Università della Svizzera italiana

Faculty of Informatics
Master's Research Scholarships (MaRS) Projects
2016-2025

Conceptual Understanding and Problem Solving in Learning to Program
Contact: Prof. Matthias Hauswirth
Learning to program is hard because programming requires (1) a solid understanding of programming language concepts and (2) mastery of the problem solving process. In the context of our research on understanding and improving how people learn to program, the Luce research group at USI (https://luce.si.usi.ch/) is developing educational approaches and platforms to study and improve both of these aspects.
To gain insight into the understanding of programming language concepts, we are developing **Expression Tutor** (https://expressiontutor.org/). This platform helps to teach and assess the understanding of expressions, one of the most prevalent concepts of programming languages.
To study and improve the problem solving process, we are developing **PyTamaro Web** (https://pytamaro.si.usi.ch/), a platform that helps to teach problem decomposition to novice Python programmers using a pure functional graphics library.
If you are interested in programming languages and in research on teaching and assessing the understanding of programming language concepts and the problem solving process, contact us to learn more.
We are looking for Master students to contribute in one or more ways to our research projects:
* Develop novel extensions of our platforms. This requires proficiency in JavaScript and experience developing with React.
* Collaborate on conducting empirical educational studies using our platforms to gain insights into learning. This requires an understanding of quantitative and/or qualitative methods of empirical research with human subjects.
* Investigate the impact of deep learning and large language models on conceptual understanding and problem solving in programming education. This requires a background in machine learning.

Fast constraint Delaunay Triangulation updates and related algorithms
Contact: Prof. Evanthia Papadopoulou
The constrained Delaunay triangulation (CDT) is a variant of the well-known Delaunay triangulation in which specified edges, sometimes called segments, are constrained to appear. These constraints have many uses, such as representing boundaries of non-convex objects, supporting better interpolation of discontinuous functions, and aiding the enforcement of boundary conditions in finite element meshes [1]. The CDT of a point-set is as close to being a Delaunay triangulation as possible subject to these constraints. The project will develop fast linear- (or near-linear) time algorithms to construct and update constraint Delaunay triangulations in cases of interest such as a monotone polygon, a simple polygon, and the general incremental CDT construction. The latter problem represents the

first goal in this project; the technique will then be expanded towards the other problems. The CDT of a simple polygon can pave the way to a fast and simple medial axis construction. The medial axis transform (MAT) is a widely-known shape descriptor used in diverse scientific areas. Constrained Delaunay triangulation are used, among others, in Delaunay refinement methods for mesh generation. The project has room for choices according to the interests of the prospective student.
The proposed algorithms are based on our recent research results (still in progress) constructing Voronoi-like graphs (and their dual) in linear expected time. They are based on the following two papers.
[1] Fast segment insertion and incremental construction of constrained Delaunay triangulations, Jonathan Richard Shewchuk and Brielin C. Brown, Computational Geometry: Theory and Applications, 2015
[2] Deletion in Abstract Voronoi diagrams in expected linear time, K. Junginger and E. Papadopoulou, SoCG 2018
https://drops.dagstuhl.de/opus/volltexte/2018/8763/pdf/LIPIcs-SoCG-2018-50.pdf

Setting a realistic simulation environment for testing self-driving cars
Contact: Prof. Paolo Tonella
Co-Supervisor: Nargiz Humbatova
In the last decade, Deep Learning (DL) solutions are adopted in a constantly growing number of domains. DL applications influence important aspects of life by tackling tasks such as fraud detection or face recognition, while some are employed in safety critical domains such as autonomous driving. This makes reliability and robustness of such systems of a high importance. While testing such systems in the real-world environment ensures the reliance of the results, it becomes increasingly expensive for complex systems. For example, providing an autonomous vehicle with all kinds of possible road situations to evaluate its safety is non-trivial and sometimes even an impossible task. Simulation testing is an efficient way to test systems like autonomous vehicles before conducting resource-demanding real-word trials.

CARLA
CARLA is an open-source autonomous driving simulator [1]. It was developed specifically to serve the needs of researches in develop-ment, training, and testing of various autonomous driving systems. CARLA comes with the flexibility in customisation of the tool and wide range of assets (urban layouts, buildings, vehicles, pedestrians, street signs, etc.) that aid generation of various testing scenarios. It is based OpenDRIVE standards to define the urban setting and Un-real Engine to run the simulation. Moreover, CARLA allows flexible setup of sensors typical to autonomous vehicles (such as cameras and LIDARs) and provides users with features such as usage of GPS coordinates, accelerations to enable training of different driving strategies. On Fig. 1 is provided an example of the same simulation scenario in CARLA simulator under 4 different weather conditions.

Project Proposal
The goal of the project is to set up and prepare a simulation envi-ronment in CARLA for testing autonomous cars. The environment should include pedestrians, vechicles, traffic signs, and other rea-sonable obstacles. It also involves adopting a Deep Learning system emulating an autonomous car that is able to successfully perform in the described environment. The last step would be to integrate the calculation of various driving quality metrics in the simulator. The examples of such metrics could be speed, lane position, acceleration, and etc. [3]. This project will prepare grounds for further studies in the domain and academic contributions.
In the frame of this project, the student will learn about state-of-the-art simulation practices in the domain of autonomous vehicle testing, will practice with popular DL frameworks and model archi-tectures suitable for the task of autonomous driving. Moreover, the student will

be able to have practice in data collection and train-ing of a DL model and later employing the trained model in the simulator.

Additional Information
The project will be carried out within the TAU research group at the Software Institute (https://www.si.usi.ch) and contribute to the PRECRIME ERC research project (https://www.pre-crime.eu). Students are supervised by researchers of the TAU group who follow them constantly and provide them with timely feedback, advice and directions. The code developed for the projects is typically released as an open source project and the results are often included in scientific publications. Both code and publication would contribute to a stronger CV of the participating student.

References
[1] 2023. CARLA. https://carla.org/
[2] Alexey Dosovitskiy, Germán Ros, Felipe Codevilla, Antonio M. López, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. ArXiv abs/1711.03938 (2017).
[3] Gunel Jahangirova, Andrea Stocco, and Paolo Tonella. 2021. Quality Metrics and Oracles for Autonomous Vehicles Testing. In 2021 14th IEEE Conference on Software Testing, Verification and Validation (ICST). 194–204. https://doi.org/10.1109/ICST49551.2021.00030

Figure 1: A street from CARLA simulator from a third-person view in 4 different weather conditions (clear day, day-time rain, daytime shortly after rain, and clear sunset)[2]

Towards an Intelligent Post-training Mutation Tool for Deep Learning Systems
Contact: Prof. Paolo Tonella
Co-Supervisor: Dr. Nargiz Humbatova
Deep Learning (DL) has become an integral part of many ground-breaking projects and products we use everyday. As quality and safety remains the main concern for the developers and users of modern products based on Artificial Intelligence, different techniques aimed at assessing their quality are of increasing interest for research community.

Mutation Testing
Mutation testing is a technique that deliberately seeds faults in form of small syntactic changes into the program under test to create a set of faulty programs called mutants. The general principle underlying this approach is the assumption that faults used by mutation testing represent the mistakes that programmers usually make. Mutation testing aims to assess the quality of a given test suite in terms of its capability to detect faults. For this, the test suite is executed on each of the generated mutants. If the result for a given mutant is different from the result of running the original program then the mutation is considered killed. The ratio of killed mutants to the overall number of generated mutants is called mutation score. The higher the mutation score, the better is the quality of the test suite.
The example in Figure 1 shows a method subtract that subtracts two integer values and returns the result. It has two mutations: in Mutant 1 the subtraction is replaced with multiplication and in Mutant 2 it is replaced with addition. If our test suite has only test0, none of the two mutations would be killed (as they all return the expected value 0) and the mutation score is 0%. If we add test case test1() to our test suite, then Mutant 1 gets killed

and the mutation score becomes 50%. Once we add test case test2(), both mutations get killed and the mutation score achieves its maximum value of 100%.

Mutation Testing for DL Systems

In traditional software systems the decision logic is often imple-mented by software developers in the form of source code. In con-trast, the behaviour of a DL system is mostly determined by the training data set and the training program, i.e. these are the two major sources of defects for DL systems. Thus, there should be a specific approach to mutation testing of DL systems. There cur-rently exist two tools that are designed specifically for performing mutation testing for DL systems. However, one of the tools is a pre-training one, which means it injects the faults into system prior to the training and thus is computationally expensive, while the second one, a post-training mutation tool, injects faults that are random and not very likely to happen in real world. Such faults usually introduce slight noise or modifications to a randomly se-lected subset of weights or change a structure of an already trained DL model by adding/deleting its layers or replacing the activation function.

Currently, mutation testing is being applied to various tasks for DL systems such as program repair, generation of optimal oracles for self-driving cars, detection of adversarial inputs, prioritisation of test inputs for the labelling, etc. Availability of a mutation tool that can inject changes that resemble the effect inflicted by real faults would be extremely useful also for these approaches as well as the advance of DL testing in general.

