.16% in Top-5 accuracy and 82.09% in Top-1 accuracy, a 26.86% improvement compared to the current FLASH framework with less wall clock time, showing its training efficiency and robustness [Chen et al., 2025b].

Research staff: J. Chen, S. Norouzi, E. Samikwa, T. Braun.

Financial support: Swiss National Science Foundation (SNSF) [Grant No. 219330]

6G-CLOUD

The 6G-Cloud project aims to design the next generation of mobile network architecture, overcoming the limitations of 5G by introducing a fully service-oriented, cloud-based, and AI-native network design. Unlike previous generations, 6G-Cloud will seamlessly integrate essential network functionalities with advanced control and management frameworks, spanning a multi-stakeholder cloud environment from the extreme edge to central clouds. This marks a significant evolution in mobile architecture, where network functions are no longer confined to isolated layers but are dynamically blended to deliver comprehensive end-to-end (E2E) connectivity. As the boundaries between the radio, core network, and application layers become increasingly fluid, 6G-Cloud will facilitate the integration, customization, and orchestration of network functions across different segments, ensuring continuous, adaptive, and flexible connectivity. By leveraging AI/ML technologies, 6G-Cloud aims to design self-adaptive, AI-driven network functions into a fully cloud-native experience, delivering the scalability, agility, and responsiveness needed to meet future technologi-

cal, economic, and societal demands. 6G will enable a wide range of new applications and use cases that 5G cannot fully accommodate, including enhanced performance and speed, massive connectivity, and the integration of AI and ML for more efficient network management.
This project will develop a service-oriented Radio Access Network (RAN) and Core Network (CN) within a cloud environment, introducing a novel approach to creating an E2E service-oriented 6G system capable of supporting seamless connectivity. Furthermore, it will target service-oriented design, building a flexible control fabric, and designing programmable control methods to coordinate and combine network functions for various network services. New network control and management application types will be developed as cloud-native functions to exploit the programmability of service-oriented network functions. Mechanisms will be created to abstract and expose information from network functions, leading to the development of intelligent NetApps for efficient resource coordination.
Building on service-oriented RAN, we investigated data-driven orchestration for distributed RAN Intelligent Controller (RIC) placement in 6G networks. To address the Controller Placement Problem (CPP) and limitations of centralized orchestration (e.g., single-point failure and scalability), we introduced decentralized orchestration using Multi-Agent Reinforcement Learning (MARL). We then extended previous work and introduced Decentralized Reinforced RAN Intelligent Controller Orchestration for 6G Networks (DERRIC) by allowing each controller agent to optimize RAN parameters, such as user transmission power, based on network observations. Through this, agents determine optimal power allocations within their domains, thereby improving the user packet delivery ratio as a reward signal [Hashemi Nezhad et al., 2025]. Subsequently, we proposed DERRIC Self-Organization orchestration (DERRIC-SO), which is a novel Reinforced Fairness-Aware Multi-Agent Self-Organization framework, where orchestrators autonomously manage duplication, relocation, and termination. This approach jointly optimizes the trade-off between throughput and fairness under time-varying system conditions [Hashemi Nezhad et al., 2025]. Currently, to preserve data privacy in MARL-based controller deployment, we are exploring Federated Reinforcement Learning (FRL), which enhances privacy, resiliency, and scalability in complex and dynamic environments.
Optimal Virtual Network Function placement through data-driven, network state–adaptive service orchestration using Deep Reinforcement Learning methods can autonomously optimize Virtual Network Function deployment and migration across the Cloud Continuum by continuously learning from network dynamics, enabling adaptive Service Function Chain placement

that minimizes latency, reduces migration overhead, and enhances overall network performance in dynamic 6G environments [Wassie et al., 2024]. We extend the previous work to overcome the scalability and visibility limitations of centralized Deep Reinforcement Learning in heterogeneous 6G environments. Distributed Multi-Agent Reinforcement Learning (MARL) can be employed for optimal Virtual Network Function (VNF) deployment and migration across the Cloud Continuum Framework. In this approach, multiple decentralized agents operate on different segments of the continuum spanning near-edge, far-edge, and extreme-edge nodes to locally monitor network conditions and collaboratively learn context-aware policies through experience sharing. This distributed decision-making enables adaptive and service-aware Service Function Chain (SFC) placement and dynamic VNF migration under time-varying traffic, improving deployment efficiency while reducing migration costs and energy consumption [Wassie et al., 2025].

Research staff: S. Wassie, E. Hasheminezhad, E. Samikwa, T. Braun.

Financial support: European Commission, Funding number 101139073

Innovative Electric Network Planning with Advanced Machine Learning Algorithms

The integration of distributed energy resources (DERs) and novel loads like electric vehicles (EVs) or heat pumps (HPs) into local energy systems (LESs), which can function in harmony with the grid, is imperative for achieving a successful transition towards sustainable energy. However, it is vital to consider price signals encompassing system costs and energy strategy objectives when determining the optimal design and operation of LESs. In this work, we aim to evaluate the potential impact of grid tariffs on LESs' optimal solutions. We introduce an initial stage of our research endeavour by formulating a linear optimization algorithm that minimizes LES energy costs. We conducted a case study in a Swiss municipality wherein 85 LESs were subjected to different grid tariff schemes (including volumetric, power peak-based, and seasonal). Our algorithm was employed to identify the optimal LESs designs and evaluate the impact at the municipality level. Based on our findings, it can be inferred that the chosen grid tariff scheme heavily influences factors such as DER & Energy Storage Systems (ESSs) size and type selection, operation procedures of LESs, profitability rates and allocation of costs related to the power grid. We discuss the implications of these findings for network operators, LES owners

and policymakers and provide some recommendations for designing grid tariffs that can promote efficient and fair integration of DERs in LESs to achieve the Swiss Energy Strategy goals and ensure the reliability and cost efficiency of the energy system.

Research staff: Y. Farhat, T. Braun, P. Favaro.

Federated and Distributed Machine Learning for Adaptive and Sustainable Intelligent Systems

Our research explores advanced federated learning approaches to improve the efficiency and adaptability of intelligent systems. We seek to develop frameworks that balance scalability, interpretability, and sustainability by tackling the challenges of heterogeneous data and constrained resources. These efforts converge toward practical solutions that extend the applicability of federated learning in real-world environments.

Our first research in this field focuses on accurate and interpretable forecasting of residential energy demand, a problem increasingly complicated by growing home electrification, the integration of distributed renewable sources, climate variability, and heterogeneous consumption behaviors. Dynamic pricing and peak penalties further increase the need for accurate, privacy-preserving, and robust forecasting solutions that enhance grid efficiency and sustainability. In this context, traditional centralized forecasting models are limited by privacy concerns, temporal variability, and their inability to adapt dynamically, which hinders their scalability and interpretability in real-world settings. To address these challenges, we propose an explainable and personalized Federated Transfer Learning (FTL) framework tailored for multivariate energy forecasting in smart homes. The framework employs a deep learning model to capture long-range dependencies and highlight relevant inputs. Households are clustered by load similarity using dynamic time warping and density-based methods, enabling cluster-specific federated models with reduced variability. The framework incorporates Bayesian optimization for hyperparameter tuning and a weighted aggregation strategy to ensure robust convergence. It supports local fine-tuning to adapt to new or rapidly changing clients while dynamically reassigning clusters to handle behavioral drifts, and provides interpretable feature importance globally and locally via SHAP. Experiments on real-world datasets demonstrate consistent forecasting accuracy and communication efficiency improvements compared to centralized and standard federated learning baselines.

Our second research addresses the challenges of heterogeneous edge

devices constrained by limited computational and communication resources by focusing on developing energy-efficient distributed machine learning (DML) techniques for heterogeneous edge environments, addressing the concerning growing energy demands of AI systems. First, this research examines Federated Transfer Learning (FTL) and Heterogeneous Federated Learning (HFL) for adaptive knowledge transfer across resource-constrained devices with limited labeled data. We introduce a novel model design that includes capability-aware model scaling and adaptive channel attention mechanisms. We aim to improve computational efficiency, communication overhead reduction, and scalability characteristics for large-scale heterogeneous edge deployments. Next, this research explores Vertical Federated Learning (VFL) for feature-partitioned collaborative learning, enabling organizations with complementary datasets to train models jointly without raw data sharing. We plan to develop communication-efficient feature selection and model compression methods for an energy-efficient learning approach. Lastly, this work investigates Federated Split Learning (FSL) for dynamic model partitioning, allowing computationally intensive model portions to execute on capable infrastructure while keeping privacy-sensitive layers closer to data sources. In that sense, we aim to plan adaptive model partitioning strategies that dynamically assign model layers based on available computational power and bandwidth constraints, creating optimized gradient synchronization mechanisms that balance update frequency with communication efficiency, and implement privacy-preserving activation sharing using lightweight encryption methods.

Research staff: I. Núñez Carrillo, S. Banik, T. Braun.

SENTI: Federated Learning for Psychological Monitoring using Smartphone Sensing Data

Digitalization presents both opportunities and challenges for human health and well-being. While smartphone use has the potential to enhance psychological well-being and support healthy behaviour, problematic usage can have adverse effects on physical and mental health. In this context, mobile sensing technology (e.g., through smartphones, smartwatches, and wearables) offers a unique opportunity to reach individuals in their daily lives. It provides indicators for measuring and predicting psycho-social states and behaviours, encompassing affective states, mental health, social interactions, and health behaviours. The SENTI project aims to leverage mobile sensing technology to better understand and pre-

dict changes in health and well-being. Ultimately, its goal is to enable individuals to leverage digital technology to monitor and promote their health and well-being in the long term.
Building on the smartphone application and data collection campaign, we also explore Federated Learning (FL) approaches to address privacy concerns inherent in handling sensitive psychological data. Since psychological monitoring often involves personal and sensitive information, it is essential to safeguard the participants' privacy while still deriving valuable insights from the data. FL offers a promising solution by enabling decentralized data analysis, where data remains on the user's device, and only aggregated model updates are shared.
In this project, we investigate FL deployments for psychological state monitoring in dense settings (hospitals, schools, workplaces) and aim to balance privacy with analytical efficacy. First, we focus on model performance with a personalized FL framework: Global-Local Adaptive personalized Federated Learning (GLAD-FL) decomposes prediction into a shared global component and client-specific adaptations, mixes them via a learned user-level coefficient, and regularizes personalization using a Rademacher complexity guided term to prevent over-fitting on label-scarce users; a personalization-aware participation schedule lowers uplink cost without sacrificing accuracy. Building on this, we address the communication layer for real deployments: Client Selection for Tree All-Reduce Federated Learning (CSTAR-FL) combines stochastic client selection with Tree All-Reduce (TAR) to minimize end-to-end aggregation time under heterogeneous wireless links, with convergence guarantees and up to 40% faster time-to-accuracy in simulations, while maintaining global accuracy [Xu et al., 2025]. Altogether, our goal is accurate, privacy-preserving FL that is both personalization-aware and communication-efficient in high-density environments.

Research staff: Z. Xu, E. Samikwa, T. Braun.

Financial support: Digitization Commission (DigiK) and Faculty of Science of the University of Bern

STEPS: Smart Tourism Evaluation, Predictions and Sustainable Development

The ongoing competition among tourism destinations has resulted in a rapid expansion of offers, leading to an increase in both the number of visitors and the utilisation of local resources, such as landscapes and infras-

tructure. While this development brings economic opportunities, it entails the risk of overstressing both natural and socio-cultural resources, particularly in sensitive alpine environments.
The inter- and transdisciplinary project STEPS establishes a link between computer science and human geography to support sustainable tourism development in alpine destinations, using Lauterbrunnen, Grindelwald and the Aletsch Region as case studies. From a geography perspective, the interplay between resource systems, stakeholder constellations, and governing institutional regimes is analysed with the objective of identifying sustainability challenges and developing indicators to measure the development.
Our research in computer science utilises advancements in big data and machine learning by developing innovative methodologies for the fusion and analysis of heterogeneous data sources, including mobile network data, booking records, and traffic sensors. Our first research proposes a framework of distributed machine learning for collaborative transport predictions among multiples transport providers ensuring privacy. Further our research focuses on developing tools to assist decision-makers in operating, adapting, and planning multiple multimodal transport systems. Specifically, the project aims to create digital twins, virtual models that simulate real-world mobility and its interactions with surrounding factors, including tourist activity and land use. The results are visually synthesised to facilitate evidence-based decision-making and proactive responses to sustainable development challenges.
The project's overarching objective is to provide systemic solutions by combining quantitative data integration, resources analysis and participatory design of institutional interventions. The project's holistic approach aims to both advance scientific understanding and provide actionable tools for stakeholders in Alpine tourism destinations, supporting digitally-enabled sustainable transformation in line with the needs of local communities, businesses and the environment.

Research staff: F. Marggi, T. Braun

Financial support: Digitization Commission (DigiK) and Faculty of Science of the University of Bern

Networking for Immersive Communications (NICO)

The growing need for immersive communications for virtual and augmented reality applications comes with extreme bandwidth, reliability, and latency requirements. NICO addresses these extreme challenges and re-

quirements through various approaches at protocol, network architecture, and implementation architecture levels. In NICO, we intend to come closer to these targets by addressing the following interrelated research challenges, which must be considered in an integrated way.

1. viewport prediction algorithms in 360° video based on sensor-based localization, tracking, and rotation/movement predictions using advanced machine learning (ML) concepts to improve users' Quality-of-Experience (QoE),
2. disruptive methods to identify QoE parameters in immersive communications combining system measurements and psychological behaviours,
3. implementation of an advanced integrated hardware and software platform, including operating system scheduling for user-level virtualization and the support of hardware accelerators for immersive communication processing to support low latency and high throughput,
4. novel mechanisms to minimize latency by (extended) caching mechanisms exploiting in-network processing for immersive communications,
5. novel network protocols based on information-centric networking and network coding to support low-latency communication,
6. original algorithms and mechanisms for service migration to minimize latency between mobile users and service entities, including the novel concept of Floating Services.

We have developed a Multi-Object Tracking algorithm and a comprehensive viewport prediction solution for real-time 360-degree videos. This solution involves a redesigned video-streaming architecture in virtual reality environments, leveraging metadata from video content analysis to optimize bandwidth usage. Experimental results demonstrate that our approach significantly reduces bandwidth consumption and shortens error recovery time, thereby enhancing the QoE for users.
Additionally, we proposed a point-to-point network architecture tailored for VR environments. This architecture identifies the optimal resource retrieval paths for both static resources, which remain unchanged, and dynamic resources, which are generated in real-time. Utilizing peer-to-peer (P2P) communication and resource sharing significantly improves the efficiency of VR video streaming and VR game communication. Our experiments indicate substantial performance improvements in both scenarios [Xing et al., 2025a].
To handle diverse user interactions and dynamic VR network conditions, we developed a novel adaptive learning approach called Personalized De-

centralized Learning (PDL) for Mobile VR networks to enhance VR applications' performance and user experience. PDL leverages decentralized learning paradigms to enable individual VR devices to collaboratively learn from their interactions and environment without relying on a central authority. Each device tailors its learning process by integrating personalisation techniques to accommodate user preferences and contextual factors, resulting in improved content delivery, reduced latency, and enhanced user satisfaction. The proposed method was evaluated through the real-world VR dataset OpenEDS, demonstrating its effectiveness in optimizing network resources and adapting to varying VR scenarios. The results highlight the potential of PDL to transform Mobile VR networks into more resilient and user-centric systems, paving the way for advanced immersive experiences.

We propose an edge caching approach for latency minimization by extended caching mechanisms [Gao et al., 2025]. It combines optimization techniques with a data-driven AI algorithm using real user viewports request trajectory datasets. This approach enables real-time decision-making with reduced computational complexity compared to offline algorithms. It enables more effective utilization of cached content among multiple users and enhancing multi-user video caching. Our method directly predicts the optimal caching strategy, eliminating the need for content popularity comparisons and considering network bandwidth and storage capacity constraints for content replacement. We propose a hierarchical L2 shared caching architecture for 360° video [Gao et al., 2025a]. This architecture introduces an L2 shared cache mechanism to optimize access efficiency and reduce latency. The global shared cache placement problem is formulated as integer linear programming, and a local cache replacement algorithm is designed by adopting a cutting-edge pruning approach. Experimental validation demonstrates the approach's effectiveness in computational efficiency, reducing latency and improving caching efficiency through partial caching.

Research staff: M. Sun, H. Xing, C. Gao, T. Braun.

Financial support: Swiss National Science Foundation Project No. 204447

Multi-Echelon Inventory Optimization

The operation of inventory systems plays an important role in the success of manufacturing companies, making it a highly relevant domain for optimization. In particular, the domain lends itself to being approached via

Deep Reinforcement Learning (DRL) models because it requires sequential reorder decisions based on uncertainty to minimize cost. We evaluated state-of-the-art optimization approaches to determine whether Deep Reinforcement Learning can be applied to the multi-echelon inventory optimization framework in a practically feasible manner to generate fully dynamic reorder policies. We investigated how it performs compared to an optimized static reorder policy, how robust it is when it comes to structural changes in the environment, and whether the use of DRL is safe in terms of risk in real-world applications. Our results show promising performance for DRL with potential for improvement in terms of minimizing risky behaviour.

Research staff: P. Hammler, T. Braun.

Testbeds

The CDS group possesses and operates a cloud infrastructure based on Dell Power Edge Servers. Currently, at the institute, we own five DELL machines: one R320, one R520, two R530, and one R540. These servers support 212 parallel threads (106 cores) and 848 GB of RAM. In addition, we have two dedicated GPU servers with 128 parallel threads (64 cores), 9 GPUs and 256 GB of RAM. We also have a dedicated NVME SSD server with 96 parallel threads (48 cores), storage of 21T, and 250 GB of RAM. We operate two external Dell PowerVault MD3800i that provide us with disk space of 35 TB in Raid 5 and Raid 6. The CDS group maintains a 13-GPU infrastructure for AI training, incorporating devices from the NVIDIA RTX 4090 series through to the NVIDIA A40. The network backbone is based on Dell N4032 switches with 48x10 GbE-T ports and an 80 Gb/s backbone connection. Together with the Lightweight Directory Access Protocol (LDAP) of the institute, our infrastructure provides the members of the CDS group with the following services:

- Mirantis OpenStack 10.0 (IaaS research cloud)
- OwnCloud (shared storage between the CDS members)
- Wiki (information dissemination for the Institute and the CDS group)
- Etherpad (collaborative real-time editor)
- SVN (collaborative version management system)

For administrator purposes, we use

- Teampass as a password management system

For monitoring our infrastructure, we use

- Nagios

The CDS group has its own IoT testbed that consists of:

- 40 MEMSIC Telsob by Crossbow (now Willow) sensors consisting of:
  - Texas Instruments 16-bit microprocessor (TI MSP 430)
  - 802.15.4 radio interface
  - Fixed Power Supply via the USB Interface
  - Temperature, humidity, and light sensor
  - 1 MB external flash

The CDS group has its own testbed for Distributed and Federated Learning in IoT that consists of:

- 8 NVIDIA Jetson nano consisting of:
  - NVIDIA JetPack and Developer Kit
  - NVIDIA Maxwell 128 NVIDIA CUDA nano cores
  - ARM Cortex-A57 MPCore processor
  - Intel Dual Band Wireless-AC 8265 - WiFi/BLE
- 10 Raspberry Pi A+/3B+/Zero consisting of:
  - 1.2GHz Broadcom BCM2837/A57 MPCore processor
  - BCM43438 wireless LAN and BLE on board
  - EMU Check power monitor device with USB data logger
- Lenovo ThinkCentre M90t (edge server) consisting of:
  - NVIDIA GeForce RTX 2060
  - Intel Core i9-10900
  - Wi-Fi 6 / 802.11ax Bluetooth i 5.1

4.4 Theses

4.4.1 Ph.D. Theses

- Pacheco, L. ”Mobility and Cloud Management with Federated and Distributed Learning”, June 2025.
- Samikwa, E. ”Resource-Aware Distributed Machine Learning for Artificial Intelligence of Things”, August 2024. URL: https://boristheses.unibe.ch/5378/

4.4.2 Master Theses

- Simon Parris: ”Decentralized Federated Anomaly Detection in Mobile Networks”, August 2025. (Supervised by: Prof. Dr. Torsten Braun, Dr. Antonio Di Maio,)

4.4.3 Bachelor Theses

- Silas Leuenberger: ”Proactive Data-Driven Node Relocation for Jammed MANETs”, August 2025. (Supervised by: Prof. Dr. Torsten Braun, Dr. Antonio Di Maio)

4.5 Awards

- Best paper award at the *26th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM25, Fort Worth)* for the paper ”Two-Stage Hybrid Edge Caching Framework for 360° VR Video.” (Chuyang Gao, Torsten Braun)
- Best student paper award at the *2024 IEEE/IFIP Network Operations and Management Symposium (NOMS24, Seoul)* for the paper ”Drift-Aware Policy Selection for Slice Admission Control” (Jesutofunmi Ajayi, Antonio Di Maio, Torsten Braun)
- Best Ph.D. Thesis award in the Joint Alumni in Computer Science of the Universities Bern, Fribourg, Neuchâtel (JAACS) 2024 for the PhD thesis entitled ”Resource-Aware Distributed Machine Learning for Artificial Intelligence of Things” (Eric Samikwa)
- Best poster award at the *2025 Bern Data Science Day* for the poster ”Leveraging Digital Technologies to Predict Lapses and their Risk Factors” (Carole Lynn Rüttimann, Dario Baretta, Corina Berli, Radim Lískovec, Zimu Xu, Eric Samikwa, Antonio Di Maio, Torsten Braun, Jennifer Inauen)

4.6 Further Activities

Organized Events

Torsten Braun

- BeNeFri Summer School, Brienz, Switzerland, September 2-4, 2024

- 2nd Workshop on Intelligent IoT Services and Applications, January 20-24, 2025, Universidad Tecnológica de Panamá, Facultad de Ingeniería de Sistemas Computacionales, Centro Regional de Veraguas, Santiago, Panama (general co-chair)
- ACM Sigcomm, August 4-8, 2024, Sydney, Australia (demo and poster co-chair)
- Eight International Workshop on Security, Privacy and Trust in the Internet of Things (SPT-IoT), In Conjunction with IEEE PerCom 2025, March 17-21, 2025, Washington DC, USA (Program Co-Chair)

Invited Presentations

Disclaimer: The invited presentations list only includes invited talks or presentations not considering conference paper presentations.

Torsten Braun

- Communication and Distributed Systems, Jilin University, Changchun, China, November 3, 2024 (invited talk)
- Distributed Machine Learning in the Internet of Things, University of Campinas, Brazil, ICoNIoT Seminar, November 14, 2024 (invited talk)
- Distributed Machine Learning for 6G Networks and Applications, 2nd Workshop on Intelligent IoT Services and Applications, January 20, 2025, Santiago, Panama (invited talk)
- Using digital sensing to predict human health and well-being – SENTI, Bern Data Science Day, May 2, 2025 (project presentation)
- Distributed Sensing and Machine Learning for Urban Comouting, 7th International Workshop on Urban Computing, DCOSS, Lucca, Italy, June 11, 2025, Keynote talk

Memberships

Torsten Braun

- Erweitertes Leitungsgremium Fachgruppe “Kommunikation und Verteilte Systeme” (KuVS), Gesellschaft für Informatik
- Member of KuVS Theses Award committee
- Kuratorium Fritz-Kutter-Fonds
- Expert for Bachelor Theses at Fachhochschule Bern

- Advisory Committee member of the Centre for Informatics and Systems of the University of Coimbra (CISUC)

Editorial Boards

Torsten Braun

- Editorial Board Member of Informatik Spektrum, Springer
- Editorial Board Member of MDPI (Multidisciplinary Digital Publishing Institute) Journal of Sensor and Actuator Networks

Conference Technical Program Committees

Torsten Braun

- 21th IEEE International Conference on Mobile Ad-Hoc and Smart Systems (MASS 2024), September 23 - 25, 2024, Seoul, South Korea
- IEEE Global Communications Conference, 8–12 December 2024, Cape Town, South Africa
- IEEE Latin-American Conference on Communications, 6–8 November 2024, Medellín, Colombia
- IEEE Consumer Communications & Networking Conference, January 10-13, 2025, Las Vegas, NV, USA
- International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT), November 26-28, 2024 Melonares, Gran Canaria, Spain
- 19th Wireless On-demand Network Systems and Services Conference, January 27-29, 2025, Hintertux, Zillertal, Tyrol, Austria
- IEEE International Conference on Communications, June 8-12, 2025, Montreal, Canada
- IFIP/IEEE Networking, May 26-29, 2025, Limassol, Cyprus
- IEEE 11th World Forum on Internet of Things, October 27-30, 2025, Chengdu, China
- IEEE/ACM International Symposium on Quality of Service, July 2-4, 2025, Gold Coast, Australia
- 7th International Workshop on Urban Computing, June 9–11, 2025, Lucca, Tuscany, Italy

- 40th ACM/SIGAPP Symposium On Applied Computing, Catania, Sicily, Italy, March 31 - April 4, 2025
- 36th International Teletraffic Congress (ITC), June 2-5, 2025, Trondheim, Norway
- ACM/IEEE Symposium on Edge Computing, December 4-7, 2024, Rome, Italy
- First Workshop on Native AI in Future Networks Including 6G (NAIFNET), June 23-27, 2025, Budapest, Hungary
- European Wireless 2024, Brno, Czech Republic, Septmber 9-11, 2024
- HES-SO AI days, January, 27-29, 2025, Geneva, Switzerland
- 8th International Workshop on Edge Systems, Analytics and Networking (EdgeSys 2025), in conjunction with ACM EuroSys 2025,

Eric Samikwa

- 39th ACM SIGCOMM Conference (SIGCOMM 2025) Posters and Demos, September 8–11, 2025, Coimbra, Portugal

Project and Person Reviewing Activities

Torsten Braun

- Luxembourg National Research Fund (CORE Panel chair)
- Deutsche Forschungsgemeinschaft (DFG)
- European Science Foundation
- Swiss National Science Foundation, Starting Grants Panel MINT-C
- Research Council of Norway, Researcher Project for Experienced Scientists (FRIPRO)
- Research Council of Norway, Evaluation of mathematics, ICT and technology in Norway (EVALMIT)
- Panel ”Seed Projects for Interdisciplinary Research”, University of Coimbra
- National Fund for Scientific and Technological Development (FONDECYT), Chile
- Zürcher Hochschule für Angewandte Wissenschaften
- Digital Futures Flagship projects, Sweden
- Research Foundation Flanders (FWO), Belgium

Journal Article Reviewing Activities

Torsten Braun

- IEEE Transactions on Network and Service Management
- IEEE Internet of Things Journal
- IEEE Journal on Selected Areas in Communications
- IEEE Transactions on Wireless Communications

Eric Samikwa

- ACM Transactions on Sensor Networks
- IEEE Transactions on Network Science and Engineering
- IEEE Transactions on Green Communications and Networking
- IEEE Transactions on Machine Learning in Communications and Networking
- IEEE Internet of Things Journal
- IEEE Transactions on Mobile Computing
- IEEE Systems Journal
- Neural Networks (Elsevier)
- Computer Networks (Elsevier)
- Cluster Computing (Springer)

Chuyang Gao

- IEEE Journal on Selected Areas in Communications

Jesutofunmi Ajayi

- IEEE/ACM Transactions on Networking
- IEEE Open Journal of the Communications Society

Ph.D. Committee Memberships

Torsten Braun

- Ioanna Angeliki Kapetanidou, Democritus University of Thrace, Greece, December 3, 2024
- Nouran Mohamed Zaghlool Arafat, German University of Cairo, Egypt, February 25, 2025
- Ana Filipa Simao de Almeida, Universidade de Aveiro, Portugal, February, 26, 2025

- Junye Li, University of New South Wales, Sydney, Australia, June 19, 2025

4.7 Publications

Disclaimer: The publication list only includes publications published or accepted during the academic year but does not include submitted papers.

Journal Papers

- Xing, Hexu; Braun, Torsten, "A Hybrid Network for Extended Reality Environments," *IEEE Transactions on Multimedia*, pp. 1–17, 2025, doi: 10.1109/TMM.2025.3590913.
- Gao, Chuyang; Braun, Torsten (2 June 2025). Dual-Engine Intelligent Caching: A Joint Optimization Framework for 360° Mobile VR Video Edge Caching. IEEE Journal on Selected Areas in Communications. IEEE 10.1109/JSAC.2025.3574596.
- Medeiros, Alisson; Di Maio, Antonio; Braun, Torsten (2025). FLATWISE: Flow Latency and Throughput Aware Sensitive Routing for 6DoF VR over SDN. IEEE Transactions on Network and Service Management, 2025. IEEE. doi: 10.1109/TNSM.2025.3600365.
- Xu, Zimu; Di Maio, Antonio; Samikwa, Eric; Braun, Torsten (April 2025). CSTAR-FL: Stochastic Client Selection for Tree All-Reduce Federated Learning. IEEE Transactions on Mobile Computing, 24(4), 3110–3129. IEEE. doi: 10.1109/TMC.2024.3507381.
- Xenakis, Dimitris; Di Maio, Antonio; Braun, Torsten (July 2025). Two-Way-Ranging-based Cooperative Localization with an Open UWB Dataset from Bern-University (TWR-CLOUD-Bern). Bern Open Repository and Information System (corresponding dataset), doi: 10.48620/90893.

Conference Papers

- Gao, Chuyang; Braun, Torsten (28 May 2025). Two-Stage Hybrid Edge Caching Framework for 360° VR Video. In the *26th Inter-*

national Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM25, Fort Worth), pp. 1-10, May 2025, IEEE (10.1109/WoWMoM65615.2025.00015).

- Ajayi, Jesutofunmi; Di Maio, Antonio; Braun, Torsten (14 Oct 2025). Hierarchical Placement Learning for Network Slice Provisioning. In the *50th IEEE Conference on Local Computer Networks (LCN)*, October 2025, IEEE.

- Ajayi, Jesutofunmi; Di Maio, Antonio; Braun, Torsten (2025). Hierarchical Placement Learning for Network Slice Provisioning. In the *IEEE Network Operations and Management Symposium (NOMS)*, 2024, IEEE. arXiv preprint arXiv:2508.06432, 2024.

- Wassie, Solomon Fikadie; Samikwa, Eric; Di Maio, Antonio; Braun, Torsten (2025). MARC-6G: Multi-Agent Reinforcement Learning for Distributed Context-Aware SFC Deployment and Migration in 6G Networks. 2025 21st International Conference on Network and Service Management (CNSM), IEEE, Oct. 2025.

- Hashemi Nezhad, Elham; Di Maio, Antonio; Braun, Torsten (14 October 2025). Hashemi Nezhad, Elham, Antonio Di Maio, and Torsten Braun. Reinforced Fairness-Aware Multi-Agent Self-Organization for 6G Radio Access Network Orchestration. In 2025 IEEE Local Computer Networks (LCN). IEEE, 2025, (https://doi.org/10.48620/88366).

- Wassie, Solomon Fikadie; Di Maio, Antonio; Braun, Torsten (2024). Deep Reinforcement Learning for Context-Aware Online Service Function Chain Deployment and Migration over 6G Networks. In the *Proceedings of the 40th ACM Symposium on Applied Computing (SAC)*, Sicily, Italy. ACM, 2025.

- Hashemi Nezhad, Elham; Di Maio, Antonio; Braun, Torsten (31 March 2025). Data-Driven Orchestration for Distributed RAN Intelligent Controller Placement in 6G Networks. In Proceedings of the 40th ACM/SIGAPP Symposium on Applied Computing (SAC '25), pp. 1371-1373, (https://doi.org/10.1145/3672608.3707972).

- Hashemi Nezhad, Elham; Di Maio, Antonio; Braun, Torsten (24 March 2025). DERRIC: Decentralized Reinforced RAN Intelligent Controller Orchestration for 6G Networks. In 2025 IEEE Wireless Communications and Networking Conference (WCNC), pp. 1-7. IEEE, 2025, (https://doi.org/10.1109/WCNC61545.2025.10978840).

- Norouzi, Sajedeh; Samikwa, Eric; Rahmani, Mostafa; Braun, Torsten; Burr, Alister (2025). Decentralized Federated Learning for GNN-Based Channel Estimation With DM-RS in O-RAN. 2025 IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN), Barcelona, Spain, 2025, pp. 1-7, doi: 10.1109/ICMLCN64995.2025.11140565.

- Norouzi, Sajedeh; Rahmani, Mostafa; Chu, Yi; Braun, Torsten; Chowdhury, Kaushik; Burr, Alister (2025). Lightweight Graph Neural Networks for Enhanced 5G NR Channel Estimation. 2025 IEEE International Symposium on Personal, Indoor and Mobile Radio Communications.

- Chen, Jinxuan; Samikwa, Eric; Braun, Torsten; Chowdhury, Kaushik (2025). FedAttention: Federated Attention-Based Fusion Learning for Multi-Modal Beamforming in IoV. In the *2025 IEEE International Conference on Communications (ICC)*, IEEE, 2025.

- Chen, Jinxuan; Samikwa, Eric; Braun, Torsten; Chowdhury, Kaushik (2025). DRFSL: Deep Reinforced Federated Split Learning for Multi-Modal Beamforming in IoV. In the *2025 IEEE 101st Vehicular Technology Conference (VTC2025-Spring)*, IEEE, 2025.

- Farhat, Yamshid; Braun, Torsten (2025). Proactive Planning of Low-Voltage Networks for the Energy Transition: Benefits Demonstrated in a Large Real Case Study. In the *2025 IEEE Power and Energy Society PowerTech (PowerTech)*, IEEE, 2025.

- Pacheco, Lucas; Braun, Torsten; Chowdhury, Kaushik; Rosário, Denis; Salehi, Batool; Cerqueira, Eduardo (2025). Dynamic Adaptive Federated Learning for mmWave Sector Selection. In the *2025 IEEE 101st Vehicular Technology Conference (VTC2025-Spring)*, IEEE, 2025.

5 Computer Graphics Group

5.1 Personnel

Head
Prof. Dr. D. Bommes ✉ david.bommes@unibe.ch
☎ +41 31 631 3301

Office Managers
D. Heinzen ✉ dragana.heinzen@unibe.ch
☎ +41 31 631 4914
P. Jordan ✉ polina.jordan@unibe.ch
☎ +41 31 684 3342

Scientific Staff
M. Heistermann ✉ martin.heistermann@unibe.ch
☎ +41 31 684 6696
D. Kalmykov ✉ denis.kalmykov@unibe.ch
L. Kamaldinova ✉ liubov.kamaldinova@unibe.ch
Dr. V. Nigolian ✉ valentin.nigolian@unibe.ch
V. Wyss ✉ valentin.wyss@unibe.ch
N. Zhou ✉ ningfeng.zhou@unibe.ch

5.2 Overview

The research activities of the Computer Graphics Group are mainly located in the area of *geometry processing*, which is one of the central topics of *computer graphics*. Geometry processing is concerned with the development of concepts and algorithms to represent, generate, analyze, and modify the shape of objects. Resulting from the physical space we live in, omnipresent classes of shapes include curves, surfaces and volumetric bodies embedded in 3D, or 4D for time-varying shapes. Nowadays, such geometric objects are fundamental in numerous disciplines, inducing a strong scientific impact of geometry processing far beyond computer graphics. Applications as for instance numerical simulation in engineering or computational geology, anomaly detection or surgery planning in medicine, shape matching in computational biology, or the design of smart materials in additive manufacturing (e.g. 3D printing) only become

feasible if accurate geometric representations of the involved shapes are available.

Currently, the group focuses on the generation of discrete geometry representations in the form of semi-structured meshes with quadrilateral elements for surfaces and hexahedral elements for volumetric objects. Such meshes combine the advantages of unstructured simplicial meshes and fully structured Cartesian grids. In contrast to previous methods, e.g. based on local operations, we focus on (global) variational formulations that enable a superior structure of the resulting meshes. There is empirical evidence that following this approach, for the first time algorithms are able to generate meshes that are comparable to manually designed ones. The variational formulation leads to involved nonlinear mixed-integer optimization problems. Hence, one goal of our research is the design of better formulations and parametrizations of the problem that pave the way for efficient solution strategies. In general, our research is driven by the idea of successively addressing the fundamental research questions that are critical from the practitioners perspective, and eventually come up with practically relevant meshing solutions.

5.3 Research Projects

HexHex: Highspeed Extraction of Hexahedral Meshes

Modern algorithms for the automatic generation of high-quality hexahedral meshes are based on the computation of an integer-grid map (IGM), which serves as an implicit representation of the output mesh. This IGM is a piecewise-linear map from the elements of a tetrahedralized input shape to the Cartesian grid, subject to certain conditions which guarantee that its iso-surfaces correspond to a valid hexahedral mesh. A mesh extraction algorithm is required to create the explicit hexahedral mesh implicitly defined by IGM iso-surfaces as a final step.

Existing state-of-the-art algorithms were designed with two goals in mind, i.e., (i) unconditional robustness and (ii) tolerance to map defects in the form of inverted or degenerate tetrahedra (i.e., IGMs which are not locally injective). With continuously improved IGM generating algorithms that can handle significantly higher mesh complexities, we require a downstream mesh extraction algorithm that is not only able to *correctly* extract meshes,

but to do so *efficiently*. The tolerance to invalid inputs was a practical necessity to process outputs of non-robust IGM creation algorithms. However, significant advancements in generating locally injective maps make this goal obsolete. Worse than that, it is now an impediment to improving runtime efficiency and memory use: Algorithms that accept non-injective IGMs must handle complex topological structures that cannot occur in injective IGMs. Data structures suited for this purpose become highly redundant if the input is injective. Restricting the algorithm scope to correctly process only locally injective input IGMs not only avoids algorithmic complexity, but more importantly, allows us to utilize specialized data structures and algorithms that fully utilize IGM properties to eliminate redundant storage and computations.

Consequently, we developed HexHex, a novel (unconditionally robust) hexahedral mesh extraction algorithm for locally injective integer-grid maps designed for maximal performance and scalability. Key contributions include a novel and highly compact mesh data structure based on so-called propellers and a conservative rasterization technique, significantly reducing the number of required exact predicate tests. HexHex not only offers lower asymptotic runtime complexities from a theoretical perspective but also lower constants, enabling practical speedups of 30$\times$ for medium-sized examples and a larger speedup for more complex inputs, specifically when the hex-to-tet ratio is large. We provide an open-source C++ reference implementation, supporting multi-core parallelization and the extraction of curved (piecewise-linear) hexahedral mesh edges and faces, which are useful for subsequent higher-order mesh generation.