Project Proposal

The goal of the project is to develop a new post-training mutation tool that would solve the limitations set by the previous approaches, i.e. to introduce smarter and fast DL-specific mutation operators that produce stable and reliable results and would facilitate the increased interest for mutation testing in DL community.

In the frame of this project, the student will learn about state-of-the-art techniques in the domain of mutation testing for DL systems, their limitations and advantages. The student will practice with most popular DL frameworks and widely-used models and datasets.

Additional Information

The project will be carried out within the TAU research group at the Software Institute (https://www.si.usi.ch) and contribute to the PRECRIME ERC research project (https://www.pre-crime.eu). Students are supervised by researchers of the TAU group who follow them constantly and provide them with timely feedback, advice and directions. The code developed for the projects is typically released as an open source project and the results are often included in scientific publications. Both code and publication would contribute to a stronger CV of the participating student.

```
// Original Program
public int subtract(int a, int b) {
    return a - b;
}

// Mutant 1
public int subtract(int a, int b) {
    return a * b;
}

// Mutant 2
public int subtract(int a, int b) {
    return a + b;
}
```

```
public void test0() {
    assertEquals(0, subtract(0, 0));
}

public void test1() {
    assertEquals(-4, subtract(-2, 2));
}

public void test2() {
    assertEquals(1, subtract(2,1));
}
```

Figure 1: Mutation Testing Example

Generating Valid Test Inputs for Deep Learning Systems
Contact: Prof. Paolo Tonella
Co-Supervisor: Matteo Biagiola

Deep Learning (DL) models are being used to address a variety of tasks, from image classification to conversational chatbots. As DL models are being deployed in production, it is of paramount importance to test their capabilities.
Several Test Input Generators (TIGs) have been proposed in the literature to test the generalization capabilities of DL models with artificially crafted inputs in order to induce a misclassification.
However, such artificially crafted inputs might not be valid, i.e., belonging to the input domain the model has been trained on. On the other hand, automatically checking the validity of the generated inputs is difficult, as validity is hard concept to formalize and, hence, automate.

Project Proposal

The goal of the project is to develop a test generation pipeline for DL models that exploit human feedback to automatically validate generated inputs. The project will focus on image classification as a DL task. An input is valid for a DL model if it is recognizable by human domain experts in the input domain, i.e., an input to which a human can confidently assign a label taken from the input domain.
For instance, let us consider the handwritten digit classification task using the MNIST dataset.

Tasks

For the realization of this project, the candidate is expected to perform a series of tasks, among which: (1) familiarize with the task of handwritten digit classification using supervised learning;(2) familiarize with existing test input generators for DL; (3) design an experiment to train a reward model using human-labeled data and possibly integrate the trained reward model into an existing test input generator.

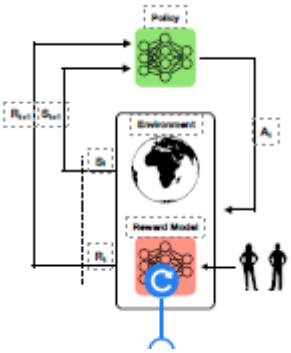

Figure 1: MNIST digits generated by different test generators.
Figure 2: Overview of the approach.

Figure 1 shows the image of a “5” digit generated by different TIGs proposed in the literature. Such inputs are misclassified by the DL model under test, i.e., the label predicted by the model is different than 5. However, the validity, and hence the significance, of some of the inputs is debatable (e.g., it is difficult to assign a label at the digit in Figure 1.c).
Since input validity is ultimately a human-related concept, the idea is to model human validity for a certain domain and use such model to guide the generation of valid inputs. Figure 2 presents a possible application of the model by integrating it into a typical reinforcement learning pipeline (i.e., reward model).

Prerequisites
We are looking for a student who is passionate about DL and mo-tivated to contribute to the project. Knowledge about DL is not required, although it is appreciated. We highly value the willing-ness to learn these technologies and adapt them depending on the project's needs. During the project, the candidate can rely upon an existing software infrastructure that integrates existing test input generators for different classification tasks (including digit classifi-cation). We will work closely with you and provide assistance when needed.

Why You Should Choose This Project
This project is unique in its multi-disciplinary nature: you will learn and use concepts in machine learning (ML), especially DL, software engineering and the ability to set up an experimental setting to rigorously compare different methods.
Upon successful completion of the project, you will have con-tributed to an open-source project that will allow developers to test their DL models using valid and hence reliable inputs. On your CV side, you will be able to show practical experience with training and above all evaluating DL models.
Further Information
The proposed work is part of the Precrime (Self-assessment Oracles for Anticipatory Testing) ERC project. The interested student can find more information about Precrime on the project's website: https://www.pre-crime.eu/. Interested in this project? Any still-unanswered questions? Drop us an email now!

Test input prioritisation for DL systems
Prof. Paolo Tonella
Co-Supervisor: Tahereh Zohdinasab
Deep Learning (DL) systems have achieved unprecedented success in solving complex tasks, including safety critical ones such as self-driving. It is crucial to guarantee the reliability of DL systems by testing how they behave with different inputs. Testing DL systems is challenging due to their huge input space, e.g., all possible road configurations for a self-driving car. Moreover, testing complex DL systems may need the execution of expensive simulations. Consequently, prioritising test inputs by choosing only 'high quality' ones, i.e., tests more likely to trigger a failure, is necessary for test cost reduction.
Multiple test prioritisation techniques have been proposed by Software Engineering researchers. However, none of them considered the unique-ness of prioritised test inputs, which guarantees the evaluation of the DL system under different circumstances.
Feature maps (i.e., maps where the inputs are positioned based on their characteristics) can report useful details about a test set, such as the features corresponding to tests that triggered misbehaviours or the probability of observing a misbehavior for each feature combination. In this work, we propose to use feature maps to prioritise diverse, critical test inputs.
The goal of this project is to propose a new approach for test pri-oritisation for DL systems based on feature maps. The first task of the candidate is to design a test prioritisation strategy that leverages feature maps. The second task is to replicate state of the art approaches for test prioritisation for DL systems and compare them to the proposed approach.

Perceptually-inspired techniques for novel displays and 3D printing
Contact: Prof. Piotr Didyk
In recent years, there has been a tremendous increase in the quality and the number of new output devices. New virtual and augmented reality headsets are being developed to provide new, interactive, and immersive ways to explore the real and the virtual worlds. 3D printers empower millions of users to create, customize, and manufacture digital models. Unfortunately, the abilities of these emerging technologies outperform the capabilities of current methods and tools for creating content.

Our group offers a wide range of projects in the field of computer graphics with a common goal of developing novel computational methods for driving novel display devices (e.g., AR/VR displays) or different digital fabrication devices (e.g., 3D printers). They involve (1) working with novel hardware, such as VR and AR displays or 3D printers, (2) investigating their essential aspects which influence the perceived quality, and (3) devising new content optimization and creation techniques which: maximize quality and human performance while overcoming computational and hardware challenges. Besides using traditional computational approaches, the projects also aim at employing new machine learning techniques, such as deep neural networks, to solve the above challenges. The examples of the projects include but are not limited to foveated rendering for VR and AR displays, content and hardware optimization for varifocal displays, appearance fabrication, and 3D printing surfaces with prescribed haptic properties. For more information about research topics see our work on https://www.pdf.inf.usi.ch/ or contact us via email.

Programming Language for Reactive Distributed Monitoring
Contact: Prof. Patrick Eugster
Co-Supervisor: Pavel Chuprikov
Existing network monitoring solutions primarily focus on querying information from the network rather than on the responses these queries would ensue. In contrast we focus on a design that combines monitoring and management in a single system. Our system aims at exploiting programmability of network devices to perform as many of the management actions as possible exactly at the point of data collection, i.e., at the switches, through small agents called seeds. To relieve the programmer from the burden of low-level resource management, we have designed a domain-specific programming language called Almanac and optimization algorithm for resource allocation and seed placement. The selected student will work on automated translation from Almanac to lower level code accessing the existing API of our system, gaining experience in programming language design and implementation. The translation would also need to include simple static analysis to produce the set of optimization constraints for the seed placement algorithm. The exact implementation framework is up to a discussion, but [1], [2], and [3] provide some examples.
[1] https://arzg.github.io/lang/
[2] https://www.stephendiehl.com/llvm/
[3] https://docs.racket-lang.org/guide/languages.html

Rust for Kernel-level Distributed Services
Contact: Prof. Patrick Eugster
Co-Supervisors: Davide Rovelli and Pavel Chuprikov
Since its conception more than 50 years ago, C has been a widely adopted programming language in a variety of applications. Today, C is still the de-facto standard for low-level programs, lying at the core of operating systems, device drivers, and network protocols. Despite its remarkable performance, C is known to have several security vulnerabilities often resulting in undetected memory bugs and data races costing millions to companies. An alternative to C is provided by Rust: a low-level systems programming language focusing on performance, reliability and robustness.
In December 2022, the Linux Kernel added support for Rust, increasing its relevance in the systems community. This makes Rust an ideal candidate to write novel secure and fast kernel-level services for future distributed systems. The goal of this project is to use Rust to integrate a simple distributed application into the Linux operating system. The student will evaluate to what extent Rust is currently usable for Linux kernel services and build a proof-of-concept service as a Linux kernel module.

Analysis of Proofs of Unsatisfiability for SMT Solvers
Contact: Prof. Patrick Eugster
Co-Supervisor: Rodrigo Otoni

Formal verification aims at providing strong guarantees about the behavior of programs and systems. It relies on logic to precisely describe the program or system in question and check any desired properties, e.g., deadlock-freedom on a concurrent system. Both academia and industry recognize the need for formal verification, with one example being the extensive use of formal verification at Amazon Web Services. Many verification approaches rely on automated reasoning engines for first-order logic. They ultimately reduce the verification problem to checking the satisfiability of first-order logic formulas. Such formulas are commonly expressed in a form suitable to satisfiability modulo theories (SMT) solvers, which makes these solvers a critical part of the verification infrastructure. Despite being used to ensure correctness, SMT solvers are known to have bugs. To tackle the problem of bugs leading a SMT solver to an incorrect result, proofs of unsatisfiability, which explain the achieved result, can be produced by the solver. The goal of this project is to analyze existing proof formats used by different SMT solvers. The student will learn the basics of SMT solving and will endeavor to understand the proof formats used by four SMT solvers – Z3, CVC5, veriT, and OpenSMT. The formats will be compared both quantitatively, via an evaluation using publicly available benchmarks, and qualitatively, via a description of the elements underpinning each proof format.

Embedding Proofs of Unsatisfiability for SMT Solvers into CHC Solving
Contact: Prof. Patrick Eugster
Co-Supervisor: Rodrigo Otoni
Formal verification aims at providing strong guarantees about the behavior of programs and systems. It relies on logic to precisely describe the program or system in question and check any desired properties, e.g., deadlock-freedom on a concurrent system. Both academia and industry recognise the need for formal verification, with one example being the extensive use of formal verification at Amazon Web Services. Constrained Horn Clauses (CHC) are a fragment of first-order logic that naturally captures many verification problems. Reasoning about first-order logic formulas boils down to satisfiability queries, commonly solved by satisfiability modulo theories (SMT) solvers. SMT solving is thus fundamental to CHC-based verification techniques. Despite being used to ensure correctness, SMT solvers are known to have bugs, which undermines all guarantees given as to the behavior of a program being analyzed. To tackle this problem, proofs of unsatisfiability for SMT solvers can be used. Recently, many SMT solvers started to have the option to produce such proofs. Verification tools, however, don't commonly benefit from these proofs to strengthen the guarantees about the final result provided to the user. The goal of this project is to embed SMT proofs of unsatisfiability into CHC-based verification. The student will learn about the basics of CHC solving and of proofs of unsatisfiability for SMT, and will then integrate the proofs into the pipeline of the Golem CHC solver.

Formal Modeling of Probabilistic Quantum Network Policies
Contact: Prof. Patrick Eugster
Co-Supervisors: Anita Buckley and Pavel Chuprikov
Quantum computing, communication, and sensing technologies offer fundamentally new ways for information processing. The objective of quantum communication is to transmit quantum states, which may be entangled, causing stronger correlations. The no-cloning theorem (i.e., qubits cannot be copied) makes quantum communication inherently secure, leading to several novel applications. The distribution of entangled qubits (Bell pairs) between distant end-nodes will be the main task of the quantum Internet of the future, and the main challenge will be scaling. We are developing QNetKAT (Quantum NetKAT), a language and logic for dealing with and reasoning about quantum networks. QNetKAT has primitives for creating and transmitting Bell pairs, together with parallel and sequential composition operators, and offers a simple way for expressing quantum network policies. In the course of this project the student will get familiar with the components of quantum networks and protocols for long distance entanglement distribution. Decoherence, losses,

and noise-errors cause stochastic behavior of quantum operations. The goal of this project is to develop the QNetKAT language with probabilistic semantics. The main tasks will consist in extending the language with new primitives for expressing probabilistic behaviors, and implementing these in the NetSquid quantum network simulation platform (https://netsquid.org/) using Python.

SelfAdapt: Self-supervised Domain Adaptation for Sensor Data
Contact: Prof. Marc Langheinrich
Co-Supervisor: Dr. Martin Gjoreski
Wearable devices, combined with Artificial Intelligence (AI) methods, can bring significant and sustainable improvements to our lives – from improved patient monitoring and decreased healthcare costs to enhanced sports performance and improved quality of life. Standard approaches involve Machine Learning (ML) techniques applied to the data captured from body-worn sensing devices. The ML techniques can be based on classical (feature-based) ML, or Deep Learning (DL) applied on the raw sensor data (end-to-end learning). A typical weakness that all ML-based HAR systems have, regardless of whether they are classical or DL-based, is the domain shift that can be caused by different sensor placements. This project will explore personalization and domain-adaptation techniques to address important challenges in wearable computing: noisy data, limited data, and domain shifts in the labels and the sensor data due to subjectivity. ML processing pipelines (including deep learning techniques) will be augmented with the latest unsupervised and self-supervised learning techniques, including contrastive learning. These advanced techniques should produce more robust and data-efficient models (i.e., requiring fewer person-specific labels). Diffusion-based approaches, could also be considered. Project tasks: 1.Overview of existing self-supervised learning approaches; (ii) Pre-process one dataset from wearable sensing systems. Example datasets include emotion recognition, activity recognition and energy expenditure estimation; (iii) Build baseline ML models using the dataset from step (ii); (iv) Develop self-supervised ML approach and compare self-supervised models with the baseline ML models from step (iii).

Physio-RECALL: Analysis of human memory and physiological signals
Contact: Prof. Marc Langheinrich
Co-Supervisor: Matias Laporte
Memory Augmentation Systems could be capable of selecting specific to-be-remembered events during an experience, e.g., by detecting the individual as distracted or unengaged. Such information then could be used to generate memory cues for the specific periods during which the user was distracted. One way used to detect the cognitive state of users, e.g. if the user is focused on a task, is through physiological signals, like electrodermal activity (EDA) and interbeat interval (IBI). The goal of this thesis is to evaluate the potential of physiological signals captured by a wrist-worn device to detect the cognitive load of individuals presented with a memory task. The final purpose is to use this work as an additional input in future Memory Augmentation Systems. The specific tasks of the project are: (i) Overview methods for physiological signal processing and review memory augmentation literature; (ii) Develop the necessary code for processing the EDA and IBI data; (iii) Design and run an experiment to test the setup and pipeline.

Media-RECALL: Analysis of human memory and audiovisual signals
Contact: Prof. Marc Langheinrich
Co-Supervisor: Matias Laporte
Memory Augmentation Systems could be capable of selecting specific to-be-remembered events during an experience, e.g., by detecting the individual as distracted or unengaged. Such information then could be used to generate memory cues for the specific periods during which the user was distracted. One way used to detect the emotional state of users, e.g. if the user is engaged on a task, is through audiovisual signals, like speech and facial

expressions. The goal of this thesis is to evaluate the potential of audiovisual signals captured by recording devices (such as microphone, cameras) to detect the cognitive load of individuals presented with a memory task. The final purpose is to use this work as an additional input in future Memory Augmentation Systems. The specific tasks of the project are: (i) Overview methods for audiovisual signal processing and review memory augmentation literature; (ii) Develop the necessary code for processing the audio and video data; (iii) Design and run an experiment to test the setup and pipeline.

MultiFed: Multimodal Federated Learning for Sensor Data
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Federated learning and its combination with differential privacy is the latest technique for building privacy-aware machine-learning models. Its primary assumption – no data leaves the local data storage, has enabled its application in a variety of privacy-sensitive domains: mobile keyboard prediction, human mobility modeling based on GPS data, modeling from electronic health records, etc. This project will investigate single modality vs. multi-modality federated models. This is an important issue for wearable sensing systems that utilize multiple sensing devices, e.g., smartphone and smartwatch. Each device, and each sensor in the devices, may have a different availability — data coming from the smartwatch may be unavailable at certain periods (e.g., while charging). To enable the collaborative learning of joint models between users with a variable data/modality availability, we will investigate several multi-modal schemes. Project tasks: (i) Pre-process one dataset from wearable sensing systems. Example datasets include emotion recognition, activity recognition and energy expenditure estimation; (ii) Build centralized multimodal and single-modal models using the dataset from step 1; (iii) Build federated multimodal and single-modal models using the dataset from step (i) and compare them with the centralized models from (ii); and (iv) Develop novel multi-modal federated learning method considering device/sensor availability, computational cost, model accuracy.