Research staff: Tobias Luc Kohler, Martin Heistermann, David Bommes

Financial support: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program, project AlgoHex, No. 853343

Singularity Optimization for High-Quality Interger-Grid Maps

High-quality hexahedral meshes are beneficial for enabling efficient and accurate numerical simulations. By introducing integer-grid maps (IGMs), hexahedral mesh generation can be reformulated as a map optimization

task, which has been investigated more thoroughly in prior research, especially with respect to minimizing distortion.

To generate IGMs, existing frame field methods first construct smooth, meshable frame fields that preserve boundary and feature alignments, and then integrate these fields. Although the frame field is expressive enough to include meshable IGMs, it also introduces complex singularity configurations that may also appear in the final hexahedral meshes. Current state-of-the-art method mainly focuses on improving the robustness and meshability of the pipeline, while the quality of the singularity structure remains a secondary consideration.

Building on the existing methods, our project extends the work by optimizing singularities to enhance the quality of the IGMs and the resulting hexahedral meshes. We propose to refine both the positions of the singularities and their potential insertion or removal, guided by the objective of reducing map distortion. Compared to previous optimization pipelines, our approach includes additional geometric optimization to allow continuous movement of the singularities. This reduce map distortion by increasing the degree of freedom available during optimization. In addition, topological changes are also enabled through local remeshing operations applied around singularities. To ensure the validity of the map, corresponding remeshing operations are applied in both the domain and codomain, preserving injectivity and respecting all feature and singular constraints. These remeshing operations expand the searching space for the singularities positions and provide guidance for their possible insertion and removal. Overall, the proposed pipeline aims to simplify the singularity structure and produce high-quality IGMs and final hexahedral meshes with reduced distortion.

Research staff: Ningfeng Zhou, David Bommes

Financial support: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program, project AlgoHex, No. 853343, and Swiss National Science Foundation Project No. 10004346

Remeshing-Enabled Bijective Tetrahedral Maps Through Non-Conformity

Volumetric mapping is a ubiquitous and difficult problem in Geometry Processing and has been the subject of research in numerous and various directions. From a geometric space, represented digitally as a mesh, we aim at transforming this mesh into another geometric space called the codomain, while maintaining the same connectivity. Our current approach is to temporarily break the codomain mesh' conformity in order to ensure a starting point with a per-element injective, but not locally injective map. Bijectivity is then obtained by recovering the codomain mesh' conformity, through successive local and/or global optimization steps. We conjecture that convergence is always possible under finite mesh connectivity modifications and explore approaches to strategically create the degrees of freedom necessary to generate a valid bijective map. This project has the potential to be applicable on its own, or as a component of a wide variety of cutting-edge methods, such as the Shrink-and-Expand method previously developed by the CGG group."

Research staff: Valentin Nigolian, David Bommes

Financial support: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program, project AlgoHex, No. 853343

Metric Fields for Hexahedral Mesh Generation

In many industrial simulation scenarios, meshing presents significant challenges due to the inherent non-uniformity of the geometry. This is especially problematic when generating hexahedral meshes, as the irregular, non-orthogonal features of such objects are difficult to conform to hexagonal structures. To address this, the concept of g-orthogonal frame fields has been proposed as a foundation for mesh generation. While this technique is well-established for quad meshing, it offers promising potential for further development in the context of meshing 3D geometries.

A frame is conventionally described by six cubic-symmetric unit vectors, which are then split into two components: the metric part g and the cross-field component. The metric part is designed to accommodate the non-uniform features of the model, allowing for targeted refinement where

necessary. This approach introduces two primary tasks: ensuring smoothness of the g-orthonormal frame field and determining the appropriate metric. While the metric selection remains an open problem in this project, the smoothness is measured by the gradient norm, which is further extended to the norm of the covariant derivative under the metric-induced connection. This extension ensures that smoothness is appropriately accounted for in the parameterization. Additionally, boundary alignment conditions are reformulated to take the metric into account.

In the numerical framework, the frames are represented by spherical harmonics, and the optimization problem for smoothness and alignment energies (including a regularization term) is solved in two stages, starting with a tetrahedral mesh as input. Current work is focused on determining the best discretization method for the metric field gradient and tackling the metric selection for different models. Future steps include exploring ways to automate and generalize the algorithm as much as possible.

Research staff: Liubov Kamaldinova, David Bommes

Financial support: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program, project AlgoHex, No. 853343

Globally Meshable Frame Fields for Hexahedral Meshing

Frame field based hexahedral mesh generation is a promising approach to achieve a first practically relevant hexahedral mesh generation algorithm, which is a longstanding and notoriously difficult problem. One remaining fundamental issue results from the fact that frame fields might exhibit topological configurations which do not exists in hexahedral meshes (or integer-grid maps). Hence, it is crucial to warrant *meshability* of the synthesized frame fields. In prior work, we have developed the theory of local meshability, which is necessary but not sufficient to ensure global meshability, i.e. a frame field topology, which is compatible with a hexahedral mesh. The ultimate goal of this project is the development of an algorithm, which converts an arbitrary given frame field into one that is globally meshable and otherwise as similar as possible.

Research staff: Denis Kalmykov, David Bommes

Financial support: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program, project AlgoHex, No. 853343

5.4 Theses

5.4.1 Ph.D. Theses

- Valentin Nigolian, "Robust Tetrahedral Maps — The Shrink-and-Expand Framework", November 2024.

5.4.2 Master's Theses

- Tobias Luc Kohler, "General Shape Reconstruction from Point Clouds using 3D Parametrizations", September 2025.
- Florin Achermann, "Remeshers for Locally Injective Maps", September 2025.
- Luca Schaller, "Pushing Beyond Up to 58 Tets per Hex", September 2025.
- Sandro Scherrer, "In-Situ 3D Visualisation of the Human Anterior Segments Based on Optical Coherence Tomography", September 2025.
- Tobias Verheijen, "Modeling Snow Compaction Bands", February 2025.

5.4.3 Bachelor's Theses

- Lino Meister, "Minimum Convex Cost Flows on Directed and Bi-Directed Networks", April 2025.
- Laura Bünemann, "Laura Bünemann", September 2025.

5.5 Awards

- Best Paper Award at Symposium on Geometry Processing 2024, Hendrik Brückler, David Bommes, Marcel Campen: *Integer-Sheet-Pump Quantization for Hexahedral Meshing*.

- Faculty Teaching Award ALL, Computer Graphics, Autumn Semester 2024.
- Faculty Teaching Award ALL, 3D Geometry Processing, Spring Semester 2025.

5.6 Further Activities

Editorial Boards

David Bommes

- Graphical Models (GMOD) Journal, Associate Editor

Conference Program Committees

David Bommes

- ACM SIGGRAPH 2024
- Symposium on Geometry Processing (SGP) 2024 & 2025
- Geometric Modeling and Processing (GMP) 2025
- Vision, Modeling and Visualization (VMV) 2025

Reviewing Activities

David Bommes

- ACM Transactions on Graphics
- ACM SIGGRAPH Asia conference
- Computer-Aided Design (CAD)
- Computer Aided Geometric Design (CAGD)
- Computer Graphics Forum (CGF)
- International Meshing Roundtable (IMR)
- Replicability Stamp
- UniBE Initiator Grants
- UniBE DocMobility

5.7 Publications

Journal Publications

- Tobias Luc Kohler, Martin Heistermann, David Bommes: “HexHex: Highspeed Extraction of Hexahedral Meshes”, ACM Transactions on Graphics, Volume 44(4), (Proc. **ACM SIGGRAPH**), 2025.
- Ryan Capouellez, Rodrigo Singh, Martin Heistermann, David Bommes, Denis Zorin: “Feature-Aligned Parametrization in Penner Coordinates”, ACM Transactions on Graphics, Volume 44(4), (Proc. **ACM SIGGRAPH**), 2025.
- Berend Baas, David Bommes, Adrien Bousseau: “Shape Approximation by Surface Reuse”, Computer Graphics Forum, Volume 44(5), (Proc. **Symposium on Geometry Processing**), 2025.
- Hendrik Brückler, David Bommes, Marcel Campen: “Integer-Sheet-Pump Quantization for Hexahedral Meshing”, Computer Graphics Forum, Volume 43(5), (Proc. **Symposium on Geometry Processing**), 2024, **Best Paper Award**.
- Yoann Coudert-Osmont, David Desobry, Martin Heistermann, David Bommes, Nicolas Ray, Dmitry Sokolov: “Quad Mesh Quantization Without a T-Mesh”, Computer Graphics Forum, Volume 43(1), 2024.

6 Computer Vision Group

6.1 Personnel

Head
Prof. Dr. P. Favaro ✉ paolo.favaro@unibe.ch
☎ +41 31 684 4451

Office Managers
D. Heinzen ✉ dragana.heinzen@unibe.ch
☎ +41 31 684 4914
P. Jordan ✉ polina.jordan@unibe.ch
☎ +41 31 684 3342

Scientific Staff
P. Acuaviva Huertos ✉ pablo.acuavivahuertos@unibe.ch
L. Cerkezi ✉ llukman.cerkezi@unibe.ch
H. Chihaoui ✉ hamadi.chihaoui@unibe.ch
A. Davtyan ✉ aram.davtyan@unibe.ch
C. Fan ✉ chenrui.fan@unibe.ch
S. Sameni ✉ sepehr.sameni@unibe.ch
A. E. Sari ✉ alp.sari@unibe.ch
L. Scharr ✉ luca.scharr@unibe.ch
V. Shipitsin ✉ viktor.shipitsin@unibe.ch
S. Stapf ✉ sebastian.stapf@unibe.ch

6.2 Overview

Our group explores how machines can learn to **see, understand, and act**. We work at the crossroads of computer vision, machine learning, and artificial intelligence, with the goal of building systems that can make sense of the world from data. For example, we design methods that allow a system to estimate the motion of objects, recover their 3D geometry, or generate navigation instructions based on visual input.
A central theme in our research is **self-supervised learning**, *i.e.*, teaching machines to learn from raw, unlabeled data without requiring human supervision. This approach makes it possible to train models on a very large scale, now an essential practice in both academia and industry. We

also investigate how models can **learn autonomously** or **cooperate with each other**, opening the door to more flexible and adaptive forms of AI.
Beyond these core directions, we pursue a broad range of topics in computer vision and computational imaging. These include **image restoration and computational photography**, which aim to enhance or reconstruct images; **motion estimation and 3D reconstruction**, which help recover dynamic and spatial structure; **image segmentation**, which allows systems to separate objects in a scene; and **optimization methods**, which provide the mathematical tools to make all of this possible.
Together, these efforts are driven by a common vision: creating intelligent systems that can interpret complex data, interact with their environment, and ultimately support new applications in science, technology, and everyday life.

6.3 Research Projects

Building Vision-Language Models on Solid Foundations with Masked Distillation

Recent advancements in Vision-Language Models (VLMs) have marked a significant leap in bridging the gap between computer vision and natural language processing. However traditional VLMs trained through contrastive learning on limited and noisy image-text pairs often lack the spatial and linguistic understanding to generalize well to dense vision tasks or less common languages. Our approach Solid Foundation CLIP (SF-CLIP) circumvents this issue by implicitly building on the solid visual and language understanding of foundational models trained on vast amounts of unimodal data. SF-CLIP integrates contrastive image-text pretraining with a masked knowledge distillation from large foundational text and vision models. This methodology guides our VLM in developing robust text and image representations. As a result SF-CLIP shows exceptional zero-shot classification accuracy and enhanced image and text retrieval capabilities setting a new state of the art for ViT-B/16 trained on YFCC15M and CC12M. Moreover the dense per-patch supervision enhances our zero-shot and linear probe performance in semantic segmentation tasks. A remarkable aspect of our model is its multilingual proficiency evidenced by strong retrieval results in multiple languages despite being trained predominantly on English data. We achieve all of these improvements without sacrificing the training efficiency through our selective application of masked distillation and the inheritance of teacher word embeddings.

Research staff: Sepehr Sameni (during an internship)

Financial support: Adobe Inc.

ViDROP: Video Dense Representation through Spatio-Temporal Sparsity

Self-supervised learning (SSL) has revolutionized image processing, but extending its success to video understanding presents unique challenges due to increased data complexity and computational demands. We introduce ViDROP (Video Dense Representation thrOugh spatio-temporal sParsity), a novel SSL architecture for video understanding that combines token dropping and masking strategies. ViDROP achieves the representational power of state-of-the-art methods like DINOv2 while maintaining exceptional computational efficiency. Our approach eliminates the need for a decoder and enables per-patch loss computation, overcoming limitations of previous video SSL methods. We propose a simple yet effective video compression technique using k-means clustering in pixel space, significantly accelerating data loading and facilitating rapid experimentation. ViDROP leverages pretrained SSL checkpoints and adapts them for optimal linear probing performance without heavy data augmentations. Our method demonstrates scalability across different network sizes and training regimes. Extensive experiments show that ViDROP achieves state-of-the-art performance on various video understanding benchmarks, including Kinetics400, SSv2, UCF101, and HMDB51. Additionally, our approach excels in temporal action detection (THUMOS14) task. These results highlight the effectiveness of our fine-grained token-level learning strategy in a domain traditionally dominated by fine-tuned SSL models, while maintaining computational efficiency and scalability.

Research staff: Sepehr Sameni, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 188690, and Adobe Award.

Boosting Unsupervised Segmentation Learning

We present two practical improvement techniques for unsupervised segmentation learning. These techniques address limitations in the resolution

and accuracy of predicted segmentation maps of recent state-of-the-art methods. Firstly, we leverage image post-processing techniques such as guided filtering to refine the output masks, improving accuracy while avoiding substantial computational costs. Secondly, we introduce a multi-scale consistency criterion, based on a teacher-student training scheme. This criterion matches segmentation masks predicted from regions of the input image extracted at different resolutions to each other. Experimental results on several benchmarks used in unsupervised segmentation learning demonstrate the effectiveness of our proposed techniques.

Research staff: Alp Eren Sari, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 188690

Unsupervised Novel View Generation with Generative Model

In this project, we aim to generate novel views of an object given only a single view of that object. We train our model on a collection of images in the wild of the same category but without information about the pose of the main object in the view and without multiple views of the same object instance. More specifically, our model is a pose-conditioned generative model. To obtain the pose encodings we leverage the emerging properties of recent self-supervised image representations, namely DINOv2. First, we center and rescale the objects in the images given the segmentation masks obtained by finding the main direction of variance in the DINOv2 feature space. We then cluster the dataset based on the Euclidean distance in the space spanned by the first 3 PCA components of the DINOv2 patch embeddings.
Finally, a conditional diffusion model is trained to generate images given the pose. At the inference time, we generate consistent novel views for all pose labels by adjusting self-attention values in the diffusion model.

Research staff: Llukman Cerkezi, Aram Davtyan, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 188690

Zero-shot Image Restoration via Diffusion Inversion

In this project, we present a novel zero-shot framework for solving Image Restoration (IR) tasks, including but not limited to image super-resolution, inpainting, compressed sensing, and blind deconvolution. Our method uses a pre-trained off-the-shelf diffusion model as an unbiased generative prior of the restored image, without requiring any extra training or network modifications. In contrast to prior work, we parameterize restored images as a deterministic function of the input noise in the diffusion model. To mitigate the substantial computational cost associated with inverting a fully unrolled diffusion model, we leverage the inherent capability of these models to skip ahead in the forward diffusion process by arbitrary time steps.

Research staff: Hamadi Chihaoui, Abdelhak lemkhenter, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 200304

When Self-Supervised Pre-training Meets Single Image Denoising

In this project, we present a self-supervised pre-training scheme for single image denoising based on a novel pretext task. Our work is inspired by the success of self-supervised learning (SSL) methods in transfer learning. These methods have been shown to be extremely effective when used to pretrain a model that is then fine-tuned on small datasets. As pretext task, we propose to train a denoising network on patches of the downsampled input image, which we treat as pseudo-clean image patches, and an adaptive noise estimator to learn the specific noise distribution of the input image. By carrying out the pre-training on the single input image, rather than on a separate dataset, we avoid the well-known noise distribution gap between images in the training dataset and the single input image used at test time.

Research staff: Hamadi Chihaoui, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 200304

Optimization of Neural Networks from Kalman Filtering Perspective

We propose KOALA++, a scalable Kalman-based optimization algorithm that explicitly models structured gradient uncertainty in neural network training. Unlike second-order methods, which rely on expensive second order gradient calculation, our method directly estimates the parameter covariance matrix by recursively updating compact gradient covariance products. This design improves upon the original KOALA framework that assumed diagonal covariance by implicitly capturing richer uncertainty structure without storing the full covariance matrix and avoiding large matrix inversions. Across diverse tasks, including image classification and language modeling, KOALA++ achieves accuracy on par or better than state-of-the-art first- and second-order optimizers while maintaining the efficiency of first-order methods.

Research staff: Aram Davtyan, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 10001278

Building a Compositional Communication Channel through Learning Representation Dynamics

This project seeks to establish a novel framework for enabling communication between homogeneous deep learning models, such as autoencoders and world models. Current machine learning systems are typically trained in isolation, limiting their ability to exchange and integrate knowledge across domains. Our approach addresses this limitation by constructing a communication channel in the shared representation space. We propose to learn this channel through autoencoding the data and learning generative denoising models of the latents, formalizing communication as the progressive refinement of a world representation initialized with noise. Compositionality, i.e. the ability to integrate information from multiple sources, is achieved through a product-of-experts formulation, allowing models to collaboratively build richer, more accurate representations from multiple observations. The proposed framework will be evaluated on diverse image and video datasets spanning faces, robotics, and other domains, demonstrating its capacity to generalize across tasks and modalities. By enabling structured exchange of latent knowledge between

models, this work has the potential to advance multi-agent learning, scalable AI systems, and cross-domain knowledge transfer.

Research staff: Aram Davtyan, Sebastian Stapf, Pablo Acuaviva, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 10001278

Adapting Video Diffusion Models for General Vision and Reasoning

This project aims to repurpose Video Diffusion Models (VDMs) into versatile learners capable of tackling both perception and reasoning tasks. Rather than altering their generative interface, we design a lightweight adaptation strategy in which diverse problems are cast as short transition sequences. Training LoRA modules on these transitions enables the frozen VDM to align with new objectives from only a handful of examples. Through this lens, VDMs demonstrate the ability to perform classical computer vision tasks such as segmentation, pose estimation, and style transfer, while also extending to abstract reasoning challenges like ARC-AGI. These results suggest that the dynamics of generative training endow VDMs with structured internal representations that can be unlocked with minimal supervision. Our approach highlights their potential as a foundation for unifying visual generation, perception, and reasoning in a single adaptable framework.