XAI-Fed: Explainable AI for Federated Models in Wearable Sensing
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Federated learning and its combination with differential privacy is the latest technique for building privacy-aware machine-learning models. Its primary assumption – no data leaves the local data storage, has enabled its application in a variety of privacy-sensitive domains: mobile keyboard prediction, human mobility modeling based on GPS data, modeling from electronic health records, etc. Artificial Intelligence (AI) methods can bring significant and sustainable improvements to our lives. However, end-users must be able to understand those systems. Unfortunately, today's groundbreaking AI methods are black-boxed (i.e., the decision model and the process are not understandable). The increased complexity of AI algorithms has made previous eXplainable AI (XAI) tools unsuitable, including the fact that most of the XAI solutions are not designed to operate under privacy constraints. This project will investigate XAI techniques compatible with privacy-aware approaches (e.g., federated learning). The focus will be on counterfactual explainers [55] for wearable sensing data. Specific project tasks are: (i) Analyze XAI tools that can operate under privacy constraints, focusing on counterfactuals; (ii) Pre-process one dataset from wearable sensing systems. Example datasets include emotion recognition, activity recognition and energy expenditure estimation; (iii) Develop machine learning models for one of the datasets in step 2, and apply existing XAI tools on the models developed, including the method for generating counterfactual explanations, BayCon; (iv) Develop XAI tool for counterfactual explanations that can operate under privacy constraints.

PrivAffect: Privacy-aware personal-video sensing for affect recognition
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Affective computing is an interdisciplinary field that aims at the development of computer science techniques that enable machines to recognize, understand and simulate human affective states. A fundamental assumption is that different mental states (e.g., emotions and stress), and different intensities of those states, manifest through physiological and behavioral changes. A variety of sensing modalities can capture these changes. Video-based sensing is one promising approach for affect recognition, however, it is also privacy intrusive. Thus, this project will develop a method for privacy-aware personal-video sensing for affect recognition. The method will utilize personal camera (e.g., smartphone or laptop camera) and will include privacy-aware features such as: to record only when the users grant permission, to record only when a specific user is in front of the camera (user identification). Once the video is collected in a privacy-aware manner, existing video-based affect recognition software will be sued to extract affect related information. Project tasks: (i) Review existing software (e.g., GitHub and Google scholar) for user identification, software for counting faces in a video, and software for Facial Action Coding System (FACS); (ii) Implement privacy-aware user identification; (iii) Implement privacy-aware method for extracting Facial Action Units (based on FACS); (iv) Test the overall processing pipeline is a small user-study (e.g., 5 to 10 participants).

Fed-CogLoad: Federated Cognitive Load Estimation
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Federated learning (FL) is a state-of-the-art machine-learning technique developed by Google, where the users' privacy is guaranteed by implementing one simple rule: "No personal data leaves the user-device". This project will investigate FL techniques for cognitive load estimate. Cognitive load can be estimated through the analysis of from pupillometry data, brain activation data (EEG), breathing rate, heart rate, heart rate variability and other related physiological responses. In the project you will: (i) overview excising datasets for cognitive load estimation; (ii) develop a centralized machine learning pipeline for cognitive load estimation; (iii) develop a FL pipeline for cognitive load estimation, (iv) compare the centralized and the FL pipeline and write a short report.

Evaluating Optimization Software as a Service with Application in Biomedical Simulations
Contact: Olaf Schenk
Any sufficiently complex system acts as a black box when it becomes easier to experiment with than to understand. Hence, black-box or derivative-free optimization has become increasingly important as systems have become more complex. In this MaRS project we will analyze Google Vizier, a Google service for performing black-box optimization that has become the de facto parameter tuning engine at Google and other tools such as the DFO optimizer from the Computational Infrastructure for Operations Research (COIN-OR), which is a project that aims to create open and high quality mathematical software. We will review requirements, infrastructure design, underlying algorithms, and advanced features such as transfer learning and automated early stopping that some service provides. These optimization services will be applied to biomedical models that are high-resolution anatomical models of humans created from MRI data. The latest generation of these models allows for detailed 3D anatomical simulations of up to 300 types of tissues and organs. Additionally, morphed anatomical models have been developed to provide variability based on anatomical descriptors like BMI or weight. Realistic biomechanical simulations of specific tissues are used to create these models.

User identification in real-world settings using wearables
Contact: Prof. Silvia Santini
Co-Supervisors: Lidia Alecci, Leonardo Alchieri, Nouran Abdalazim
Wearable devices provide the means for continuously identify users by performing biometric recognition. The highly subject-dependent nature of physiological data (e.g., heart rate, electrodermal activity, skin temperature) makes it possible to identify the user by tracking those traits over time, as it is already done nowadays with the fingerprint or iris recognition. This project aims to understand to what extent user identification is possible and if different condition can hamper that capability. After performing data exploration on an existing data, provided by the supervisor, the student will develop a machine learning model to identify users from wearable data. In addition, he/she will test the efficiency of common anonymization techniques, such as noise addition and synthetic data generation.

EDA Lateralization during sleep
Contact: Prof. Silvia Santini
Co-Supervisors: Lidia Alecci, Leonardo Alchieri, Nouran Abdalazim
Advances in wearable technologies have made possible the ubiquitous sensing of physiological signals, like Blood Volume Pressure (BVP) or ElectroDermal Activity (EDA). However, some signals, like EDA, can change the value that is recorded depending on which side of the body the sensor is placed on. The medical literature has explored extensively this phenomenon, called Lateralization, in both EDA and other physiological recordings, but it is not clear the impact it might have on wearable-based applications. In this thesis, the objective is thus to analyse the lateralization effect from physiological signals recorded while people are sleeping. As such, the first a data collection will be run, to integrate some existing dataset. Then, the student will analyse the raw signal as well as extract some hand-crafted features. A Machine Learning task will be then investigated, to gauge the impact of lateralization, e.g., higher or lower performance of the classifier depending on which side the data is trained on

Personalized Sleep Quality Prediction Model Using Wearable Devices
Contact: Prof. Silvia Santini
Co-Supervisors: Lidia Alecci, Leonardo Alchieri, Nouran Abdalazim
Mobile and wearable devices can nowadays be used as highly available, non-invasive tools to monitor human behavior. Their ubiquitous characteristics encourage their employment in personal health monitoring systems. Such systems provide the user with continuous feedback about daily behavior, productivity, well-being, etc.
Daily sleep is considered a pivotal factor in daily routine due to its role in the rejuvenate of brain and body from accumulated daily basis fatigue. The continuous quantification of sleep quality helps in assessing human health and life patterns. However, the interpersonal variability hinders this task by creating a significant gap between the user-reported sleep quality and the objective sleep quality.
The goal of this thesis is to analyze the impact of different personal characteristics on the user-reported sleep quality. Then, to make use of this knowledge along with physiological signals from wearable devices to build a personalized sleep quality predication model.

A tangible interface for controlling capture and sharing of personal data (RECALL.A)
Contact: Prof. Marc Langheinrich
The goal of this project is to design and prototyping a tangible device for controlling the capture and exchange of “lifelog” data – e.g., photos captured by a wearable camera, or audio recorded by a wrist-worn audio-capture device – with other, co-located peers. The overall vision of lifelogging is that captured experiences can help us remember better our past, and thus improve our cognitive skills and overall memory performance. By supporting the dynamic exchange of such captured data when co-located with other (e.g., in a meeting) we can also have access to our own experiences from someone else’s vantage point. The

to-be-designed device would allow one to control both the capture and the subsequent sharing of experiences not only in a tangible way, but also act as a social marker that would allow all parties involved in a meeting to understand when their discussions would be captured and shared. For example, recording would only take place if the device is placed on a table; exchange with others would only proceed if other devices are placed next to it; shaking the device would delete the last minute of captured data, etc. The first stage of this project will study the different requirements for a privacy friendly data recording device. The second part will be about applying those requirements and constructing a first prototype of a physical recording gadget. Willingness to learn advanced prototyping skills (3D printing, embedded systems development) is required, prior actual experience with embedded systems programming (Arduino or Raspberry Pi) is a plus.

API metrics and pattern visualization
Contact: Prof. Cesare Pautasso
APIs are at the center of microservice architectures as they decouple clients from service providers by explicitly describing the interface contracts connecting them. In this project you will contribute to the development of a tool for statically measuring API specifications (of different languages, e.g., OpenAPI/Swagger/RAML) and visualizing API landscapes so that various microservice API patterns can be observed in the wild.

Application of machine learning in energy market analysis
Contact: Prof. Olaf Schenk
Co-Supervisor: Juraj Kardos
Efficient energy trading relies on high-fidelity price prediction systems with short response times that enable producers, consumers, and traders to make informed decisions in real time. One of the aspects of the problem that provide insight into the market dynamics is understanding how the price perturbations of the underlying fuel prices influence the energy prices. These dynamics can be computed deploying global sensitivity analysis frameworks, which are computationally very expensive tasks. In this project, we will study the potential of artificial neural networks to accelerate the convergence of global sensitivity analysis frameworks, including development of model surrogate aimed to speed up generation and evaluation of the samples, or train a network to predict the final sensitivity indices. A prerequisite for this project is previous experience with applying machine learning techniques to solve scientific problems and experience with major machine learning software packages.

Approximating order-k Voronoi diagrams using clusters of sampled points
Contact: Prof. Evanthia Papadopoulou
The Voronoi diagram is a versatile space partitioning structure finding applications in various areas of science and engineering.
Given a family S of sites, a Voronoi diagram (VD) partitions the plane into regions, one for each site in S. In the nearest VD, all points belonging to a region have the same nearest neighbor. Analogously, in the farthest VD, all points in a region have the same farthest neighbor. A generalized definition is that of the order-k Voronoi diagram, where points belonging to a region have the same k nearest neighbors.
When the sites are not simple objects, and it is difficult to compute the exact nearest Voronoi diagram, a common approach is to sample the sites, compute the nearest VD of all points and keep the edges of the VD which are equidistant to different sites. The larger the sample of the sites the better the approximated diagram resembles the exact one. Unfortunately, this simple and efficient approach works only for the nearest VD and fails for any other order-k VD.
We propose a method of approximating order-k Voronoi diagrams by using Color Voronoi Diagrams (CVD), where each site is a cluster of points. More specifically, the concept is to sample a cluster of points from each site and to construct the order-k CVD, Such a diagram approximates the exact order-k diagram and can be accompanied by theoretical guarantees.

The goal of this project is to implement such construction algorithms and to experiment with different datasets. Work can then be extended to theoretical guarantees for the quality of the approximation or to the design of more efficient algorithms for special input sets.

ASQ
Contact: Prof. Cesare Pautasso
ASQ is a Web platform for delivering interactive lectures in traditional brick and mortar classrooms. It allows presenters to pose questions inside HTML slides and reveive answers from the audience ASQ has grown to a platform that scales to hundrends of students, featuring its own plugin system for advanced question types and HTML5 presentation frameworks (reveal and impress) and a data-driven presenter control view with live statistics. ASQ has been succesfully used both at USI and other universities from which we have gathered a lot of valuable feedback. We are looking for motivated students that will help us make ASQ an indispensable tool for teachers and presenters in general that are interested to ask questions to everyone attending their talk. There are a lot of different areas to work on from data analysis and visualization to live programming and Web presentation development tools.
More information: http://asq.inf.usi.ch

Automated Program Analyses of Student Programs
Contact: Prof. Matthias Hauswirth
Are you interested in programming languages? Would you like to hack compilers or program analysis tools? Then this MARS internship may be for you.
The Luce research group is looking for Master students interested in joining our new research project on Conceptual Change in Learning to Program.
We are developing program analyses that syntactically and semantically analyze student source code. We then run these analyses on code snippets and programs the students submit, to identify flaws in their conceptual understanding of language features. The results of these analyses inform the design of educational materials and assessments and enable educational environments that automatically adapt to the students' current understanding.
This project is related to two other projects offered by Luce ("Qualitative Analysis of Programming Interview Videos" and "Educational Technology for Learning to Program"). Unlike in these other projects, in this project the focus is on program analysis design and implementation development. Knowledge about compilers and programming languages as well as excellent programming skills are a must, and experience writing some kind of program analysis (such as one might get in a compilers course) is an advantage.

Blockchain, blockchain, blockchain
Contact: Prof. Fernando Pedone
Blockchain has gained much traction in recent years. From a topic restricted to specialized circles, it has made it to the general press with daily headlines, including many scandals. Beyond the hype, blockchain fostered the development of sophisticated distributed algorithms and models. And many interesting issues remain unaddressed. In this project, the student will team up with a group of talented researchers from USI and a leading blockchain enterprise to help advance the state of the art in the field.

Collaborative Economy Practices and Communities in Switzerland – A Case Study
Contact: Prof. Marc Langheinrich
The general principle of "collaborative consumption" enables the effective and efficient coordination, acquisition, distribution, and sharing of many kind of different resources, e.g., vehicles, housing, or fertile land. Apart from the well-known for-profit sharing services such as Airbnb, Uber, and TaskRabbit, an increasing amount of community groups and organizations have established not-for-profit cooperatives that often prioritize environmental, social, and cultural values within their local communities.

The goals of the master project are (1) to conduct an empirical research study in the form of in-depth interviews and observations in both commercial and non-for-profit organizations (e.g. in the context of sharing personal artifacts and/or bike sharing); and (2) to compare and contrast sharing practices throughout these services and provide a comprehensive interpretation of the results. Particularly, we would be interested in utilizing practice-based approaches (e.g. Shove's theory of social practice) within the data analysis. This project requires strong analytical skills and a willingness to learn about a novel and emerging research field. Experience with empirical research is a plus (e.g., ethnography, seminar work in human-computer interaction) though supervising guidance is available. Proficiency in any of the Swiss official languages (mostly German, French, or Italian) is an asset.

Compile-time Verification of Fault-tolerant Distributed Systems
Contact: Prof. Patrick Eugster
Co-Supervisor: Dr Pavel Chuprikov
Software defects cost our IT-centered society exorbitant amounts of money. To make matters worse, driven by the advent of paradigms such as cloud computing, blockchains, and the Internet of things, software has been becoming increasingly distributed, i.e., its execution spans many processes. Besides having to avoid "conventional" intrinsic defects in the actual software, programmers now have to cater for partial failures, e.g., the possibility that certain processes or hosts fail while others continue to operate. Catering for these requires complex protocols, leading to highly error-prone code. Traditional "full-depth" verification of programs involve lengthy verification processes requiring much manual effort and expert knowledge, and are thus easily left out of the loop.
We have adapted a very recent technique for partial software verification, so-called session types, to real-life distributed systems, enabling the verification of fundamental properties in the interaction of distributed components in a lightweight fashion integrated with compilation of programs. While striving for partial verification, session types attempt to capture much richer information on programs than typical typing disciplines derived from type theory which focus on data types or other simple properties of program variables or statements. Session types are a form of behavioral typing, which, as the latter name suggests, capture behavioral properties of program code including ordering of operations for interaction between distinct components/processes.
Our session typing discipline is based on an event-driven programming model widely used in distributed systems, and is currently prototyped in Scala. To make type-based verification tractable while dealing with the dynamic nature of asynchronous distributed systems, the model introduces a number of abstractions -- besides sessions/subsessions these include role sets, role parameterization, and role promotion, which the programs adhere to. While an initial evaluation based on applying our typing discipline to the cluster manager core component of Apache Spark shows only moderate performance overheads compared to the unverified vanilla version (<10%), the programming model is still rather basic. The goal of this project will be to apply the prototype to further distributed middleware systems, in the process not only identifying limitations, but also proposing practical extensions for programmers of individual components and implementing runtime optimizations.

Computational Display and Fabrication
Contact: Prof. Piotr Didyk
In recent years, there has been a tremendous increase in the quality and the number of new output devices. Standard 2D screens are being replaced with 3D stereoscopic and multiscopic high-resolution displays. New virtual and augmented reality headsets are being developed to provide new, interactive, and immersive ways to explore the real and the virtual worlds. 3D printers empower millions of users to create, customize, and manufacture digital models. Unfortunately, the abilities of these emerging technologies outperform the capabilities of current methods and tools for creating content.
The Perception, Display, and Fabrication Group at USI Lugano offers a wide range of projects in the field of computer graphics with a common goal of developing novel

computational methods for driving novel display devices (e.g., AR/VR displays) or different digital fabrication devices (e.g., 3D printers). They involve (1) working with novel hardware, such as VR and AR displays or 3D printers, (2) investigating their essential aspects which influence the perceived quality, and (3) devising new content optimization and creation techniques which: maximize quality and human performance while overcoming computational and hardware challenges. Besides using traditional computational approaches, the projects also aim at employing new machine learning techniques, such as deep neural networks, to solve the above challenges. The examples of the projects include but are not limited to foveated rendering for VR and AR displays, content and hardware optimization for varifocal displays, appearance fabrication, and 3D printing surfaces with prescribed haptic properties. For more details on the research topics, please check our website, or contact us via email.

Converting complex polygons into simple polygons
Contact: Prof. Kai Hormann
Most algorithms in computational geometry for 2D polygons (e.g. polygon partition, Boolean operations, and offset curves) assume the input to be simple, that is, without self-intersections. Your task will be to develop and to implement an algorithm that takes as input an arbitrary complex polygon, with possible self-intersections and degeneracies, and generates as output a set of simple polygons that together enclose the same interior than the complex polygon. The ideal candidate for this project has good programming and mathematical skills and enjoys geometric thinking and reasoning.