Research staff: Pablo Acuaviva, Aram Davtyan, Sebastian Stapf, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 10001278

E-COVER: Events-Conditioned Video Reconstruction via Flow Matching

Event-based vision is a novel approach to Computer Vision using a non-traditional type of sensor. This field is not mature (no more than 15 years of research have passed since the appearance of the first commercially

available event sensor). The main difference lies in the way the scene is recorded: with the help of an event camera, we can only register the difference in the logarithms of the intensities (illuminations) of each pixel independently. This gives us the ability to record with extremely high dynamic range and frame rate. With qualitative image intensity reconstruction as a preprocessing step, we can use current state-of-the-art algorithms working with conventional images as input and improve the quality of the entire system while maintaining the benefits of event sensors. We propose a novel approach to reconstruct image intensities (in the form of a video sequence) from only event streams. Reconstructing images from events is challenging because multiple image sequences could yield the same sequence of events. We handle the inherent ambiguity of this task by using a diffusion model that learns to generate a distribution of images corresponding to the same events. Another challenge in addressing this task is that events are generated only when objects move or, if they emit light, change their color. Thus, to still take into account such non-moving structures, existing methods use auto-regressive models to accumulate the events history into a single current frame and predict the next frame based on the current frame and the arrival of new events. However, because diffusion models are slow at generating frames, they are not suitable for training in an auto-regressive manner. As a solution, we propose a much simpler, yet effective, diffusion feed-forward model where the history of the events is directly handled by accumulating a fixed number of events, regardless of when they were generated in the past. We call our method **E**vents-**co**nditioned **V**id**e**o **R**econstruction via Flow Matching, or, in short, E-COVER. In order to validate the effectiveness of E-COVER, we conduct extensive experiments on multiple datasets, designed specifically for event-based vision. We demonstrate that E-COVER achieves on-par or even superior results compared to prior work, despite using a simple feed-forward model.

Research staff: Viktor Shipitsin, Paolo Favaro

Financial support: CVG

Self-Supervised Dense Features for Optical Flow and 3D Estimation

Optical flow refers to the task of estimating how every point in an image moves from one video frame to the next. It provides a detailed descrip-

tion of object motion and is a cornerstone of many applications, from autonomous navigation and robotics to video analysis and augmented reality. Closely related is the estimation of 3D structure, where reliable matches across images are essential for recovering depth and scene geometry. In this project, we go beyond directly predicting optical flow and instead focus on learning dense visual features that can be used as building blocks for both motion and 3D estimation. Our approach is based on a novel form of self-supervised learning, which allows models to learn without the need for human annotations. The key idea is to train the model so that regions of an image that are close to each other also have similar feature representations, while regions that correspond across different images maintain consistent representations despite changes in viewpoint, lighting, or object motion. These learned features make it easier to identify which parts of two images belong together, a step that lies at the heart of both optical flow and 3D reconstruction. Once established, such correspondences can be refined into precise motion fields or used to compute depth through multi-view geometry. By creating robust and transferable visual features, this project aims to improve motion estimation and 3D reconstruction in challenging scenarios, while also contributing useful tools for related areas such as image restoration, computational photography, and segmentation.

Research staff: Luca Scharr, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 200304

GEM: A Multimodal World Model for Driving and Ego-Vision Tasks

Autonomous driving and other ego-vision tasks require predictive models that not only replicate the visual distribution of environments, but also provide fine-grained control over agents and objects in the scene. We introduce GEM (Generalizable Ego-Vision Multimodal world model), a generative diffusion-based framework trained on more than 4000 hours of multimodal data across driving, egocentric human activities, and drone flights. GEM predicts future frames from a reference view by conditioning on sparse DINOv2 features, human poses, and ego-trajectories. This enables three complementary forms of controllability: ego-motion control, object-level scene composition (including insertion and motion of objects), and

pedestrian pose manipulation. Beyond RGB, GEM generates paired depth maps to enhance spatial understanding. To support long-horizon prediction, we devise an autoregressive noise scheduling strategy that stabilizes temporal consistency across hundreds of frames. We further introduce the Control of Object Manipulation (COM) metric to assess controllability. Experimental evaluations on Nuscenes and OpenDV show GEM achieving strong FVD improvements, outperforming prior work in both quality and controllability. Our approach demonstrates that multimodal, multidomain world models can serve as scalable backbones for decision-making in safety-critical driving tasks, where precise and interpretable scene control is essential.

Research staff: Sebastian Stapf, Aram Davtyan, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 10001278

Enhancing World Models with Memory for Long-Horizon Consistency

While diffusion-based world models excel at short-horizon video prediction, they degrade over long rollouts due to bounded context windows and stateless training. Inspired by the interplay between memory and imagination in humans, we propose *Memory Adaptation*, a modular test-time training approach that equips pretrained world models with persistent memory. A lightweight adapter network is trained online to reconstruct past frames, encoding them into its parameters as associative memory. During inference, we combine the outputs of the frozen base model and the memory adapter using a probabilistic product-of-distributions guidance scheme. This ensures generated futures remain faithful to both recent and distant past events, even when the past is no longer in the visible context. Our method improves long-term consistency on synthetic and real-world datasets, achieving significant gains in FVD, LPIPS, and semantic recall (DINOSIM). Qualitatively, Memory Adaptation enables models to correctly recall objects or layouts when revisiting scenes, a task where baselines fail. By introducing memory without retraining large diffusion models, our approach offers a scalable path to temporally coherent, interactive video world models, crucial for downstream applications like embodied agents, robotics, and simulation.

Research staff: Sebastian Stapf, Pablo Acuaviva Huertos, Aram Davtyan, Paolo Favaro

Financial support: Swiss National Science Foundation Project No. 10001278

PlantDream: Dual-Guided 4D Plant Modeling via Neural Radiance Fields

In this project, we address the challenging problem of generating realistic 4D plant growth sequences from minimal multi-view static images and text prompts. The key challenge lies in capturing biologically plausible temporal dynamics while maintaining spatial consistency across different viewpoints. Our approach leverages Neural Radiance Fields (NeRF) as the underlying 3D representation due to their superior ability to model complex volumetric structures and handle view-dependent effects naturally. In contrast to existing 4D generation methods that rely solely on Score Distillation Sampling (SDS), we introduce a novel dual-guidance framework that synergistically combines explicit pixel-level supervision with diffusion-based priors. We achieve this by fine-tuning CogVideoX1.5-I2V on a carefully curated dataset of plant growth videos to generate high-quality reference view sequences, while employing SDS for robust cross-view geometry regularization. Our method consists of three main components: (1) canonical NeRF pretraining from multi-view images at t=0 to establish the initial plant structure, (2) time-aware deformation field learning to capture realistic growth dynamics, and (3) joint optimization using both reference-view pixel supervision and novel-view SDS losses to ensure temporal coherence and multi-view consistency. Extensive experiments demonstrate that our framework can model various plant species and growth stages, producing temporally coherent 4D sequences with biologically plausible botanical dynamics that significantly outperform existing baseline methods.

Research staff: Chenrui Fan, Llukman Cerkezi, Paolo Favaro

Financial support: UniBe DigiK2025 “The Human Being in Digital Transformation”, and CVG

6.4 Theses

6.4.1 Ph.D. Theses

- Aram Davtyan, “Learning Object Interactions via Efficient and Controllable Generative Models”, December 2024.

6.4.2 Master’s Theses

- Darya Ardan, “Analysis of the Curvature of the Trajectories in the Concept of Flow Matching for Generative Models”, September 2024
- Timo Blattner, “Reconstructing Highly Folded Cortices A Few-Shot Learning Approach to Investigate Universal Brain Folding”, June 2025
- Chenrui Fan, “PlantNerf: Towards 4D Future Growing Plants Via LDM and Nerf”, June 2025

6.4.3 Bachelor’s Theses

- Kevin Lim, “Developing a UI for Video Inpainting”, October 2024
- Martí Farré Farrús, “Optimization of OmniMotion, a tracking algorithm”, June 2024

6.5 Memberships

Paolo Favaro

- Member of IEEE
- Member of ELLIS
- Member of the Board of the Swiss Center for Augmented Intelligence (SCAI)
- Member of the Swiss AI Initiative

6.6 Further Activities

Evaluation Committes

Paolo Favaro

- SNF Ambizione Committee 2025

- Technion, IIT, Israel 2025

Invited Talks

Paolo Favaro

- “Towards Building Controllable World Models”, AMLD, EPFL, February 2025
- “Towards Building Controllable World Models”, Keynote, AWS ML Workshop, Berlin, September 2025.
- “Advancing Autonomous Learning Models that Learn, Adapt, and Predict from Real Data”, Swedish Symposium on Image Analysis and Deep Learning, KTH, Sweden, March 2025

Presentations

Seminars Given by External Speakers

Timo Blattner, ”Reconstructing Highly Folded Cortices - A Few-Shot Learning Approach to Investigate Universal Brain Folding”, Jun. 27, 2025

Christian-Alexandru Botocan, ”Task Arithmetic for Removing Backdoor Attacks in Multi-Modal Foundation Models”, May 23, 2025

Omri Avrahami, ”Machine Learning Tools for Content Synthesis and Editing”, Apr. 11, 2025

Abdelhak Lemkhenter, ”Learnings from Building Video Game World Models”, Mar. 21, 2025

Conference Program Committees and Reviews

Paolo Favaro

- CVPR 2025 Area Chair
- NeurIPS 2025 Area Chair
- ICCV 2025 Area Chair
- ICLR 2025 Area Chair

Aram Davtyan

- ICLR 2025 Reviewer
- CVPR 2025 Reviewer

- ICCV 2025 Reviewer
- NeurIPS 2025 Reviewer

Publications

- K. Zhang, Y. Deng, J. Ma, P. Favaro, “Adapting Dense Matching for Homography Estimation with Grid-based Acceleration”, in the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2025.
- M. Hassan, S. Stapf, A. Rahimi, P. M B Rezende, Y. Haghighi, D. Brüggemann, I. Katircioglu, L. Zhang, X. Chen, S. Saha, M. Cannici, E. Aljalbout, B. Ye, X. Wang, A. Davtyan, M. Salzmann, D. Scaramuzza, M. Pollefeys, P. Favaro, “GEM: A Generalizable Ego-Vision Multimodal World Model for Fine-Grained Ego-Motion, Object Dynamics, and Scene Composition Control.”, A. Alahi, in the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2025.
- S. Sameni, S. Jenni, P. Favaro, “ViDROP: Video Dense Representation through Spatio-Temporal Sparsity”, in the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) eLVM Workshop, 2025.
- A. Davtyan, L. T. Dadi, V. Cevher, and P. Favaro, “Faster Inference of Flow-Based Generative Models via Improved Data-Noise Coupling.”, in International Conference on Learning Representations (ICLR), 2025.
- A. Davtyan, S. Sameni, B. Ommer, and P. Favaro, “CAGE: Unsupervised Visual Composition and Animation for Controllable Video Generation.”, in AAAI Conference on Artificial Intelligence, 2025.
- H. Chihaoui, A. Lemkhenter and P. Favaro, “Blind Image Restoration via Fast Diffusion Inversion”, in The 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
- H. Chihaoui and P. Favaro, “When Self-Supervised Pre-Training Meets Single Image Denoising”, in IEEE International Conference on Image Processing (ICIP), 2024.

7 Cryptology and Data Security Group

7.1 Personnel

Head
Prof. Dr. C. Cachin ✉ christian.cachin@unibe.ch
☎ +41 31 684 8560

Office Manager
B. Choffat ✉ bettina.choffat@unibe.ch
☎ +41 31 684 8426

Scientific Staff
Dr. J. Alupotha ✉ jayamine.alupotha@unibe.ch
(until 30.04.2025)
Dr. P. Schneider ✉ philipp.schneider2@unibe.ch
(until 31.01.2025)
V. Bammert ✉ vivien.bammert@unibe.ch
(as of 01.11.2024)
M. Barbaraci ✉ mariarosaria.barbaraci@unibe.ch
A. Cimatti ✉ annalisa.cimatti@unibe.ch
(as of 01.11.2024)
D. Lehnherr ✉ david.lehnherr@unibe.ch
M. Senn ✉ michael.senn@unibe.ch
J. F. Villacis Llobet ✉ juan.villacis@unibe.ch
F.-X. Wicht ✉ francois-xavier.wicht@unibe.ch

External Ph.D. Students
Hongyan Sun ✉ hongyan.sun@hslu.ch

7.2 Overview

The Cryptology and Data Security Group broadly investigates security and privacy in a digital world. Concrete topics include cryptographic protocols, distributed consistency, consensus, and cloud-computing security, with applications to blockchains, distributed ledger technology, cryptocurrencies, and their economics.

Security and privacy are at stake in the information society, threatened by the enormous developments in networks, cloud, and mobile. Information technology has already revolutionized many aspects today's life. Finding a balance between the practical convenience of being "always online", current business practices, the changing demands of society, and the privacy and security concerns of individual people represents one of the great open questions of our time. Cryptography and data security provide techniques to answer this question.
Our research spans cryptography, data security, and distributed algorithms, with emphasis on cryptography, secure protocols, and consensus for distributed systems and blockchain networks. Since 2023, we have been a member of *IC3: The Initiative for Cryptocurrencies and Contracts* (`https://www.initc3.org`).

7.3 Research Projects

Consensus protocols

Blockchains are trustworthy distributed networks, maintained by many nodes using distributed a consensus process and relying on cryptographic techniques. Some systems regulate which nodes may participate in the consensus protocol, but others leave this completely open and operate without any central authority. Cryptocurrencies rely on blockchains and hold substantial value.
The participating nodes validate transactions executed by the network and append information to the blockchain, which also takes the form of a ledger. A distributed *consensus protocol* tolerating faults and adversarial attacks ensures that the nodes agree on a unique order in which entries are appended. Advanced cryptographic algorithms play an important role for achieving consistency and privacy. Guaranteeing security and safety for blockchain networks requires mathematical models and sophisticated arguments, drawn from theoretical computer science, cryptography, and the theory of distributed computing.
Based on earlier work on Byzantine-fault tolerant (BFT) consensus for, we are exploring consensus protocols and security mechanisms, and apply them to blockchain systems.
The current research broadly addresses two questions.

Generalized and asymmetric trust. In traditional consensus protocols, all involved nodes adhere to a global, symmetric trust assumption, typically

only expressed through bounds on the number of faulty processes. Most systems deployed today thus count the influence of all nodes in the same way, for example, as when relying on the assumption that a strict 2/3-majority of the nodes are correct. Moreover, all nodes in the system make the same trust assumption.
However, whom one should trust may be a subjective choice. Going beyond this symmetric notion of trust, we are currently exploring models for *asymmetric trust*. Several practical blockchain systems have recently suggested to let each participating node express its own subjective trust assumption. In this model of *asymmetric trust*, every node is free to choose which others it trusts and which ones might collude against it.
In this year we have further refined the understanding of when asymmetric quorum systems offer advantages compared to shared, symmetric quorums. In particular we identified scenarios and protocols, in which the asymmetric model reduces to the symmetric model.
We have also formulated the first DAG-based consensus protocol with asymmetric trust. DAG-structured consensus protocols offer a novel way to structure such protocols and generally perform better than traditional algorithms.

Research staff: Michael Senn, Juan Villacis, Christian Cachin.

Financial support: Swiss National Science Foundation (SNSF), grant agreement Nr. 200020_219403 (Emerging Consensus).
IC3: The Initiative for Cryptocurrencies and Contracts (`https://www.initc3.org`).

Protocols and models for consensus. The consensus protocol used by a blockchain network determines its speed and its security. Whereas the performance can be measured through experimentation, assessing the security requires formal models and analyses. For gaining confidence in a protocol, mathematical insight into their structure must be gained.
We have analyzed several existing protocols and discovered weaknesses in them. We are also interested in developing additional features for such protocols: one prominent aspect is fairness. A blockchain network processing financial transactions should be resistant to front-running, which is a type of insider attack that harms innocent users. An *order-fair consensus* protocol prevents this as much as possible.
In collaboration with LTG, we have studied notions of belief, which are relevant for representing distributed knowledge. Simplicial interpretations of

modal logic have received attention in the literature recently. While notions of (distributed) knowledge have been well investigated in this context, it has been open how to model belief in simplicial models. We have introduced polychromatic simplicial complexes, which naturally impose a plausibility relation on states. We were able to define various notions of belief based on this.
We have also analyzed the Avalanche protocol, used in the Avalanche blockchain, and extended the approach to other models. In particular we have explored probabilistic, gossip-based consensus algorithms that are based on simple local decision rules. We have analyzed their consistency and the time they need to reach stable consensus.

Research staff: David Lehnherr, Philipp Schneider, Michael Senn, Juan Villacis, Christian Cachin.

Financial support: Swiss National Science Foundation (SNSF), grant agreement Nr. 200020_219403 (Emerging Consensus).
Grant from Ava Labs (`https://www.avalabs.org`).

Distributed cryptography

Distributed programs running without trusted coordinator, such as the *smart contracts* executed by a blockchain network, cannot perform cryptographic operations today because no single node can hold a secret key. As one faulty node alone may leak any secrets it knows, keys cannot simply be distributed among the participating nodes. *Distributed cryptography*, also known as *threshold cryptography*, provides well-known methods to secure cryptosystems in the model of distributed and fault-tolerant replicated computations on nodes subject to Byzantine faults.
We have completed the performance analysis of Thetacrypt, a distributed service for threshold cryptography that was developed by the team. Thetacrypt decouples the implementation details of threshold cryptography and the underlying system. It makes cryptographic primitives application independent and promotes seamless integration into various distributed protocols. We have also shown how Thetacrypt can be integrated with the BFT-SMaRT consensus platform and evaluated the throughput of the solution.

Research staff: Mariarosaria Barbaraci, Michael Senn, Christian Cachin.

Financial support: Ripple University Blockchain Research Initiative (UBRI) and donation from Ripple Impact Fund (`https://ripple.com/impact/ubri/`).

Digital currencies for central banks

Cryptocurrencies have shown how to realize a secure equivalent of money in a purely digital way. They challenge the role of traditional currencies, which are issued by central banks. Central banks have therefore started to investigate digital currencies, and many are exploring how to issue a central-bank digital currency (CBDC) to consumers for this purpose. By making such Retail CBDCs accessible to households, the existing relationships between central banks, commercial banks, and people will undergo fundamental changes. This interdisciplinary research project addresses the topic from two perspectives: law and computer science, and runs in collaboration with Mirjam Eggen and team at the Institute for Civil Law.
In this context we have addressed the tradeoff between integrity and privacy for open, blockchain-based payment systems. These systems must ensure, on the one hand, that transactions preserve privacy and do not reveal on which earlier transaction they depend. On the other hand, hiding such information imposes a significant storage overhead due to limitations in current systems, and adds an ever-growing amount of data to the blockchain. In current work we have introduced a privacy-preserving payment scheme that mitigates this problem by structuring the dependency information in a particular way, which ensures that it may be removed again in most situations. This reduces the storage overhead and allows to scale privacy-preserving cryptocurrencies to higher performance.

Research staff: François-Xavier Wicht, Christian Cachin.

Financial support: Office for Digitalization (Digitalisierungskommission, DigiK), University of Bern.

Privacy-preserving identities

Protocols and methods to tackle digital identities are evolving fast. Today it is fundamental to safeguard user privacy and equip service providers with mechanisms enforcing Sybil resistance, i.e., preventing a single entity from showing as many. Current approaches, such as anonymous credentials

and self-sovereign identities, typically rely on trusting that identity providers or identity registries will not to track users.
We have introduced the new cryptographic notion of Anonymous Self-Credentials (ASC) and show two implementations of the concept. ASC enable users to maintain their privacy within an anonymity set while allowing service providers to obtain Sybil resistance. We have formulated User-Issued Unlinkable Single Sign-On (U2SSO), which is an identity management system implemented from ASC that solely relies on a public identity registry to immutably store identities. A U2SSO solution allows users to generate unlinkable child credentials for each SP using only one set of master credentials.

Research staff: Jayamine Alupotha, Mariarosaria Barbaraci, François-Xavier Wicht, Christian Cachin.

Financial support: Academic Research Award from Stellar Development Foundation (SDF), (`https://stellar.org/foundation`).