Cost-based Mechanism Selection for Secure Cloud Computing
Contact: Prof. Patrick Eugster
Co-Supervisors: Shamiek Mangipudi, Dr Pavel Chuprikov
The value of what can be derived from customer data is being increasingly recognized by many industries, information is becoming the new currency. At the same, the amount of data generated has been growing exponentially, and many organizations have turned to the cloud in their search for cost-effective information processing. On the account of that, there is a huge demand for processing of sensitive data using third-party untrusted computational resources. While there are both software (homomorphic encryption) an hardware (secure enclaves) techniques with the potential to perform such processing without leaking any information, they have their own constraints and overheads, so that there is no single universal solution. For the best performance different techniques must be combined, but it quickly becomes hard to reason about end-to-end security for non-experts, which data analysts writing queries usually are not.
We have designed a system, called Hydra, that supports a multitude of security mechanisms nicely decoupling privacy policies from the business logic of the queries. To guarantee compliance with a chosen privacy policy, Hydra introduces a lambda-calculus-based domain specific language (DSL), equipped with a type system which ensures the absence of insecure information flows in the system — especially valuable with data propagating through different heterogeneous components/security mechanisms. Our system Hydra is based on Apache Spark — one of the hottest open source projects boasts contributions from over 1200 developers spread across 300 companies, making it a unique code base to learn and experience the length and breadth of system design principles. More precisely Hydra is integrated with the Apache Spark SQL processor to provide users with the familiar query abstraction. The current limitation of Hydra is that the choice of the security mechanism is hardcoded.
In the course of project, the selected student will be working on the query optimization part of Hydra (which uses standard Spark SQL extension points). The goal will be to use empirical performance measurements for different security mechanisms combined with Spark execution metrics in order to develop heuristics that guide the mechanism choice at the query transformation phase.

Design and Evaluation of a Smartphone App to Support Sharing Physical Objects (SHA21.C)
Contact: Prof. Marc Langheinrich
Many persons are willing to contribute to the community by sharing objects they own, such as household items, tools and media items. The goal of the project is to develop an application that supports this by connecting lenders and borrowers in an easy way. Beyond the obvious features of a standard "classifieds" app (i.e., a potential lender provides information about objects that he is willing to lend; the application provides an easy way for the borrower to find the items he is interested in and, once found, contact the lender) the app should focus on supporting the interaction inherent in such physical lending, i.e., the physical exchange of the item in question. For example, a simple "bump" could allow a borrower to acknowledge receipt of an item, or an embedded NFC tag in the item itself could be used. Optionally, the app should be evaluated within a short study with several participants. Willingness to learn qualitative research methods in Human-Computer Interaction, as well as basic iOS programming skills required; strong Web programming skills an asset. Hardware such as a mobile phone and a smartwatch will be provided.

Educational Technology for Learning to Program
Contact: Prof. Matthias Hauswirth
Are you interested in web applications? Would you like to contribute to sites with highly interactive components, rich visualisations, and powerful backends? Then this MARS internship may be for you.
The Luce research group is looking for Master students interested in joining our new research project on Conceptual Change in Learning to Program.
We are developing novel learning platforms, from clicker systems that allow students to produce open-ended answers to questions, to platforms that enable mastery learning, boost metacognitive development, and provide novel collaborative and social learning support. We want to extend these platforms to better capture and track the conceptual understanding of students who are learning to program. This work will enable the construction of more efficient and effective learning environments.
This project is related to two other projects offered by Luce (“Automated Program Analyses of Student Programs” and “Qualitative Analysis of Programming Interview Videos”). Unlike in these other projects, in this project the focus is on web development. Knowledge of JavaScript is a must, and experience developing with React is an advantage.

EEG-RECALL: Analysis of the influence of distractions on human memory EEG Signals and distraction
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Memory Augmentation Systems could be capable of selecting specific to-be-remembered events during an experience, e.g., by detecting the individual as distracted or unengaged. Such information then could be used to generate memory cues for the specific periods during which the user was distracted. One way used to detect the cognitive state of users, e.g. if the user is being interrupted, is through physiological signals, like electroencephalography (EEG). The goal of this thesis is to evaluate the potential of EEG Signals captured by an OpenBCI device to detect distractions of individuals presented with a memory task. The final purpose is to use this work as an additional input in future Memory Augmentation Systems. The specific tasks of the project are: (i) Overview methods for EEG signal processing and review memory augmentation literature; (ii) Develop the necessary code for processing the EEG data; (iii) Design and run an experiment to test the setup and pipeline.

End-to-end security based on heterogeneous mechanisms
Contact: Prof. Patrick Eugster

Co-Supervisors: Shamiek Mangipudi, Dr Pavel Chuprikov
The value of what can be derived from customer data is being increasingly recognized by many industries, information is becoming the new currency. At the same, the amount of data generated has been growing exponentially, and many organizations have turned to the cloud in their search for cost-effective information processing. On the account of that, there is a huge demand for processing of sensitive data using third-party untrusted computational resources. While there are both software (homomorphic encryption) and hardware (secure enclaves) techniques with the potential to perform such processing without leaking any information, they have their own constraints and overheads, so that there is no single universal solution. For the best performance different techniques
must be combined, but it quickly becomes hard to reason about end-to-end security for non-experts, which data analysts writing queries usually are not.
We have designed a system, called Hydra, that supports a multitude of security mechanisms nicely decoupling privacy policies from the business logic of the queries. To guarantee compliance with a chosen privacy policy, Hydra introduces a lambda-calculus-based domain specific language (DSL), equipped with a type system which ensures the absence of insecure information flows in the system — especially valuable with data propagating through different heterogeneous components/security mechanisms. Our system Hydra is based on Apache Spark — one of the hottest open source projects boasts contributions from over 1200 developers spread across 300 companies, making it a unique code base to learn and experience the length and breadth of system design principles. More precisely Hydra is integrated with the Apache Spark SQL processor to provide users with the familiar query abstraction. The current limitation of Hydra is that the choice of the security mechanism is hardcoded.
The different security mechanisms used in Hydra — both hardware, such as enclave-based mechanisms, and software, such as partially homomorphic cryptographic primitives which span both symmetric and asymmetric crypto schemes — provide different performance characteristics and guarantees under different attacker models. The goal of this project is to extend Hydra with further security mechanisms that it can choose from, such as hardware-based ones (e.g., AWS Nitro enclaves, AMD SEV) and/or software-based ones (e.g., novel symmetric partially homomorphic encryption schemes).

Evaluating a Human Memory Augmentation App (RECALL.C)
Contact: Prof. Marc Langheinrich
Technology always had a direct impact on what humans remember. In the era of smartphones and wearable devices, people easily capture information such as pictures and videos on a daily basis which can help them evoke memories for reminiscing or simply to help one remember a past event. The ubiquity and increasing use of such devices and technologies produce a sheer volume of pictures and videos that, in combination with additional contextual information, could significantly improve one's ability to recall a past experience or future event. The successful MSc student candidate will design, develop, deploy and evaluate a mobile app that collects and displays such information on mobile devices (Android and/or iOS) for assisting memory recall. As part of the MSc project, the student may help fine-tune and improve an existing set of software (alternatively develop novel modules as needed), combine them with state-of-the-art hardware (E4 wristband, Narrative Clip 2, last-generation smartphone) into a fully functional prototype, and trial the entire system in a week-long study (which includes participant recruitment, participant briefing and de-briefing, data collection, and data analysis). The student will be part of an international research project and will gain invaluable experience in running scientific studies with real participants in real settings. Basic mobile programming skills (Android and/or iOS) and willingness to learn about human subject research required; knowledge of human-computer interaction methodologies and visual design skills are a plus.

Evaluating Secure Personal Memory Sharing with Co-Located People (RECALL.B)
Contact: Prof. Marc Langheinrich

Using smartphones and wearable devices people now can fully log their life in pictures, audio or even video recordings. Such data - "life-logs" - can help evoke past memories and potentially improve our overall cognitive abilities. One interesting opportunity in highly networked environments is the ability to share parts of one's life-logs with others, in order to benefit from recordings of each-other (e.g., by having access to a third person view of oneself in a meeting). In order to avoid any privacy violations, life-logs should only be shared with co-located people – as soon as people leave or join the meeting, the exchange of lifelogs should be stopped or initiated, respectively. In prior work we have designed an initial prototype of this system, running on several Nexus 5X smartphones. The aim of this project is to evaluate the Android app - called ‘MemShare’, in order to understand usability requirements and use, and to further refine the overall system (including the backend server and web-based control and inspection tools) based on collected feedback and observed use. This project required a highly motivated student that will co-design, develop, and trial the MemShare system. Basic Android programming skills and willingness to learn about human subject research required; knowledge of human-computer interaction methodologies and visual design skills are a plus.