Erasure codes

Erasure codes facilitate efficient, fault-tolerant storage for large volumes of data across a set of distributed nodes. These codes also reduce the communication overhead in distributed, fault-tolerant broadcast and consensus protocols. Traditional methods assume that all nodes are equally likely affected by errors, such that only the number of missing storage nodes is considered. Modern blockchain networks, however, use more general trust specifications expressed through sets of qualified nodes, also called quorums.
This project develops *generalized erasure codes* and investigates corresponding space-efficient storage and communication protocols. With these codes, the sets of subsets of nodes qualified to reconstruct the data are specified by set systems, and not only by their cardinality.

Research staff: Vivien Bammert, Annalisa Cimatti, Christian Cachin.

Financial support: Academic Research Award from Stellar Development Foundation (SDF), (`https://stellar.org/foundation`).

7.4 Theses

7.4.1 Master's Theses

- "Adapting the PBFT Consensus Algorithm to the Asymmetric Trust Setting." Jonathan Bernhard, Spring 2025.
- "Privacy Preserving Decentralized Identities." Marko Cirkovic, Fall 2024.

7.4.2 Bachelor's Theses

- "Privacy-preserving identity management authentication flow from DIDs." Yanis Berger, Spring 2025.
- "Privacy-preserving credentials with BBS+." Lukas Leuba, Spring 2025.
- "Exploring the Consensus Mechanism of the BNB Smart Chain." Daniele De Jeso, Fall 2024.
- "Balancing Privacy and Public Policy Objectives: A study of eCash applied to the Digital Euro." Joel D. Auerbach, Fall 2024.

7.5 Awards

- Best student paper award at the *32nd International Colloquium On Structural Information and Communication Complexity (SIROCCO 2025)* for the paper "Simplicial Belief" (David Lehnherr).

7.6 Further Activities

Talks

Jayamine Alupotha

- "The Cost of Privacy in Cryptocurrencies" UBRI Connect, Zurich, Switzerland, Sept. 2024.
- "Anonymous Self-Credentials and their Application to Single-Sign-On." Swiss Crypto Day 2025, Università della Svizzera italiana, Lugano, April 2025.

Mariarosaria Barbaraci

- "Enhancing User Privacy in Digital Identity Management through Distributed Ledger Technology." Presented at IC3 Winter Retreat, Engelberg, January 2025.

Annalisa Cimatti

- "Dynamic-FROST: Schnorr Threshold Signatures with a Flexible Committee" Workshop on Cryptographic Tools for Blockchains, Madrid, Spain, May 2025.

Christian Cachin

- "Consensus protocols." Summer School, Deep Dive into Blockchain, University of Zurich, Zurich, Switzerland, July 2025.
- "Consensus in blockchains: Theory and practice." **Keynote Talk**, IEEE International Conference on Blockchain and Cryptocurrency, Pisa, Italy, June 2025.
- "Avalanche consensus." **Keynote Talk** at Lisbon Blockchain Winter School (LxBWS), Lisbon, Portugal, Feb. 2025.
- "What is distributed trust? Some answers and open questions." Presented at IC3 Winter Retreat, Engelberg, Switzerland, January 2025.
- "Asymmetric and heterogeneous trust." WIDE Workshop on Reliable Distributed Systems and Blockchain, INRIA, Rennes, France, Dec. 2024.
- "Asymmetric and heterogeneous trust." Workshop "A Fistful of Distributed Algorithms," Telecom Paris, Palaiseau, France, Dec. 2024.
- "An overview of consensus protocols for cryptocurrencies." UBRI Connect, Zurich, Switzerland, Sept. 2024.
- "Consensus in blockchains: Overview and recent results." a16z crypto Research Seminar, New York, USA, Aug. 2024.

David Lehnherr

- "Simplicial Belief." Workshop, From Complex to Simple, Technical University of Vienna, Vienna, Austria, September 2024.

Michael Senn

- "Structured Asymmetric Quorum Systems and their Applications." Presented at IC3 Winter Retreat, Engelberg, Switzerland, January 2025.

Juan Villacis

- "An Asymmetric DAG-based Consensus Algorithm." ROBUST 2025 Workshop on Resilient Operations – Byzantine Fault Tolerance and

State-Machine Replication, KIT Karlsruhe, Germany, March 2025.

François-Xavier Wicht

- "Toxic Decoys: A Path to Scaling Privacy-Preserving Cryptocurrencies." Privacy Enhancing Technologies Symposium (PETS), Washington D.C., USA, July 2025.
- "Toxic Decoys: A Path to Scaling Privacy-Preserving Cryptocurrencies." Monerokon, Prague, Czech Republic, June 2025.
- "Toxic Decoys: A Path to Private Blockchain Scalability." BSA Conference, EPFL, Lausanne, Switzerland, Mar. 2025.
- "How Privacy Impacts Storage In Cryptocurrency." BSA Conference, EPFL, Lausanne, Switzerland, Mar. 2025.
- "A Scalable Privacy-Preserving Payment Scheme." IC3 Winter Retreat, Engelberg, Switzerland, Jan. 2025.

Societies and Steering Committees

Christian Cachin

- Member of Steering Committee for ACM Conference on Advances in Financial Technologies (AFT), 2019–.

Conference Program Committees

Jayamine Alupotha

- Member of Program Committee for the 25th International Conference on Financial Cryptography and Data Security (FC 2025), Miyakojima, Japan.

Christian Cachin

- Member of Program Committee for 7th Conference on Advances in Financial Technologies (AFT 2025), Pittsburgh (PA), USA.
- Member of Program Committee for Crypto 2025, Santa Barbara, USA.
- Member of Program Committee for 8th Science of Blockchain Conference (SBC 2025), Berkeley (CA), USA.
- Member of Program Committee for 45th IEEE International Conference on Distributed Computing Systems (ICDCS 2025), Glasgow, UK.

- Member of Program Committee for 12th Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC 2025), Rotterdam, The Netherlands.

Philipp Schneider

- Member of Program Committee for 44nd Symposium on Principles of Distributed Computing (PODC 2025), Santa María Huatulco, Mexico.

Reviewing and Thesis Examiner

Christian Cachin

- Member of jury for doctoral exam of Manuel José Ribeiro Vidigueira, School of Computer and Communication Sciences, EPFL, Lausanne (Switzerland), Title: *Techniques for Identifying Elusive Corner-Case Bugs in Systems Software*, 2025.
- *Rapporteur (external referee)* for doctoral thesis of Arthur Rauch, Université de Rennes, Title: *Towards More Scalable and Privacy-Preserving Distributed Asset Transfer Systems*, 2024.
- *Rapporteur (external referee)* for doctoral thesis of Luciano Freitas de Souza, Telecom Paris & Institut Polytechnique de Paris, Title: *Achieving Accountability, Reconfiguration, Randomness, and Secret Leadership in Byzantine Fault Tolerant Distributed Systems*, 2024.
- Examiner for Ph.D. defense of Pouriya Zarbafian, University of Sydney (Australia), Title: *Optimising Order-Fairness in Blockchains using Ordering Linearizability*, 2024.
- Examiner in the Ph.D. defense comittee of Joachim Neu, Stanford University (USA), Title: *Internet-Scale Consensus In The Blockchain Era*, 2024.

7.7 Publications

Journal Papers

- C. Cachin and F.-X. Wicht, “Toxic decoys: A path to scaling privacy-preserving cryptocurrencies,” *Proceedings on Privacy Enhancing Technologies*, vol. 2025, no. 4, pp. 926–943, 2025.
- C. Cachin, D. Lehnherr, and T. Studer, “Synergistic knowledge,” *Theoretical Computer Science*, vol. 1023, p. 114902, 2025.

- A. Cimatti, F. De Sclavis, G. Galano, S. Giammusso, M. Iezzi, A. Muci, M. Nardelli, and M. Pedicini, “Dynamic-FROST: Schnorr threshold signatures with a flexible committee,” *Journal of Mathematical Cryptology*, vol. 19, no. 1, 2025.

Conference Papers

- I. Amores-Sesar, C. Cachin, J. Villacis, and L. Zanolini, “DAG-based consensus with asymmetric trust,” in *Proc. 44th ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 151–161, June 2025.
- C. Cachin, D. Lehnherr, and T. Studer, “Simplicial belief,” in *Proc. Structural Information and Communication Complexity (SIROCCO)* (U. Schmid, ed.), vol. 15671 of *Lecture Notes in Computer Science*, pp. 176–193, Springer, 2025.
- R. de Laage, P. Yuhala, F.-X. Wicht, P. Felber, C. Cachin, and V. Schiavoni, “Practical secure aggregation by combining cryptography and trusted execution environments,” in *Proc. Distributed and Event-based Systems (DEBS)*, pp. 152–163, June 2025.
- Z. Wang, O. Alpos, A. Kavousi, H. W. H. Wong, S. Y. Chau, D. V. Le, and C. Cachin, “DSKE: Digital signatures with key extraction,” in *Topics in Cryptology – CT-RSA 2025* (J. Garcia-Alfaro *et al.*, eds.), vol. 15598 of *Lecture Notes in Computer Science*, pp. 149–173, Springer, 2025.
- M. Senn and C. Cachin, “Asymmetric failure assumptions for reliable distributed systems,” in *Proc. 12th Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC)*, pp. 8–14, Apr. 2025.

Preprints and Other Publications

- I. Amores-Sesar, C. Cachin, and J. Villacis, “Weaker assumptions for asymmetric trust.” e-print, arXiv:2509.09493 [cs.DC], 2025.
- J. Alupotha, M. Barbaraci, I. Kaklamanis, A. Rawat, C. Cachin, and F. Zhang, “Anonymous self-credentials and their application to single-sign-on.” Cryptology ePrint Archive, Paper 618, Apr. 2025.
- M. Barbaraci, N. Schmid, O. Alpos, M. Senn, and C. Cachin,

"Thetacrypt: A distributed service for threshold cryptography." e-print, arXiv:2502.03247 [cs.CR], 2025.

- P. Schneider, "Byzantine fault tolerant protocols with near-constant work per node without signatures." e-print, arXiv:2501.05377 [cs.DC], 2025.

8 Logic and Theory Group

8.1 Personnel

Head
Prof. Dr. T. Studer ✉ thomas.studer@unibe.ch
☎ +41 31 684 3984

Office Manager
B. Choffat ✉ bettina.choffat@unibe.ch
☎ +41 31 684 8426

Scientific Staff
Dr. R. Randrianomentsoa ✉ rojo.randrianomentsoa@unibe.ch
(since 01.07.2025)
D. dos Santos Gomes ✉ djanira.dossantosgomes@unibe.ch
(since 15.04.2025)
A. Feuilleaubois ✉ armand.feuilleaubois@unibe.ch
V. Müller ✉ valentin.mueller@unibe.ch
(since 01.02.2025)
B. Sierra Miranda ✉ borja.sierra@unibe.ch

External Ph.D. Students
L. Zenger ✉ lukas.zenger@unibe.ch

Visiting Researchers:
Dr. S. Horvat (Nov 2024 – March 2025)
Dr. B. Bentzen (Jan 2025)
Dr. M. Mojtahedi (Feb 2025)
Prof. Dr. H. van Ditmarsch (April 2025 – May 2025)
Dr. J. Luo (May 2025)

8.2 Overview

The Logic and Theory Group (LTG) focuses on theoretical computer science and mathematical logic. In particular, we are interested in modal logic and its applications as well as in proof theory and mathematical foundations.

8.3 Research Projects

Modalities in Substructural Logics: Theory, Methods and Applications

Modal logics are a family of formal systems based on classical logic which aim at improving the expressive power of the classical calculus allowing to reason about “modes of truth”. The aim of the present proposal is to put forward a systematic study of substructural modal logics, understood as those modal logics in which the modal operators are based upon the general ground of substructural logics, weaker deductive systems than classical logic. Our aim is also to explore the applications of substructural modal logics outside the bounds of mathematical logic and, in particular, in the areas of knowledge representation; legal reasoning; data privacy and security; logical analysis of natural language.

Research staff: All members of the research group

Financial support: Horizon 2020, MSCA-RISE

Proof and Model Theory of Intuitionistic Temporal Logic

Intuitionistic logic enjoys a myriad of interpretations based on computation, information or topology, making it a natural framework to reason about dynamic processes in which these phenomena play a crucial role. Yet there is a large gap to be filled regarding our understanding of the computational behaviour of intuitionistic temporal logics. The aim of this project is to cement our understanding of intuitionistic temporal logics by developing their model theory based on dynamic topological systems, and their proof theory based on prominent paradigms such as Gentzen-style calculi as well as cyclic proofs.

Research staff: L. Zenger, T. Studer

Non-wellfounded and Cyclic Proof Theory

Fixed points and recursive definitions as well as proofs by induction are essential concepts in mathematics and computer science. In recent years, a novel approach to formal proofs representing such concepts has become popular, namely so-called cyclic proofs. These are not circular, but contain

an infinite number of repetitions of a certain proof figure. These infinite repetitions then lead to a logically valid proof.
Cyclic proofs not only derive propositions from axioms, but also provide a formal counterpart to proofs by infinite descent, i.e. proof branches that are not well-founded but satisfy a global correctness condition.
In this project, we aim at a better understanding of the structural properties of non-wellfounded and cyclic proof systems. We develop syntactic cut elimination methods for (fragments of) the modal mu-calculus. We will also develop novel modal logics for non-monotonic inductive definitions. In the last part of the project we will investigate closure ordinals of positive and non-monotonic modal fixed point operators.

Research staff: A. Feuilleaubois, B. Sierra Miranda, T. Studer

Financial support: Swiss National Science Foundation (No. 214820)

Epistemic group attitudes

We plan to develop and investigate new notions of group knowledge in multi-agent systems. The two classic forms of group knowledge are common knowledge and distributed knowledge. We will devise novel variants of these notions that are important for distributed computing and study their logical and mathematical properties.One part of our project deals with eventual common knowledge (ECK), which is a form of common knowledge that is relevant for the epistemic analysis of many applications, e.g., distributed ledger protocols. However, ECK needs to be better understood from a mathematical and logical perspective. We will develop axiomatic systems for ECK and investigate the meta-logical properties of ECK, like compactness, finite model property, and so on. We will then study the relationship of task solvability and the properties of the corresponding form of ECK. In particular, we will establish the precise connection of ECK and the Firing Rebels with Relay problem, which is a particularly interesting case in distributed computing.The second part of the project is concerned with novel variants of distributed knowledge. Simplicial models are an essential tool in the area of distributed computing. Recently, it has been observed that simplicial models also provide an interesting semantics for multi-agent epistemic logic. These models put the focus on the local states of the agents instead of the global state of the system (represented as a possible world). This simplicial semantics makes it possible to represent new forms of distributed knowledge. We will provide sound and complete

logical systems for these kinds of distributed knowledge. Further, we will use this approach to develop models for distributed data structures.Current simplicial models only support notions of knowledge. However, in the presence of Byzantine agents, knowledge can often not be obtained, only notions of belief. We will adapt simplicial models to support belief and belief dynamics. Further, we will develop and study an epistemic logic for simplicial sets. This is the most general form of simplicial models. Yet, their logical analysis is still missing. We will then also explore the relationship of simplicial sets to belief and belief dynamics, in particular to plausibility models and neighborhood frames.

Research staff: D. dos Santos Gomes, V. Müller, R. Randrianomentsoa, T. Studer,

Financial support: Swiss National Science Foundation (No. 10000440)

8.4 Theses

8.4.1 Bachelor's Theses

- M. Kulowatz: Monitoring Sensing Devices in Medical Environments
- K. Lautenschlager: Feature Model Approach to Mathematical Tasks
- L. D'Arcangelo: Incompleteness in Arithmetic Theories
- A. Muttappillil: Introduction to Circuit Complexity

8.5 Awards

Valentin Müller

- MSc thesis prize of the Dutch Association for Logic and Philosophy of the Exact Sciences (VvL), December 2024

8.6 Further Activities

Editorial Boards

Thomas Studer

- Member of the Editorial Board of Springer book series on Computer

Science Foundations and Applied Logic

Invited Talks

Thomas Studer

- Justification Logic - Introduction and Recent Developments. Seminar talk Zhejiang Universtiy, Hangzhou, China, August 2024
- Simplicial Models for Epistemic Logic. Seminar talk Peking University, Beijing, China, September 2024
- Non-wellfounded and cyclic proofs. Seminar talk Tsinghua University, Beijing, China, September 2024
- Simplicial Complexes for Epistemic Logic. Seminar talk University of Florence, Florence, Italy, November 2024
- Simplicial Models for Epistemic Logic. Seminar talk East China Normal University, Shanghai, China, June 2025
- Strategic Knowledge. Seminar talk Zhejiang University, Hangzhou, China, June 2025

Valentin Müller

- The Method of Higher-Level Rules in the Proof Theory of Inquisitive Logic. Online seminar talk at Annual VvL Seminar, on the occasion of receiving the VvL MSc thesis prize, Eindhoven University of Technology, Eindhoven, the Netherlands, December 2024

Borja Sierra Miranda

- Non-wellfounded master modality: from cut admissibility to cut elimination. Seminar talk National Research University Higher School of Economics, Moscow, Russia, October 2024
- Cut elimination for a non-wellfounded system for the master modality. Seminar talk University of Lisbon, Lisbon, Portugal, December 2024
- La solidez de los razonamientos inválidos: demostraciones circulares y mal fundadas. Seminar talk Instituto de Matemáticas Universidad de Sevilla, Sevilla, Spain, May 2025
- Teoría de la demostración de progreso local y la lógica de la demostrabilidad. Seminar talk Instituto de Matemáticas Universidad de Sevilla, Sevilla, Spain, May 2025
- Local-progress proof theory: admissibility implies eliminability. Seminar talk University of Zagreb, Zagreb, Croatia, March 2025

Technical and Research Committees

Thomas Studer

- Presidium Member of the Platform Mathematics, Astronomy and Physics of the Swiss Academy of Sciences (until December 2024)
- Swiss Delegate to the International Federation for Information Processing Technical Committee 1 (Foundations of Computer Science)
- Swiss Delegate to the International Union of History and Philosophy of Science and Technology (until December 2024)
- Board member of the Swiss Society for Logic and Philosophy of Science
- Member of the Committee for the Promotion of Young Talents (Kommission Nachwuchsförderung) of ScNat
- Member of the Kantonale Maturitätskommission, Hauptexperte Informatik
- Member of the Kommission Gymnasium–Hochschule
- Member of the SNSF Postdoc.Mobility Evaluation Panel MINT-C
- Member of the Steering Committee Workshop on Logic, Language, Information and Computation
- Member of the computer science curriculum group, Canton of Bern

Borja Sierra Miranda

- Board Member of the Swiss Graduate Society for Logic and Philosophy of Science

Organized Events

Thomas Studer

- Logic and Application, Inter University Centre Dubrovnik, Dubrovnik, Croatia, 23-27 September 2024

Borja Sierra Miranda

- Swiss Graduate Society for Logic and Philosophy of Science (SGSLPS) Meeting: Model Theory, University of Bern, Bern, Switzerland, 4 December 2024
- Swiss Graduate Society for Logic and Philosophy of Science (SGSLPS) Meeting: Duality Theory, University of Bern, Bern, Switzerland, 16-17 June 2025

PhD Committee Memberships

Thomas Studer

- *Universality of Büchi Automata by Graph Neural Networks*, Christophe Stammet, University of Fribourg
- *Cut-elimination in non-wellfounded proofs*, Esaïe Bauer, Université Paris Cité

8.7 Publications

- *Jieting Luo*, *Thomas Studer*, and Mehdi Dastani. Tailoring explanations through conversation. Journal of Logic and Computation, Volume 35, Issue 4, Oxford University Press, 2025.
- Christian Cachin, David Lehnherr, *Thomas Studer*. Synergistic knowledge. Theoretical Computer Science, Volume 1023, 2025.
- Atefeh Rohani, *Thomas Studer*. Explicit non-normal modal logic. Journal of Logic and Computation, Volume 35, Issue 6, Oxford University Press, 2024.
- Federico L. G. Faroldi, Meghdad Ghari, Eveline Lehmann, *Thomas Studer*. Consistency and Permission in Deontic Justification Logic. Journal of Logic and Computation, Volume 34, Issue 4, 2024. p.640-664.
- Christian Cachin, David Lehnherr, *Thomas Studer*. Simplicial Belief. In: Schmid, U., Kuznets, R. (editors) Structural Information and Communication Complexity, SIROCCO 2025. Lecture Notes in Computer Science, Volume 15671, Springer, 2025. p. 176–193.
- Christian Cachin, David Lehnherr, *Thomas Studer*. Simplicial Belief. Advances in Modal Logic 2024 Short Paper, 2024.
- Roman Bögli, Atefeh Rohani, *Thomas Studer*, Christos Tsigkanos, Timo Kehrer. Temporal Logics Meet Real-World Software Requirements: A Reality Check. In: Proc. 2025 IEEE/ACM 13th International Conference on Formal Methods in Software Engineering (FormaliSE), 2025. p. 74–85.
- *Borja Sierra Miranda*, *Thomas Studer*, *Lukas Zenger*. Coalgebraic proof translations for non-wellfounded proofs. In: Agata Ciabattoni, David Gabelaia, Igor Sedlár (editors) Advances in Modal Logic 202, 2024. p. 527-548.