Formal Analysis of Smart Contracts
Contact: Prof. Natasha Sharygina
Transactions of smart contracts in a decentralized network (e.g. the Ethereum blockchain) are executed by miners that execute the contract functions requested by each transaction sender in exchange for a fee. The fee is proportional to the amount of computing resources necessary to complete the transaction and it is expressed in an abstract quantity called gas. The sender specifies the maximum amount of gas for each transaction and miners execute the transaction code until such limit is reached. Miners either complete the transaction returning the unused gas back to the sender, or abort the transaction when the execution exceeds the gas limit. In both cases miners keep the actual gas used for the execution as a compensation for mining the transaction. In general the cost of a transaction depends on the unknown state of the contract, exposing the sender to the risk of setting the gas limit not high enough to complete the transaction, and therefore leading to a sure money loss. Furthermore contracts cannot be changed once deployed, and predicting the gas needs considering all future possible scenarios the contract could get is a hard task. The goal of this project is to reduce the problem of finding the worst-case gas consumption of contracts to the optimization problem in the Satisfiability Modulo Theories (SMT) context. The SMT problem is the decision problem of determining whether a logical formula is satisfiable, given that some of the variables have an interpretation with respect to combinations of first-order background theories. The work, including implementation and experimentation, will be carried out as an extension of frameworks and tools such the novel parallel version of SMT solvers OpenSMT (http://verify.inf.usi.ch/opensmt) and the C model checker HiFrog (http://verify.inf.usi.ch/hifrog) currently being developed by the Verification Group at USI, and the induction-based model checker Z3-Spacer (https://github.com/Z3Prover/z3) by Microsoft research and others. We are looking for a motivated student who wants to improve his/her knowledge on software verification applied to smart contracts. This project will give the student an excellent overview of a quickly developing field while being sufficiently approachable. Prior knowledge of C++ and Solidity languages is desirable.

Fed-CogLoad: Federated Cognitive Load Estimation
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Federated learning (FL) is a state-of-the-art machine-learning technique developed by Google, where the users’ privacy is guaranteed by implementing one simple rule: “No personal data leaves the user-device”. This project will investigate FL techniques for cognitive load estimate. Cognitive load can be estimated through the analysis of from pupillometry data, brain activation data (EEG), breathing rate, heart rate, heart rate variability and other related physiological responses. In the project you will: (i) overview excising

datasets for cognitive load estimation; (ii) develop a centralized machine learning pipeline for cognitive load estimation; (iii) develop a FL pipeline for cognitive load estimation, (iv) compare the centralized and the FL pipeline and write a short report.

Federated Clustering
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Federated learning (FL) is a state-of-the-art machine-learning technique developed by Google, where the users' privacy is guaranteed by implementing one simple rule: "No personal data leaves the user-device". By default, FL has been developed for supervised learning (e.g., classification models). This project will investigate FL techniques for unsupervised learning (e.g., clustering models). For the project you will: (i) overview methods for distributed clustering and federated clustering; (ii) implement one federated clustering method and evaluate it in iid and non-iid settings; (iii) summarize the results and propose future work

Geometric properties of indirect Pythagorean hodograph curves
Contact: Prof. Kai Hormann
An offset curve is a curve of fixed distance away from a given planar curve along its normal direction. They arise in a variety of applications, which include CNC machining, railway design and shape blending. However, for polynomial curves that are widely used in computer-aided design and manufacturing, their offset curves may not be polynomial or rational. Only a special set of polynomials curves has rational offsets, and they are called "Pythagorean hodograph curves". For these curves, there exists a polynomial such that this polynomial, together with the x- and the y-component of the derivative (or hodograph) of the given curve form a "Pythagorean triple", hence the name. An extension of this concept are those curves that do not have a polynomial Pythagorean hodograph, but can have a rational Pythagorean hodograph, after a properly chosen reparameterization, and hence they are called indirect Pythagorean hodograph curves. While the geometric properties of Pythagorean hodograph curves are well understood by now, little is known so far about indirect Pythagorean hodograph curves, and the aim of this project is to change that. We are therefore looking for a highly motivated student with a thorough background in Mathematics (esp. Algebra, Analysis, Geometry, and Numerics) and basic programming skills.

Highly parallelizable public blockchains
Contact: Prof. Patrick Eugster
Distributed ledgers have recently emerged as a promising technology that can strengthen the trust relationships between its users. These ledgers are conceived as auditable systems that record the actions of the participants to enable accountable behaviors. Blockchains, in particular, are highly replicated distributed ledgers that rely on cryptographic primitives to prevent record tampering. Blockchains are often used in public environments designed specifically not to contain any authority that on one hand can tremendously help with the performance of the system, but on the other hand might compromise the security of the system due to its greater power.
The lack of authority in public blockchains forces its participants to rely on agreement protocols to establish the state of the blockchain. These agreement protocols have prohibitive costs when deployed on large-scale networks and must yet be instanced every time new records are added to the chain. One approach to increase the throughput of such large-scale systems is called sharding: the system is divided into several partially independent sub-systems to limit the scope of coordination happening on the entire system. The selected student will take part in designing a sharded public blockchain, and performing evaluations to assess its performance. Understanding the inner workings of distributed systems and that of public blockchains in particular (e.g., Bitcoin, Ethereum), including their incentive mechanisms, is a clear advantage.

Investigating the dichotomy of sharing practices in virtual and physical realms: from theoretical overview to design considerations (SHA21.A)
Contact: Prof. Marc Langheinrich
Online social networks have made sharing personal experiences with others - mostly in form of photos and comments - a common activity. Nowadays the scope of user-generated and shared content on the net varies vastly from personal media to individual preferences to physiological information (e.g. in form of daily workouts). Popular "sharing economy" services (e.g. AirBnB, Uber) and connected devices are expanding the set of "things" one can share. Given that a new generation of sharing services is about to emerge, it is of crucial importance to understand how traditional sharing practices inform and support designers of those services. This project will look into consolidating the existing body of work on both sharing personal digital content (e.g., on social networking sites, through photo sharing apps) and personal physical possessions (e.g., apartment sharing). The project aims to identify commonalities and differences between digital and non-digital context sharing, in particular summarizing existing research on motivations to share, audience management, privacy and trust issues, and user experience requirements. If possible, these findings should be connected to contemporary theories of social psychology and practice theory. The final results of this project would be: (1) a comprehensive account of the existing body of knowledge on content and resource sharing practices; (2) a set of design considerations that allows designers and developers to build future sharing services to enable sharing activities bridging virtual and physical realms. This projects requires strong analytical skills and a willingness to learn about a novel and emerging research field. Experience with interdisciplinary research literature a plus (e.g., seminar work in human-computer interaction) though supervising guidance is available.

Just Share It: A Decentralized Autonomous System to Support Sharing Physical Objects Using Blockchain and Smart Contracting
Contact: Prof. Marc Langheinrich
Many persons are willing to contribute to the community by sharing objects they own, such as household items, tools and media items. For example in Switzerland an online service pumpipumpe.ch provides a set of stickers for a mailbox to let people see what household items one can borrow from their neighbours. However, the service does not support the actual act of sharing those items – how borrower and lender meet, agree, and exchange. "Just Share It" is meant as an application that provides such a service, by connecting lenders and borrowers through mobile technology.
A potential lender provides information about objects that he is willing to lend. The application provides an easy way for borrowers to find the items they are interested in. Once a borrower has found an item, the application provides a way for the lender and the borrower to communicate and come to an arrangement. An underlying layer of blockchain-driven smart contracting technology facilitates online contractual agreements (e.g., to record sharing transactions). In order to maintain a positive and friendly environment, “Just Share It” will also need to provide means by which users can build up trust. For example, the application should allow the borrower to leave a short feedback about the experience with the item in the form of a short notice and a picture.
This project is part of the Swiss National Science Foundation funded SHARING21 research project, where we are looking into new ways of supporting sharing both digital information and physical objects. The goal of this master project is to implement and eventually evaluate the "Just Share It" application on a mobile platform, incorporating blockchain and smart contracting technologies (e.g., using Etherium, an open-source distributed computing platform). Strong Web programming skills (Ajax frameworks) are required, basic Solidity and/or C++ programming skills are an asset, prior experience with iOS and/or multi-platform smartphone app development (e.g. Ionic framework) is desirable. Hardware such as a mobile phone and a smartwatch will be provided.

On the farthest-segment Voronoi diagram: predicates and robust computation
Contact: Prof. Evanthia Papadopoulou
The Voronoi diagram is a versatile space partitioning structure with numerous applications in diverse areas of science and engineering. The basic concept is simple: given n simple geometric objects in a space, called sites, their Voronoi diagram divides the space into regions such that the Voronoi region of a site s is the locus of points closer to s than to any other site.
In this project we will focus on the farthest Voronoi diagram of line segments and lines in the plane. It is well known that the farthest-segment Voronoi diagram has a tree structure of complexity linear in the number of the input segments. We want to study the algorithmic predicates involved in computing this diagram robustly. If successful this project may lead to a new package in CGAL - the Computational Geometry Algorithms Library, https://www.cgal.org. The project, however, will focus on the construction algorithm and its predicates.
The ideal candidate must have good programming skills, good algorithmic skills, and an interest in geometric computing. Some mathematical skills are also a plus.
More information about the algorithm and robust computation.