9 Pattern Recognition Group

9.1 Personnel

Head
PD Dr. K. Riesen ✉ kaspar.riesen@unibe.ch
☎ +41 31 684 4998

Office Manager
B. Choffat ✉ bettina.choffat@unibe.ch
☎ +41 31 684 8426

Scientific Staff
Dr. A. Gillioz ✉ anthony.gillioz@unibe.ch
K. Dobler ✉ kalvin.dobler@unibe.ch
F. Leonardi ✉ francesco.leonardi@unibe.ch

External PhD Students
B. Fankhauser ✉ benjamin.fankhauser@unibe.ch
C. Masanti ✉ corina.masanti@unibe.ch
F. Jafari ✉ fereshteh.jafari@unibe.ch

9.2 Overview

The Pattern Recognition Group (PRG) conducts research on algorithms and complex data structures in the fields of pattern recognition, machine learning, and data science. The group's expertise is built on two complementary pillars:

- Graph-based representations for intelligent information processing.
- Document analysis using novel methods from natural language processing (NLP) and large language models (LLMs).

With the rapid advancement of storage technologies and data acquisition methods, the volume of available data in science and industry is growing at an unprecedented pace. At the same time, much of this data is inherently complex, making graphs one of the most powerful and suitable data structures for modeling and analysis. In parallel, the increasing importance of unstructured text and documents requires advanced language technologies to unlock and process the knowledge contained within. These de-

velopments highlight the need for robust and efficient methods that help humans to interpret and manage large, complex datasets. The PRG addresses this challenge by developing novel graph-based algorithms for pattern recognition and data science as well as innovative document analysis techniques leveraging NLP and LLMs—always with the goal of providing feasible and reliable solutions that advance both theory and practice.

9.3 Research Projects

AI-Supported Human-Computer Interaction with Large Volumes of Historical Documents

This project is in the area of computational analysis of documents using *machine learning* and *natural language processing* (in particular, *large language models* (LLMs)). It is based on a collaboration between the applicants' research groups and the research center *Diplomatic Documents of Switzerland* (*Dodis*), which is the center of excellence for the study of history of Swiss foreign policy.
The overarching goal of the project is to research novel algorithms to automate and support the process of studying, understanding, selecting, and editing documents at Dodis. In particular, we strive for a unique solution that allows both experts and interested laypersons to have a written conversation with an artificial system which has in turn access to Dodis' large document corpus. We aim for a framework that is composed of two building blocks. First, a *conversational artificial intelligence* (CAI) that accurately answers user questions and follow-up questions and challenges incorrect assumptions, and second a high-quality *summarization* component that provides accurate and useful summaries of documents avoiding hallucinations.
Both fields of research (i.e., summarization and CAI) gained momentum in the last years and elaborated models are available. However, to further push the boundaries of current understanding and implement a solution that is actually usable for Dodis, substantial efforts in research are required. First, although Dodis will contribute an extensive corpus of real-world documents including ground truth transcripts, we face the major challenge that state-of-the-art OCR is prone to errors for historical typewritten documents and is largely unable to extract (implicit) meta-information from the documents that could be valuable for subsequent research (e.g., the document type or recipient). Second, existing (locally executable) models for summarization produce questionable summaries that are hardly

usable for Dodis' research purposes. Moreover, while public LLMs such as GPT provide grammatically and linguistically sound summaries, the results often do not reflect the historically important aspect of the document (moreover, the summaries are still not precise enough, i.e. they contain hallucinations). Third, we observe that open platforms like Dodis are based on the idea of keyword search which are not able to answer specific questions in natural language. On the other hand generic CAI platforms provide very vague or even false answers to specific questions on Dodis documents (as they do not have good access to the necessary documents). It is a major research challenge of this project to enable natural language interactions with a novel CAI system based on the large amount of documents available. The aim is to provide highly accurate and verified answers to specific user questions without hallucinations.
We start our research by selecting and preparing large-scale ground truth training sets for building the envisaged systems. We will first research and empirically evaluate existing models for all tasks and systematically document all limitations (e.g., fabricated facts or similar). Major contribution of the present project is then to develop and research novel approaches that use large knowledge graphs to guide and improve both the summarization and CAI processes. In the experimental phase we use both automatic evaluation metrics and human-based evaluations. Last but not least, the writing of scientific papers and doctoral theses is also part of the present project (two full-time PhD students and one postdoctoral researcher are scheduled).
The significance and impact of the proposed project is manifold (ranging from unique and novel data sets for document analysis to more robust document workflows using LLMs). One of the most important impacts is, however, that the present project might result in radically different services, which can be interpreted as first — yet significant – step towards a more natural human-computer interaction with large archives of historical documents.

Financial support: Swiss National Science Foundation Project No. IC00I0-231354 - 10002280

Research staff: K. Riesen (PhDs and Postdocs TBD)

Neural Networks for Learning Graph Embedding and Graph Matching

This project is concerned with graph-based pattern recognition. Research in this area can be roughly divided into three time periods, viz. the eras of *graph matching*, *graph kernels*, and *graph neural networks*. The overall objective of the present project is to develop and research robust methods which combine the best ideas and methods that emerged from these three eras. This way, we aim to introduce novel graph-based methods that significantly exceed the current state of the art, both in terms of speed and accuracy. Due to their power and flexibility, graphs indeed play a pivotal role in many disciplines in science and industry and the research of efficient and accurate graph-based methods has become a crucial challenge. Therefore, both the relevance and potential impacts of the proposed project can be considered very high. We envisage three relatively independent lines of research that will be investigated by two PhD students and one post-doctoral researcher.

- **Line of Research 1**: We will employ existing graph embedding frameworks, more specifically the resulting vector space embedding, as ground truth for designing, developing, and training of novel graph neural networks. The aim is to learn graph embeddings that minimize the difference of the learned and the ground truth embedding. Major benefit of this procedure is that the embedding of unseen graphs can be accomplishment incomparably faster than with the original techniques.
- **Line of Research 2**: We will learn graph matchings using graph neural networks. To this end, we plan to omit the global pooling layer and directly infer node mappings based on the learned node embeddings. Initial ideas include exploring clustering and/or assignment algorithms that can be applied to the learned node embeddings. The advantage of this method is that it derives explicit mappings of the substructures of the underlying graphs (which is both valuable and often necessary in applications).
- **Line of Research 3**: We plan to explore new (interdisciplinary) applications for graph-based pattern recognition. In particular, we seek to address data challenges where graphs are not currently used, even though they would be the natural approach for formal pattern representation. We believe that this will enable us to define breakthrough and pioneering algorithms in various scientific fields.

Along the three lines of research, several open research questions need

to be answered to push forward the frontier of current knowledge. These questions range from technical issues (e.g., Which architectures are best suited for the target tasks?), over conceptual issues (e.g., Which novel applications are suitable for graph-based methods?), to empirical issues (e.g., To what extent do the proposed methods improve and accelerate the current state of the art in graph-based pattern recognition?).

Financial support: Swiss National Science Foundation Project No. 219388

Research staff: K. Dobler, F. Leonardi, A. Gillioz, K. Riesen

Spatio-temporal graph convolutional networks - a novel deep learning approach to forecasting river temperatures

The Federal Office for the Environment (FOEN) analyses several environmental aspects of Switzerland. The monitoring of water temperatures over long time periods belongs to one of the most important tasks of the Hydrology Division of FOEN. At the moment, the Hydrology Division maintains approximately 80 metering stations. Stations measure various parameters such as water temperature, discharge, water level etc. The sampling frequency is in most cases consistently regulated at a ten minutes interval. Although the initiation date of monitoring varies between 1971 and 2015 as more stations were installed over time, concise water data is available over several decades. Besides the federal metering stations, the cantons of Switzerland maintain more than 700 additional stations. As is often the case in Switzerland, measurement policies vary in the different cantons. Ongoing climatic change and thereby resulting adverse effects to all lifeforms and society are a major concern in almost all countries. Since several years rising river water temperatures are being observed. Hence, several studies have been conducted to model and forecast the river temperatures. As far as we could identify none of the current studies focus on the connectivity of rivers or try to model a large part if not the entire Swiss river (water) network as a collective. We propose a novel deep learning approach to modelling the Swiss river network with a spatio-temporal graph. Spatio-temporal graphs are graph structures where the node and/or edge features are allowed to change over time. Using the temporal information of the spatial graph, temperature for a future time step can be forecasted. The key idea of graph spatial-temporal networks is to consider spatial dependency and temporal dependency simultaneously. In this project we de-

velop novel and robust methods, algorithms, data structures, and heuristics that go beyond current understanding in graph-based machine learning on a concrete physical system that is fundamentally important to our society.

Financial support: Swiss National Science Foundation (Practice-to-Science project in collaboration with Bern University of Applied Sciences)

Research staff: B. Fankhauser, K. Riesen

Detecting Real-Word Errors with Large Language Models: Opportunities and Limitations

Professional proofreading companies are typically focused on orthography, grammar, punctuation, and typography in text documents. Due to linguistic problems (e.g. ambiguities) combined with a specific lack of training data, unresolved challenges for automated proofreading exist. Even with the introduction of transformer-based language models, which significantly advanced research in error detection and error correction, some significant research challenges remain, comprising detection of so-called real-word errors, compliance with specific typesetting rules, or company-specific spelling guidelines. The major aim of a recent research project1 (in which the applicant made a significant contribution as a scientific employee) was to research and solve some of the identified scientific obstacles. One of the main outcomes of this project is a unique data set of annotated text documents and an associated framework for error-detection.
The basis of this data set is given by a unique collection of around 80'000 manually corrected documents (written by native speakers from different industries, such as pharmaceuticals, banking, insurance, retail, communications, and more). Using this corpus, a multilingual data set of erroneous and correct text sentences has been extracted. The primary purpose of this data set is to support the research and development of automated error detection systems, especially in the multilingual setting where high-quality data sets are scarce.
Using this large corpus of text data, the applicant of the present proposal adapted various well-established language models and conducted several thorough experimental evaluations. This analysis assessed the classification accuracy measured over different error types in various settings. Interestingly, while some error types can be detected with an accuracy exceeding 80%, it turns out that the recognition of some specific error types is

still very difficult to solve (even with the most elaborated models available). The applicant also proposed completely novel approaches to systematically incorporate synthetic data into the training process. To this end, ensemble learning methods for language models have been proposed (in particular, an adaptation of the Boosting technique to language model learning). Although the applicant was able to make substantial progress in research and build a reliable and advanced (semi-)automated proofreading framework, there are still many open research questions. First, there is considerable room to further refine the innovative Boosting approach by integrating targeted fine-tuning. Second, we see great potential in exploiting the embeddings of the language models used for classification with traditional machine learning models. Third, we think that prompt engineering with linguistic knowledge could generate a powerful model (e.g., by inserting examples of false positives to the prompt), and we see the opportunity to use relatively simple post-processing rules to correct certain recurring errors that language models typically generate.

Financial support: Hasler Foundation Project No. 2025-02-17-292.

Research staff: C. Masanti, K. Riesen

9.4 Theses

9.4.1 Master's Theses

- M. Streilein: VAGUE: Variational Autoencoder for Graph Understanding and Expansion
- M. Biner: New Systems for Graph Classification: Using learned graph embeddings to obtain support vector machine kernels

9.4.2 Bachelor's Theses

- D. Demir: Towards Reliable Fact-Verification Systems
- L. Hertzberg: Between Phones and Video-Chats – A Computer-Mediated Communication Experiment
- R. Brentani: Die Wirkung von UX-Gesetzen auf die Usability und User Experience interaktiver Webschnittstellen
- D. Rhyn: The Core of Change: Frequent Subgraph Discovery in Software Evolution Graphs

- C. Semelet: BERT for Error Detection: Comparing Fine-Tuning and Embedding Classification
- G. A. Zwahlen: A Reinforcement Learning Approach to Graph Edit Distance Estimation
- N. Neeb: Rule-based Text Correction

9.5 Awards

- Recognition of Outstanding Teaching Achievements for the lecture ”Programmierung 1” in Autumn Semester 2024, among the top 7 courses offered by the Faculty of Sciences at the University of Bern. (Kaspar Riesen)
- Faculty Teaching Award for the year 2024 (Kaspar Riesen)

9.6 Further Activities

Invited Talks

Kaspar Riesen

- Graph-based Pattern Recognition – Latest Developments of the Bern Group – Invited talk at the working group on machine learning with graphs at INSA ROUEN

Editorial Boards

Kaspar Riesen

- Associate editor for Pattern Recognition, 2015–, Elsevier.

Reviewing Activities

Kaspar Riesen

- *Trans. on Pattern Analysis and Machine Intelligence*
- *Pattern Recognition*
- *Journal of Pattern Recognition and Artificial Intelligence*
- *Pattern Recognition Letters*
- *Data Mining and Knowledge Discovery*

- *Parallel Computing*
- *Journal on Computing*
- *Information Sciences*

Conference Program Committees

Kaspar Riesen

- Member of Program Committee for the 18th International Conference on Document Analysis and Recognition 2024
- Member of Program Committee for the 14th IAPR-TC-15 International Workshop on Graph-Based Representations in Pattern Recognition 2025

Talks

Corina Masanti

- ”Boosting Language Models for Real-Word Error Detection” at ICPRAM 2025

Benjamin Fankhauser

- ”Leveraging LSTM Embeddings for River Water Temperature Modeling” at ANNPR 2024
- ”Exploring a Graph Regression Problem in River Networks” at GbRPR 2025

Kalvin Dobler

- ”Learning Graph Matching with Graph Neural Networks” at ANNPR 2024
- ”A Geometric Perspective on Graph Similarity Learning Using Convex Hulls” at GbRPR 2025

Francesco Leonardi

- ”Dissimilarity-Based Graph Embedding: An Efficient GAT-based Approach” at ICPR 2024

Fereshteh Jafari

- ”Predicting Photovoltaic Power Output Using LSTM: A Comparative Study Using both Historical and Climate Data” at ICPRAM 2025

Hannes Thurnherr

- ”Neural Decompiling of Tracr Transformers” at ANNPR 2024

9.7 Publications

Journal Publications

- Anthony Gillioz, Kaspar Riesen:
  Normalized graph compression distance - A novel graph matching framework. Pattern Recognit. Lett. 190: 97-104 (2025)
- Mathias Fuchs, Kaspar Riesen:
  Fast approximate maximum common subgraph computation. Pattern Recognit. Lett. 190: 66-72 (2025)

Refereed Conferences

- Kalvin Dobler, Kaspar Riesen: A Geometric Perspective on Graph Similarity Learning Using Convex Hulls. GbRPR 2025: 103-112
- Benjamin Fankhauser, Vidushi Bigler, Kaspar Riesen: Exploring a Graph Regression Problem in River Networks. GbRPR 2025: 203-213
- Corina Masanti, Hans Friedrich Witschel, Kaspar Riesen: Boosting Language Models for Real-Word Error Detection. ICPRAM 2025: 318-325
- Fereshteh Jafari, Joseph Moerschell, Kaspar Riesen: Predicting Photovoltaic Power Output Using LSTM: A Comparative Study Using both Historical and Climate Data. ICPRAM 2025: 733-740
- Kalvin Dobler, Kaspar Riesen: Learning Graph Matching with Graph Neural Networks. ANNPR 2024: 3-12
- Hannes Thurnherr, Kaspar Riesen: Neural Decompiling of Tracr Transformers. ANNPR 2024: 25-36
- Benjamin Fankhauser, Vidushi Bigler, Kaspar Riesen: Leveraging LSTM Embeddings for River Water Temperature Modeling. ANNPR 2024: 283-294
- Francesco Leonardi, Kaspar Riesen: Dissimilarity-Based Graph Embedding: An Efficient GAT-based Approach. ICPR (10) 2024: 361-374

10 Digital Sustainability Group

10.1 Personnel

Head
PD Dr. M. Stürmer ✉ matthias.stuermer@unibe.ch
☎ +41 76 368 8165

Office Manager
B. Choffat ✉ bettina.choffat@unibe.ch
☎ +41 31 684 8426

Scientific Staff
L. Rolshoven ✉ luca.rolshoven@students.unibe.ch

Other Staff
Dr. B. Hitz-Gamper ✉ benedikt.hitz@unibe.ch
Dr. J. Schaerer ✉ jakob.schaerer@abilium.com

10.2 Overview

Since 2021, PD Dr. Matthias Stürmer is the head of the Institute of Public Sector Transformation in the Business School at the Bern University of Applied Sciences (BFH). He stays at the Institute for Computer Science of the University of Bern for the lecture “Digital Sustainability”. A new SNSF Sinergia project on natural language processing in sustainable public procurement started in April 2024 with a new doctoral student, Luca Rolshoven. Dr. Benedikt Hitz was in charge of the programming course for non-computer science students until end of 2024. Since 2025, Dr. Jakob Schaerer is responsible of the programming course

10.3 Research Projects

Sinergia - Sustainable Public Procurement (SPP)

This four-year interdisciplinary SNSF research project ”How environmental and social public procurement affects sustainability transformation in the

public and private sector: The role of the revision of the public procurement law in Switzerland" is structured in three sub-studies that directly build on each other. In Study 1, the computer science researchers will operationalize the identification of SPP criteria in tender documents using NLP methods. By extracting relevant information on procurement related and bidder related criteria as well as other data out of tender documents, they bring new insights into an enormous corpus of unstructured text data. In study 2, law scholars will use the data extracted in Study 1 and evaluate it in order to assess its impact on the public sector, i.e. on the SPP implementation practices of contracting authorities in various regions and sectors in Switzerland. Economists in study 3 examine the impact of public procurement on the private sector and how procurement reform affects that impact. Compared to Study 2, which focuses on the impact on tender documents - the demand side of the public procurement market - the focus here is on the impact of the reform on business sustainability - the supply side of the public procurement market. In general, the public sector is expected to play a central role in the transformation to greater sustainability.

Research staff: L. Rolshoven, M. Stürmer.

Financial support: SNSF project No. 10000100

10.4 Further Activities

Memberships

Matthias Stürmer

- President of the Digital Impact Network
- President of CH Open
- Advisory Board Member of Opendata.ch
- Member of Smart Capital Region
- Managing Director of Parldigi (Parlamentarische Gruppe Digitale Nachhaltigkeit)

11 Software Engineering Group

11.1 Personnel

Head

Prof. Dr. T. Kehrer	✉ timo.kehrer@unibe.ch ☎ +41 31 684 4618

Office Manager

B. Choffat	✉ bettina.choffat@unibe.ch ☎ +41 31 684 8426

Scientific Staff

PD Dr. C. Tsigkanos	✉ christos.tsigkanos@unibe.ch (until 31.03.2025)
Dr. S. Panichella	✉ sebastiano.panichella@unibe.ch (since 01.09.2024)
Dr. R. Darooei	✉ reza.darooei@unibe.ch (since 01.09.2024)
Dr. M. Ohrndorf	✉ manuel.ohrndorf@unibe.ch
Dr. A. Rohani	✉ atefeh.rohani@unibe.ch (until 31.03.2025, again since 01.06.2025)
J.-A. Bard	✉ jan-andrea.bard@unibe.ch
R. Bögli	✉ roman.boegli@unibe.ch
A. Boll	✉ alexander.boll@unibe.ch
S. Hernández Goicochea	✉ sandro.hernandezgoicochea@unibe.ch (since 01.02.2025, previously Research Assistant)
A. Javadi	✉ ali.javadi@unibe.ch (since 01.04.2025)
R. Macháček	✉ roman.machacek@unibe.ch (since 01.09.2024)
S. Mazraehkhatiri	✉ sajad.mazraehkhatiri@unibe.ch (since 01.01.2025)
J. Spieler	✉ jonas.spieler@unibe.ch (since 01.02.2025)
R. Uysal	✉ ramazan.uysal@unibe.ch (since 01.06.2025)

External Ph.D. Students

C. Birchler	✉ christian.birchler@unibe.ch
A. Schultheiß	✉ alexanderschultheiss@pm.me (until 06.06.2025)
T. Sutter	✉ thomas.sutter@unibe.ch
P. Valenzuela	✉ pablo.valenzuela@unibe.ch
A. Duc Vu	✉ ahn.vu@unibe.ch

Research Assistants

J. Abdullah	✉ judi.abdullah@students.unibe.ch
D. Richard	✉ david.richard@students.unibe.ch (since 01.10.2024)
Y. van Dok	✉ yael.vandok@unibe.ch

Intern

P. Aryan	✉ prakash.aryan@unibe.ch (from 01.04.2025 until 31.07.2025)

11.2 Overview

”Software Engineering Research: One of the Most Critical Endeavors for the Digital Society”
– Open letter published by Informatics Europe in 2025, co-signed by the Swiss Informatics Research Association (SIRA) and other leading associations in computer science

The core message of this letter underscores the indispensable role of software in driving innovation across various sectors – from healthcare to finance, transportation, and education - while highlighting the growing challenges posed by our increasingly complex, software-intensive systems. To meet these challenges, software engineering research must advance fundamental principles that ensure security, correctness, reliability, performance, scalability, safety, and sustainability in these systems.

At the Software Engineering Group (SEG), a research group established in 2022 and headed by Prof. Dr. Timo Kehrer at the Institute of Computer Science at the University of Bern, we embrace this mission from multiple perspectives, contributing to building the foundations for a sustainable digital future through research and teaching.

11.3 Research Projects

Merge++: Turning software merge conflicts into conflict-induced variability

Duration: 2023 - 2027
Financial support: Swiss National Science Foundation (SNSF)

Merging is an essential operation in software version control that reconciles concurrent changes to multiple working copies of a shared development artifact. However, available merge techniques often report considerable amounts of merge conflicts that must be resolved manually, bothering developers in their daily work and hampering continuous integration. Most research on software merging has tackled this problem by increasing the accuracy of conflict detection, yet with limited impact.

In Merge++, we aim to develop foundations that enable a novel paradigm of software version control that temporarily tolerates merge conflicts. The key idea is to take a radically different view on merge conflicts. Instead of enforcing their resolution upon each merge, conflicts shall be considered as a source of software variability that represents different ambitions of developers that shall be made explicit. A merged artifact that comprises such conflict-induced variability may vary its behavior. These different behaviors may be investigated. Eventually, based on the analysis results, developers may take informed decisions to permanently bind or even eliminate conflict-induced variability, which yields a deferred conflict resolution.

Our approach to realize this research vision is to leverage concepts and techniques from research on software product lines, which deals with the systematic handling of software variability. In a nutshell, merge conflicts whose resolution is to be deferred shall be transparently converted to conflict-induced variability, while developers may continue working on individual projections of the integrated version. The binding of accumulated conflict-induced variability shall be informed by adapting concepts from software product-line analysis. The question of how to seamlessly integrate such systematic handling of conflict-induced software variability to efficiently deal with merge conflicts in the context of software version control constitutes an open scientific problem to be solved by Merge++.

We summarize our two overall research goals as follows:

- RG1: Writing working copy changes into a central repository or the main development branch may lead to conflicts in case of concurrent modifications. Instead of enforcing time-consuming conflict resolutions upon each merge attempt, we aim to decouple the writing of changes from conflict resolution by temporarily tolerating conflicts. While the conflicts may be resolved later, continuous development on working copies shall be possible in the meantime.
- RG2: Deferring the resolution of merge conflicts not only accelerates development, but bares the potential for advanced (semi-)automated conflict resolution. By centrally accumulating conflict-induced variability which may be analyzed at the repository site, we aim at synthesizing informed conflict resolution recommendations and guided exploration of resolution spaces, thereby facilitating effective conflict resolution upon request.

VariantSync II: Automating the Synchronization of Software Variants

Project partners: Prof. Dr. Thomas Thüm (TU Braunschweig)
Duration: 2024 - 2027
Financial support: Swiss National Science Foundation (SNSF)

Software product lines have been actively researched for more than two decades with the goal of systematic software reuse through managed software variability. Nevertheless, practitioners frequently rely on ad-hoc reuse based on a principle which is known as clone-and-own, where new variants of a software family are created by copying and adapting an existing variant. However, if a critical number of variants is reached, their maintenance and evolution becomes impractical, if not impossible.

With VariantSync I, we are closer to the vision of bridging the gap between clone-and-own and software product lines by automating the synchronization of software variants. However, during the project we identified two major challenges not being envisioned nor addressed yet by VariantSync I. First, developers that have developed ad-hoc variants in the past cannot immediately specify domain knowledge with 100% confidence and hence we need to incorporate uncertainty in domain knowledge. Second, even though the vision of VariantSync is to automate the synchronization of variants, we cannot expect variants to be 100% synchronized and need

to incorporate that variants drift away from each other over time. Both challenges are crucial for our vision as described below and are planned to be investigated in VariantSync II.

This work is a continuation of VariantSync I (see *INF Annual Report 2023/2024*).

SwarmOps: Human-sensing based MLOps for Collaborative Cyber-Physical Systems

Duration: 2025 - 2029
Financial support: Swiss National Science Foundation (SNSF)

Using AI-enabled Cyber-Physical Systems (AI-CPSs) such as Self-Driving Cars (SDCs) and Unmanned Aerial Vehicles (UAVs) for real-world tasks is now considered the most credible setup in most applications. However, current state-of-the-art solutions for AI-CPSs have a limited perception of the environment, which poses significant challenges in real-world applications, such as mapping disaster areas at nuclear power plants or locating people stranded after natural disasters. This project addresses this main challenge.

InnoGuard: Hybrid and Generative Intelligence for Trustworthy Autonomous Cyber-Physical Systems

Project partners: Simula Research Laboratory (NO), OsloMet - storbyuniversitetet (NO), Università degli Studi del Sannio (IT), Stichting VU (NL), Universidad de Málaga (ES)
Duration: 2024 - 2027
Financial support: Horizon - Marie Skłodowska-Curie Actions-funded Doctoral Networks

The increasing use of Autonomous Cyber-Physical Systems (ACPS), such as self-driving cars, which rely on AI, holds great promise for future AI developments. However, concerns arise due to the lack of robust quality assurance and dependability in current ACPS solutions. Supported by the Marie Skłodowska-Curie Actions programme, the InnoGuard project aims to develop novel methods for ACPS quality assurance by creating tailored

training programmes for early-stage researchers. The project seeks to automate ACPS quality assessment and evolve system behavior. Additionally, it will enhance environmental sustainability, energy efficiency, and the trustworthiness of AI methods for ACPS, while improving overall reliability and legal compliance.

Safe-2-Fly: Advancing UAV Reliability and Societal Trust Through Integrated Testing and Formal Verification

Duration: 2025 - 2026
Financial support: Hasler Foundation

This project aims to improve the safety and reliability of unmanned aerial vehicles (UAVs) by bridging the gap between practical testing methods and theoretical formal verification. While testing helps identify failures through simulations or real-world trials, and formal methods mathematically prove system correctness, each has limitations when used alone. This project proposes a hybrid approach that uses specification mining to automatically extract formal behavioural rules from UAV simulation data.

The project focuses on analyzing time-series execution traces from UAV simulations and field tests to extract meaningful behavioral properties using Temporal Logics. These specifications define both expected and anomalous behaviors, such as safe versus failed landings, and are essential for runtime verification, misbehavior prediction, and test generation. By combining expertise in formal methods and UAV development, Safe-2-Fly aims to develop an efficient and scalable framework for monitoring and verifying UAV behavior in both simulated and real-world environments, ultimately advancing the safety and trustworthiness of autonomous flight systems.

TRUST-VR: AI-Empowered Virtual Reality for Professional Soft Skills Training - A Feasibility Study in Pediatric Care

Duration: 2025 - 2026
Financial support: Hasler Foundation

Virtual Reality (VR) holds transformative potential for medical education. However, its current applications have focused on technical skill training,

whereas its potential for cultivating communication competencies remains underexplored. This limits the broader pedagogical impact of immersive technologies in healthcare training. In this project, we aim to address this gap through the development of an AI-empowered VR training environment which enables realistic and adaptive simulations of emotionally complex clinical interactions.

While our overall research vision is to advance the usage of ICT in professional training by extending immersive technologies towards human-centered, emotionally intelligent skills development, we will limit the scope of TRUST-VR to pediatric care. Given that interactions with frightened children are emotionally charged and require careful trust-building strategies, we consider this as an ideal case to explore both the technical feasibility and the didactic effectiveness of our overall vision.

The project builds on an interdisciplinary research collaboration with the Bern University Hospital (Inselspital), Division of Pediatric Emergency Medicine (Dr. med. et MME Isabelle Steiner and Dr. med. Fabrizio Romano)

AERIALIST: Bridging the Reality Gap in Testing Unmanned Aerial Vehicles

Duration: 2024 - 2025
Financial support: Hasler Foundation

In the last decade, there has been a growing interest in Unmanned Aerial Vehicles (UAVs) or drones, leading to significant technological advancements in avionics. Thanks to artificial intelligence and machine learning, UAVs are becoming more autonomous. However, a major challenge is the limited observability, testability, and predictability of their behavior, which can lead to fatal crashes, even involving humans. Current testing approaches involve expensive and unscalable field testing, while simulation-based testing shows promise in being cheaper and safer. However, simulations may not fully mirror real-world performance, leading to a concern called the Reality Gap.

The AERIALIST project aims to enable simulation-based testing and test automation for UAVs by addressing the Reality Gap. It investigates runtime monitoring for early detection of misbehaviors, reduces simulation-based

testing costs by predicting outcomes without running simulations, and extracts lower-level tests from simulated ones. The project also aims to use logged data from misbehaviors and failing tests to enhance UAV performance in new scenarios.

Swiss-Africa Cybersecurity Community: A Focus on Education, Research and Knowledge Building

Project partners: International project consortium, led by Prof. Dr. Bettina Schneider (FHNW University of Applied Sciences and Arts)
Duration: 2024 - 2026
Financial support: Movetia

Cybersecurity, including social engineering tactics, poses a significant threat to Switzerland. While Swiss Higher Educational Institutions (HEIs) have made commendable efforts to enhance cybersecurity education, the focus often remains on technical measures, neglecting the critical human and cultural aspects. Despite Switzerland's international outlook, cybersecurity awareness regarding the African continent is limited, even as technological adoption there accelerates, bringing new vulnerabilities. Countries like Cameroon and Kenya face challenges such as mobile money fraud, driven by socio-economic factors, linguistic diversity, and distinct regulatory landscapes. Therefore, researchers, educators, and students from Swiss HEIs must understand cybersecurity's cultural dimensions and stay informed on trends across various African regions to strengthen Switzerland's defense. This goal can be achieved by fostering deeper collaboration between Swiss and African cybersecurity stakeholders.

The project aims to build a Swiss-African cybersecurity community via a digital platform to promote cross-cultural exchange, enhance cybersecurity skills, and support collaborative research between HEIs from both regions. Through this community, Swiss educators and students can broaden their understanding of cybersecurity's cultural dimensions, which will inform and improve their curriculums and defenses against global cyber threats. The platform will serve as a space for dialogue, trust-building, and the foundation for deeper partnerships, including potential exchanges. Initially, 5 Swiss and 6 African partners will engage in key activities: creating the platform, piloting research and workshops, and ensuring community sustainability through continuous evaluation and improvement.

CaSSIS-CheckTC: Mission Telecommands Trace Checking

Project partners: Prof. Nicolas Thomas (Space Research and Planetology, Physikalisches Institut, University of Bern)
Duration: 2024 - 2025
Financial support: University of Bern

Telecommands are instructions sent from ground control centers to manage scientific instruments during a mission. Verifying telecommand sequences before uplinking them to the spacecraft involves checking intended sequences against requirements that such sequences need to satisfy. Such telecommand requirements, elicited from mission operators, reflect operational constraints, mission requirements and safety protocols in place. This is a largely manual process that involves domain knowledge and expertise by mission officers.

The goal of CaSSIS-CheckTC consists of providing appropriate tool support to assist mission operators in verifying TC sequences in an automated way, allowing them to catch errors before sending telecommands to the spacecraft. We seek to first (i) elicit and then (ii) express desired TC requirements as temporal logic formulae in order to (iii) enable their automated verification via trace checking. When operators prepare a new telecommand sequence (reflecting, e.g., some new experiment), they can use appropriate CaSSIS-CheckTC tooling to check it for validity against stated requirements. Any deviations, such as missing or out-of-order commands, can then be flagged as anomalies, indicating potential errors that could lead to undesired instrument states.

This project will leverage expertise of the Planetary Imaging Group in the Space Research and Planetology Division of the Physikalisches Institut and the Software Engineering Group at the Institute of Computer Science, towards the verification of CaSSIS flight software. This is a goal that is highly interdisciplinary in nature, since expertise of both groups is involved.

This work is a continuation of CaSSIS-Verif (see *INF Annual Report 2023/2024*).

RUNVERSPACE: Runtime Verification for Space Software Architectures

Duration: 2024 - 2025
Financial support: Swiss National Science Foundation (SNSF)

Recent technological advances have led to powerful and cost-effective on-board computers running software built with open frameworks targeting small-scale missions such as CubeSats and Unmanned Aerial Vehicles. While offering great flexibility, this also opens new potential sources of failures, posing challenges for requirements assurance as development may take place in fast cycles, at higher-levels of abstraction, across multiple teams and culminating into reusable component-based architectures.

RUNVERSPACE addresses the systematic engineering of contemporary small-scale flight- and space- software, arguing that integrating runtime verification facilities is crucial for increased operational requirements assurance. Addressing fundamental challenges of bringing such a goal to runtime has the potential of significant disruption. RUNVERSPACE is based on the premise that cutting-edge software engineering research leveraging applied formal methods is direly needed for the next generation of space applications. RUNVERSPACE will focus on defining suitable programming abstractions and specification notations in tandem with development of architectural support, with an overall goal of demonstrating the potential that runtime verification can bring.

11.4 Theses

11.4.1 Ph.D. Theses

- A. Schultheiß: Towards Managed Clone-and-Own: Automating Matching and Patching

11.4.2 Master's Theses

- S. Anthamatten: Empirical Analysis of Java Software Evolution - Breaking Changes and Migrations in Git Histories

- D. Zeidan: Automated Generation of Code Contracts: Semantic Analysis using Mutation Based Testing
- A. Moisil: Enhancing Code Quality - Customized Source Code Refactoring by Example
- C. Wu: Explaining GitHub Actions Failures with Large Language Models: Challenges, Insights, and Limitations
- S. Hérnandez Goicochea: Addressing Developer Information Needs in GitHub Actions
- M. Mumtaz: Test Decomposition for Cost-effective Testing of UAVs

11.4.3 Bachelor's Theses

- K. Lautenschlager: A Feature Modeling Approach to Mathematical Tasks
- J. Meier: Developing a Hybrid Static-Dynamic Analysis Tool for Third-Party Library Detection
- O. Mingard: Java Conflicting Chunk Derivability and Structured Merge Algorithms
- A. Müller & V. Meier: An Algorithmic Approach for Solving the Grouping Problem
- B. Jossen: The Potential of Principal Component Analysis and Term Frequency-Inverse Document Frequency in Range Queries on N input models
- M. Amon: Segmentation of Source Codes for Minimal Reproducible Examples
- D. Richard: Presortedness-Based Prediction of the Optimal Sorting Algorithm

11.5 Activities

Scientific Boards and Steering Committees

Timo Kehrer

- CHOOSE: Swiss Group for Software Engineering, Special Interest Group (SIG) of the Swiss Informatics Society: Presidency
- VaMoS: International Working Conference on Variability Modelling of Software-Intensive Systems: Steering Committee member

- Re:Volution: International Workshop on Reverse Variability Engineering and Evolution of Software-Intensive System: Steering Committee member
- GReTA: International Seminar Series on Graph Transformation Theory and Applications: Scientific Committee member

Sebastiano Panichella

- International Conference on Software Testing, Verification and Validation (ICST): Steering Committee member
- International Workshop on Search-Based and Fuzz Testing (SBFT): Steering Committee member
- International Workshop on NL-based Software Engineering (NLBSE): Steering Committee member
- Journal of Software: evolution and process: Editorial Board Member
- Transactions on Software Engineering and Methodology: Editorial Board Member
- Empirical Software Engineering (EMSE): Review Board Member
- ACM TOSEM Board of Distinguished Reviewers: Review Board Member

Conference Chairing and Organization

Christos Tsigkanos

- Software-Defined Everything: Engineering Dependable Ubiquitous Systems Symposium 2025: Co-Chair
- International Joint Conferences On Theory and Practice of Software (ETAPS) 2025: Proceedings Chair

Sebastiano Panichella

- International Conference [beginpenalty=10000]on Software Testing, Verification and Validation (ICST) 2025: General Co-Chair, Tool Competition Co-Chair of UAV and Self-driving car competitions
- International Workshop on Search-Based and Fuzz Testing (SBFT) 2025 (Collocated with ICSE 2025): Tool Competition Co-Chair of UAV, SDC, and Python competitions
- International Workshop on Natural Language-Based Software Engineering Workshop (NLBSE) 2025 (Collocated with ICSE 2025): Program Co-Chair

Sajad Mazraehkhatiri

- International Conference on Software Testing, Verification and Validation (ICST) 2025: Web Co-Chair, Tool Competition Co-Chair of UAV competition
- International Workshop on Search-Based and Fuzz Testing (SBFT) 2025 (Collocated with ICSE 2025): Tool Competition Co-Chair of UAV competition

Christian Birchler

- International Conference on Software Testing, Verification and Validation (ICST) 2025: Web Co-Chair, Tool Competition Co-Chair of Self-Driving Car competition
- International Workshop on Search-Based and Fuzz Testing (SBFT) 2025 (Collocated with ICSE 2025): Tool Competition Co-Chair of Self-Driving Car competition

Program Committees

Timo Kehrer

- ICST: International Conference on Software Testing, Verification and Validation: 2025
- MoDELS-WS: International Conference on Model Driven Engineering Languages and Systems, Workshop Track: 2024
- EASE: International Conference on Evaluation and Assessment in Software Engineering: 2025
- SEAMS: International Symposium on Software Engineering for Adaptive and Self-Managing Systems: 2025
- VaMoS: International Working Conference on Variability Modelling of Software-intensive Systems: 2025

Christos Tsigkanos

- ECSA: European Conference on Software Architecture: 2025
- ESEC/FSE: ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Demonstrations Track): 2025

Sebastiano Panichella

- 2030 Software Engineering-track (Collocated with ICSE 2025): 2025
- ICSE: International Conference on Software Engineering: 2025

- FSE: International Conference on the Foundations of Software Engineering: 2025

Roman Bögli

- EASE: International Conference on Evaluation and Assessment in Software Engineering (Main Research Track): 2025

Christian Birchler

- ICST: International Conference on Software Testing, Verification and Validation (Short Papers, Vision and Emerging Results Track): 2025
- ASE: International Conference on Automated Software Engineering (Industry Showcase Track): 2024

Reviewing Activities

Timo Kehrer

- IEEE Transactions on Software Engineering
- ACM Transactions on Software Engineering and Methodology
- Journal of Software: Practice and Experience (Wiley)
- Journal of Systems and Software (Elsevier)
- Journal of Computer Languages (Elsevier)
- IEEE Transactions on Reliability
- Software and Systems Modeling (Springer)
- Science of Computer Programming (Elsevier)
- Automated Software Engineering (Springer)

Sebastiano Panichella

- Empirical Software Engineering (Springer)
- IEEE Transactions on Software Engineering
- ACM Transactions on Software Engineering and Methodology
- Software: Practice and Experience (Wiley)

Reza Darooei

- Scientific Reports (Nature Portfolio)
- Artificial Intelligence (Elsevier)
- International Ophthalmology (Springer)
- International Journal of Machine Learning and Cybernetics (Springer)
- BMC Medical Imaging

- Journal of Medical Signals and Sensors (Medknow Publication)
- Advanced Biomedical Research (Medknow Publication)

Manuel Ohrndorf

- International Conference on Evaluation and Assessment in Software Engineering (EASE): 2025
- Innovations in Systems and Software Engineering (Springer)
- ACM Transactions on Software Engineering and Methodology

Roman Bögli

- International Conference on Software Testing, Verification and Validation (ICST): 2025
- International Working Conference on Variability Modelling of Software-intensive Systems (VaMoS): 2025
- International Conference on Evaluation and Assessment in Software Engineering (EASE): 2025
- ACM Transactions on Software Engineering and Methodology

Alexander Boll

- International Conference on Evaluation and Assessment in Software Engineering (EASE): 2025
- International Working Conference on Variability Modelling of Software-intensive Systems (VaMoS): 2025

Sajad Mazraehkhatiri

- Empirical Software Engineering (Springer)

Christian Birchler

- IEEE Transactions on Software Engineering
- ACM Transactions on Software Engineering and Methodology
- Science of Computer Programming (Elsevier)

Thomas Sutter

- Computers & Security (Elsevier)
- AI Perspectives & Advances (Springer)

Pablo Valenzuela

- ACM Transactions on Software Engineering and Methodology
- Journal of Software: Evolution and Process (Wiley)
- International Conference on Mining Software Repositories (MSR): 2025
- International Conference on Evaluation and Assessment in Software

Engineering (EASE): 2025

Presentations

Timo Kehrer

- ”Mining Domain-Specific Edit Operations from Model Repositories with Applications to Semantic Lifting of Model Differences and Change Profiling”, Software Engineering (SE) 2025, Karlsruhe DE

Christos Tsigkanos

- ”Leitmotifs on Space Software Engineering: From Specification to Architecture to Verification”, Invited Talk at Gran Sasso Science Institute (GSSI), 2025, L’Aquila IT

Sebastiano Panichella

- ”ICST Tool Competition 2025 - Self-Driving Car Testing Track”, ICST’25, Naples IT
- ”ICST Tool Competition 2025”, ICST’25, Naples IT
- ”Development and Testing of AI-enabled Systems: Reflection on Automation and Practices”, Research Talk at the University of Lugano, 2025, Lugano CH

Manuel Ohrndorf

- ”Automated Generation of Code Contracts - Generative AI to the Rescue?”, GPCE’24, Pasadena, California USA

Roman Bögli

- ”A Systematic Literature Review on a Decade of Industrial TLA+ Practice”, iFM’24, Manchester UK
- ”Community-Driven Variability”, FOSD’25, Köthen DE
- ”Temporal Logics Meet Real-World Software Requirements: A Reality Check”, FormaliSE’25, Ottawa CA
- ”Beyond Software Families: Community-Driven Variability”, FSE’25, Trondheim NO
- ”Cryptographic Agility in the Realm of Post-Quantum-Secure Algorithms”, Swiss-African Cybersecurity Community Workshop, 2025, Bern CH

Alexander Boll

- ”Towards Semi-Automated Merge Conflict Resolution: Is It Easier Than We Expected?”, Software Engineering (SE) 2025, Karlsruhe DE

Sajad Mazraehkhatiri

- ”ICST Tool Competition 2025 – UAV Testing Track”, ICST’25, Naples IT
- ”CertiFail at the ICST 2025 Tool Competition–Self-Driving Car Testing Track”, ICST’25, Naples IT
- ”NN-SDCTest at the ICST 2025 Tool Competition-Self-Driving Car Testing Track”, ICST’25, Naples IT

Christian Birchler

- “ICST Tool Competition 2025 – Self-Driving Car Testing Track”, ICST’25, Naples IT

Thomas Sutter

- ”Security and Privacy Concerns on Android Phones”, Swiss-African Cybersecurity Community Workshop, 2025, Bern CH
- ”Uninstallable by Design”, OWASP Switzerland Community Meetup, 2025, Winterthur CH
- ”Uninstallable by Design: The Role of Pre-installed Apps in Android’s Security Landscape”, Computer Security Day, IEEE Computer Society, Madras, 2024, Taramani IN

Pablo Valenzuela

- ”The Hidden Costs of Automation: An Empirical Study on GitHub Actions Workflow Maintenance”, SCAM’24, Flagstaff, Arizona USA
- ”Run Failure Diagnostic”, AIware Leadership Bootcamp 2024, Toronto CA
- ”Explaining GitHub Actions Failures with Large Language Models: Challenges, Insights, and Limitations”, ICPC’25, Ottawa CA

Ahn Duc Vu

- ”Towards Generating Contracts for Scientific Data Analysis Workflows”, WORKS’24, Atlanta, Georgia USA

Outreach and Other Activities

- Timo Kehrer gave a talk for interested students at the Gymnasium Kirchenfeld, Bern, entitled ”Was tun, wenn mein neues Auto mehr Konfigurationen als Atome im Universum hat?”.
- Timo Kehrer served as expert for the Matura Examination Commission Canton Bern.

- Timo Kehrer served as expert for the Swiss Agency of Accreditation and Quality Assurance.
- Sebastiano Panichella served as Management Committee Member for COST Action CA22137 ((European Cooperation in Science and Technology - ROAR-NET) (Working Group 4 (Optimisation under Uncertainty) at the Second General Meeting in Sarajevo, Bosnia & Herzegovina.
- Sebastiano Panichella was invited to and participated in the Huawei research meeting on Software Engineering challenges at the University of Lugano CH.
- Sebastiano Panichella co-organized 2 international PhD summer schools in the context of the InnoGuard EU Research Project (Marie Skłodowska-Curie Actions-funded Doctoral Networks). The host institutions were University of Sannio, Italy, Simula Research Laboratory, Norway, and Oslo University, Norway.
- Roman Bögli is a member of the Prüfungskommission des Schweizerischen Informatik Zertifikats (SIZ).
- Yael van Dok served as Student Volunteer at the 18th IEEE International Conference on Software Testing, Verification and Validation (ICST) 2025.

11.6 Publications

Conference Papers

- Nitish Patkar, Aimen Fahmi, Timo Kehrer, Norbert Seyff: What is a Feature, Really? Toward a Unified Understanding Across SE Disciplines. IEEE/ACM 47th International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER) 2025: 41-45
- Roman Bögli, Atefeh Rohani, Thomas Studer, Christos Tsigkanos, Timo Kehrer: Temporal Logics Meet Real-World Software Requirements: A Reality Check. International Conference on Formal Methods in Software Engineering (FormaliSE) 2025: 74-85
- Pablo Valenzuela-Toledo, Chuyue Wu, Sandro Hernández, Alexander Boll, Roman Machácek, Sebastiano Panichella, Timo Kehrer: Explaining GitHub Actions Failures with Large Language Models: Challenges, Insights, and Limitations. 2025 IEEE/ACM 33rd International Conference on Program Comprehension (ICPC): 286-297

- Christian Birchler, Sajad Khatiri, Pouria Derakhshanfar, Sebastiano Panichella, Annibale Panichella: Single and Multi-objective Test Cases Prioritization for Self-driving Cars in Virtual Environments. Software Engineering (SE) 2025: 13
- Alexander Boll, Yael Van Dok, Manuel Ohrndorf, Alexander Schultheiß, Timo Kehrer: Towards Semi-Automated Merge Conflict Resolution: Is It Easier Than We Expected? Software Engineering (SE) 2025: 20
- Christof Tinnes, Timo Kehrer, Mitchell Joblin, Uwe Hohenstein, Andreas Biesdorf, Sven Apel: Mining Domain-Specific Edit Operations from Model Repositories with Applications to Semantic Lifting of Model Differences and Change Profiling. Software Engineering (SE) 2025: 21
- Christian Birchler, Tanzil Kombarabettu Mohammed, Pooja Rani, Teodora Nechita, Timo Kehrer, Sebastiano Panichella: How Does Simulation-Based Testing for Self-Driving Cars Match Human Perception? Software Engineering (SE) 2025: 41
- Roman Bögli, Alexander Boll, Alexander Schultheiß, Timo Kehrer: Beyond Software Families: Community-Driven Variability. ACM SIGSOFT Conference on the Foundations of Software Engineering (FSE) 2025: 571-575
- Elias Kuiter, Thomas Thüm, Timo Kehrer: Teach Variability! A Modern University Course on Software Product Lines. 19th International Working Conference on Variability Modelling of Software-Intensive Systems (VaMoS) 2025: 1-10
- Christian Birchler, Stefan Klikovits, Mattia Fazzini, Sebastiano Panichella: SBFT Tool Competition 2025 - CPS-SDC Regression Testing Track. International Workshop on Search-Based and Fuzz Testing (SBFT) 2025
- Sajad Khatiri, Tahereh Zohdinasab, Prasun Saurabh, Dmytro Humeniuk, Sebastiano Panichella: SBFT Tool Competition 2025 - UAV Testing Track. International Workshop on Search-Based and Fuzz Testing (SBFT) 2025
- Christian Birchler, Stefan Klikovits, Mattia Fazzini, Sebastiano Panichella: ICST Tool Competition 2025 - Self-Driving Car Testing Track. 18th IEEE International Conference on Software Testing, Verification and Validation (ICST) 2025: 801-804
- Prakash Aryan, Sajad Khatiri: NN-SDCTest at the ICST 2025 Tool

Competition - Self-Driving Car Testing Track. 18th IEEE International Conference on Software Testing, Verification and Validation (ICST) 2025: 813-814

- Sajad Khatiri, Tahereh Zohdinasab, Prasun Saurabh, Dmytro Humeniuk, Sebastiano Panichella: ICST Tool Competition 2025 - UAV Testing Track. 18th IEEE International Conference on Software Testing, Verification and Validation (ICST) 2025: 815-818
- Ali Javadi, Christian Birchler: TGen-UQ at the ICST 2025 Tool Competition - UAV Testing Track. 18th IEEE International Conference on Software Testing, Verification and Validation (ICST) 2025: 821-822
- Sandra Greiner, Noah Bühlmann, Manuel Ohrndorf, Christos Tsigkanos, Oscar Nierstrasz, Timo Kehrer: Automated Generation of Code Contracts: Generative AI to the Rescue? ACM SIGPLAN International Conference on Generative Programming: Concepts & Experiences (GPCE) 2024: 1-14
- Roman Bögli, Leandro Lerena, Christos Tsigkanos, Timo Kehrer: A Systematic Literature Review on a Decade of Industrial TLA+ Practice. 9th International Conference on Integrated Formal Methods (iFM) 2024: 24-34
- Alexander Boll, Timo Kehrer, Michael Goedicke: SMOKE: Simulink Model Obfuscator Keeping Structure. 27th International Conference on Model Driven Engineering Languages and Systems (MoDELS) 2024: 41-45
- Anh Duc Vu, Timo Kehrer: Towards Generating Contracts for Scientific Data Analysis Workflows. SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis, 2024: 2048-2055
- Pablo Valenzuela-Toledo, Alexandre Bergel, Timo Kehrer, Oscar Nierstrasz: The Hidden Costs of Automation: An Empirical Study on GitHub Actions Workflow Maintenance. IEEE International Conference on Source Code Analysis and Manipulation (SCAM) 2024: 213-223
- Sandra Greiner, Alexander Schultheiß, Paul Maximilian Bittner, Thomas Thüm, Timo Kehrer: Give an Inch and Take a Mile? Effects of Adding Reliable Knowledge to Heuristic Feature Tracing. 28th ACM International Systems and Software Product Line Conference (SPLC) 2024: 84-95
- Paul Maximilian Bittner, Alexander Schultheiß, Benjamin Moosherr,

Jeffrey M. Young, Leopoldo Teixeira, Eric Walkingshaw, Parisa Ataei, Thomas Thüm: On the Expressive Power of Languages for Static Variability. Object-oriented Programming, Systems, Languages, and Applications (OOPSLA) 2024: 307

Journal Papers

- Christian Birchler, Sajad Khatiri, Pooja Rani, Timo Kehrer, Sebastiano Panichella: A Roadmap for Simulation-Based Testing of Autonomous Cyber-Physical Systems: Challenges and Future Direction. ACM Transactions on Software Engineering and Methodology 34(5): 152:1-152:9 (2025)
- Alessio Gambi, Sebastiano Panichella: Preface for the special issue on SBFT'23: Search-Based and Fuzz Testing - Tools. Science of Computer Programming (Elsevier) 239: 103180 (2025)
- Timo Kehrer, Robert Haines, Guido Juckeland, Shurui Zhou and David. E. Bernholdt, Do Research Software Engineers and Software Engineering Researchers Speak the Same Language?. IEEE Computing in Science & Engineering 27, 2: 18-26 (2025)
- Sajad Khatiri, Andrea Di Sorbo, Fiorella Zampetti, Corrado Aaron Visaggio, Massimiliano Di Penta, Sebastiano Panichella: Identifying safety-critical concerns in unmanned aerial vehicle software platforms with SALIENT. SoftwareX 27: 101748 (2024)

Others

- Pooja Rani, Jan-Andrea Bard, June Sallou, Alexander Boll, Timo Kehrer, Alberto Bacchelli: Can We Make Code Green? Understanding Trade-Offs in LLMs vs. Human Code Optimizations. CoRR abs/2503.20126 (2025)
- Thomas Sutter, Ariane Trammell, Timo Kehrer: Uninstallable by Design: The Role of Pre-installed Apps in Android's Security Landscape. ERCIM News 2024(139) (2024)

11.7 Awards

- Timo Kehrer: Best Paper Award at VaMoS 2025, the International Working Conference on Variability Modelling of Software-Intensive

Systems, for the work ”Teach Variability! A Modern University Course on Software Product Lines” by Elias Kuiter, Thomas Thüm and Timo Kehrer.

- Timo Kehrer, Alexander Schultheiß: Best Research Paper Award at SPLC 2024, the International Systems and Software Product Lines Conference, for the work ”Give an Inch and Take a Mile? Effects of Adding Reliable Knowledge to Heuristic Feature Tracing” by Sandra Greiner, Alexander Schultheiß, Paul Maximilian Bittner, Thomas Thüm, Timo Kehrer.
- Sebastiano Panichella: Best reviewer award at FSE 2025, the ACM International Conference on the Foundations of Software Engineering.
- Sebastiano Panichella: Most Influential 5-Years Journal Paper on Software Testing at ICST 2025, the International Conference on Software Testing, Verification and Validation, for the work ”How developers engage with static analysis tools in different contexts” (published in Empirical Software Engineering (Springer)) by the Best Journal Committee.
- Sebastiano Panichella: Elevation to the grade of IEEE Senior Member: Only 10% of IEEE’s more than 450,000 members hold this grade, 2024.
- Jonas Spieler: 1st Place at the SAPIENCE Drone Competition 2 in Huntsville, Alabama, with the team Klagenfurt, AU (Prof. Stephan Weiss), 2025.
- Alexander Schultheiß: Distinguished Artifact Award at OOPSLA 2024, Object-oriented Programming, Systems, Languages, and Applications (collocated with SPLASH 2024), for the work ”On the Expressive Power of Languages for Static Variability - Artifact”, by Paul Maximilian Bittner, Alexander Schultheiß, Benjamin Moosherr, Jeffrey M. Young, Leopoldo Teixeira, Eric Walkingshaw, Parisa Ataei, Thomas Thüm.

12 Administration

University:

D. Bommes:	Member of the Research Commission (Forschungskommission)
T. Braun:	Member of Focus Group eXtended Reality
T. Studer:	Member of *Kantonale Maturitätskommission*
	Member of *Kommission Gymnasium – Hochschule*

Faculty:

D. Bommes:	Joint Master in Computer Science of the Universities of Bern, Fribourg and Neuchâtel: President of the Branch Committee
	Member of the Strategy Board; Faculty delegate
T. Braun:	Member of the Bern Data Science Initiative (BeDSI)
A. Di Maio:	Member of the Strategy Board
P. Favaro:	Member of the Bern Data Science Initiative (BeDSI)
	Committee Member of the Data Science Lab (CAS Program Management); Fachbereich INF-MATH-STAT representative
T. Kehrer:	Member of the Strategy Board; Deputy Faculty delegate
	Member of the Quality Assurance and Development expert committee
	Equal Opportunities Commission; Department Delegate
T. Studer:	Member of the Strategy Board
	Representative of Oberer Mittelbau in faculty meetings

Institute:

D. Bommes:	Director of Studies
C. Cachin:	Deputy Managing Director of INF
	Representative to CUSO Doctoral School in Computer Science
P. Favaro:	Managing Director of INF
T. Kehrer:	Deputy Director of Studies
	Member of Hauskommission Engehalde

K. Riesen:	Member of Library Committee on behalf of INF Representative to CUSO Doctoral School in Computer Science
T. Studer:	Member of Hauskommission Exakte Wissenschaften