Model-predictive control in power systems with renewables
Contact: Prof. Olaf Schenk
Co-Supervisor: Juraj Kardos
Due to transition to clean energy sources including renewables, the operation planning of the power grid becomes increasingly complex problem that needs to consider very long time horizons. The resulting computational complexity of these very large problems becomes prohibitive using classical optimization approaches. A typical solution is to apply various heuristics, one of the widely adopted being the model-predictive control. Instead of solving the full optimization problem over the entire time horizon in one coupled solve, several smaller problems with shorter time spans are solved iteratively until the desired prediction range is covered.
In this project, we want to evaluate different systems configurations of the control algorithm, including various presolve and multi-scale strategies. The improvements of the fine-tuned algorithm will be evaluated using metrics such as the objective function of the underlying problem, the error introduced by the heuristic, or the improvements in terms of the final operational costs. An interesting extension of the problem could be designing an auto-tuning system strategy based on e.g. an artificial neural network that would predict optimal settings for the algorithm.

Neural Style Transfer-based Testing of Autonomous Driving Systems
Contact: Prof. Paolo Tonella
Co-Supervisor: Andrea Stocco
DNN-based autonomous driving systems are tested using either (1) physics-based simulators that model the interactions with the environment; or, (2) data-driven simulators that use GANs to use or reproduce a real-world stream of driving data. The former approach lacks photo-realism, whereas the second approach is immaterial with respect to the physical interactions with the environment.
The goal of the thesis is to tackle the gap between these two approaches and investigate the potential for hybridization, by implementing the possibility of using real-world data within a physics-based simulator. Approaches such as CycleGAN will be adopted to process simulated driving data and translate them into real-world driving data (from a given distribution), which will be sent back to the simulator. The results of the enhanced simulations will be validated through in-field testing with a real small scale self-driving car based on the Donkey Car framework.
In summary, the questions the thesis will aim to answer are: Can we enhance a driving simulator to use real-world data? Do the testing results obtained using such enhanced

simulator correlate with those of in-field testing? The results of the thesis are expected to increase the degree of support available to engineers in self-driving cars development and testing, and give the student experience on the following topics: supervised learning, neural translation, sim2real testing, empirical analysis.

Online data-center modeling
Contact: Prof. Robert Soulé
The open position will focus on the design of a common data model and representation for the state of an operational data center. The model will be populated and driven by logs, traces, and configuration information; queried by operators to determine global properties of the system (such as traffic matrices), and drive online workload-driven simulations to explore the effects of configuration changes. The open position will involve research in data representation, language design, and distributed data processing.

Physio-RECALL: Analysis of human memory and physiological signals
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Memory Augmentation Systems could be capable of selecting specific to-be-remembered events during an experience, e.g., by detecting the individual as distracted or unengaged. Such information then could be used to generate memory cues for the specific periods during which the user was distracted. One way used to detect the cognitive state of users, e.g. if the user is focused on a task, is through physiological signals, like electrodermal activity (EDA) and interbeat interval (IBI). The goal of this thesis is to evaluate the potential of physiological signals captured by a wrist-worn device to detect the cognitive load of individuals presented with a memory task. The final purpose is to use this work as an additional input in future Memory Augmentation Systems. The specific tasks of the project are: (i) Overview methods for physiological signal processing and review memory augmentation literature; (ii) Develop the necessary code for processing the EDA and IBI data; (iii) Design and run an experiment to test the setup and pipeline.

Prioritising Test Inputs for Deep Learning Systems via Mutation Analysis
Contact: Prof. Paolo Tonella
Co-Supervisor: Nargiz Humbatova
The recent success of Deep Learning (DL) in performing complex, human-competitive tasks, such as artificial vision, speech recognition and natural language processing, are making DL based components an integral part of advanced software systems. When such systems involve life, business or ethics critical activities, the quality of the DL components they use becomes a major concern. Due to this, there is an increasing interest in the research community for different techniques designed to assess the quality and reliance of DL-based systems.
Typically, the correctness of the prediction power of a neural network is evaluated by using a set of inputs called test set that the network has not seen while being trained. One of the biggest challenges related to the process of compiling an adequate and comprehensive test set is the cost of the manual labelling of new inputs. In fact, in order to be able to cover all the main scenarios and corner cases, a test set is often composed of a large number of inputs. Moreover, as the labelling is a manual process, it usually requires the involvement of multiple persons with domain specific knowledge to ensure the correctness of the results. Thus, there is an increased need of approaches that would prioritise the inputs for labelling, to facilitate a faster and efficient testing pipeline.
Mutation testing is a testing technique that deliberately seeds faults in the form of small syntactic changes into the program under test to create a set of faulty programs called mutants. The general principle underlying this approach is the assumption that faults used by mutation testing represent the mistakes that programmers usually make. As mutation testing aims to assess the quality of a given test suite (test set in case of DL systems) in terms of its

capability to detect faults, it can serve as a great measure to evaluate and rank test inputs in the input prioritising problem.
The goal of the project is to improve the state of the art in test input prioritisation by designing and implementing a new tool based on assessing the 'usefulness' of test inputs via mutation analysis. Such a tool would employ the most efficient mutation operators available in the literature to reveal the mutation killing ability of test inputs and integrate novel approaches such as Explainable Artificial Intelligence (AI) for smarter input mutation that would give developers the intuition on to which degree the information provided in the input is utilized by a model. The characteristics of inputs obtained through such an analysis would enable the knowledgeable ranking of inputs for labelling. The last step of the project would be to compare the performance of the newly implemented tool with the existing ones.

PrivAffect: Privacy-aware personal-video sensing for affect recognition
Contact: Prof. Marc Langheinrich
Co-Supervisor: Martin Gjoreski
Affective computing is an interdisciplinary field that aims at the development of computer science techniques that enable machines to recognize, understand and simulate human affective states. Video- based sensing is one promising approach for affect recognition, however, it is also privacy intrusive. Thus, this project will develop a method for privacy-aware personal-video sensing for affect recognition. The method will utilize personal camera (e.g., smartphone or laptop camera) and will include privacy-aware features such as: to record only when the users grant permission, to record only when a specific user is in front of the camera (user identification). Once the video is collected in a privacy-aware manner, existing video-based affect recognition software will be sued to extract informative. The specific tasks of the project are: (i) check existing software (e.g., on GitHub) for user identification, software for counting faces in a vide, and software for Facial Action Coding System (FACS) [1]; (ii) implement privacy-aware user identification; (iii) implement privacy-aware method for extracting Facial Action Units (based on FACS); (iv) test the overall processing pipeline is a small user-study (e.g., 5 to 10 participants).

Production optimization through water-front control using adjoint gradient-based techniques
Contact: Prof. Olaf Schenk
The optimization of oil production is a tedius and computational intensive process that requires the solution of time dependent nonlinear set of partial differential equations describing the flow of hydrocarbons in anisotropic porous media. Optimization of production is usually performed using either gradient free techniques like genetic, particle swarm algorithms, or gradient-based techniques where the gradients are computed through the solution of the adjoint problem. Optimization using gradients converges much faster than gradient-free techniques resulting in significant saving in computational time but it usually gets trapped to poor local optima. It is known that the optimal solution of the production optimization problem in homogeneous reservoirs requires equal arrival times of the water-front from the injector wells to the production wells. The aim of this project is to achieve production optimization by a redefinition of the objective function, which is usually defined to be the cumulative oil recovery, so that water-fronts can be controled directly to arrive simultaneously at the production wells. This project requires a highly motivated student that will co-design, develop, and implement the adjoint gradient-based method for the particular objective functions in a compositional reservoir flow simulator. Strong C++ programming skills are required as well as experience in reservoir simulation and compositional flow models.

Qualitative Analysis of Programming Interview Videos
Contact: Prof. Matthias Hauswirth

Are you interested in qualitative research, and specifically in trying to “debug” people’s conceptual understanding (or misunderstanding)? Are you a proficient programmer, and like to help others in learning to program? Then this MARS internship may be for you.
The Luce research group is looking for Master students interested in joining our new research project on Conceptual Change in Learning to Program.
We are conducting “mastery checks” of students in undergraduate programming courses. The students are asked to explain concepts, and to perform various programming tasks. We record these checks on video, over the course of an entire semester. We then use modern qualitative data analysis software to study the videos and to analyze how the students’ conceptual understanding of specific programming language features and of programming strategies changes over time. The goal is the identification of learning trajectories which will inform the design of educational materials and assessments.
This project is related to two other projects offered by Luce (“Automated Program Analyses of Student Programs” and “Educational Technology for Learning to Program”). Unlike in these other projects, in this project the focus is on qualitative research. Interest in gaining a profound understanding of how people learn difficult concepts is a must, and experience using qualitative data analysis methods (such as one might get in a qualitative research course) is an advantage.

Query optimization for graph databases
Contact: Prof. Robert Soulé
The position will focus on investigating language and system support to optimize graph database queries. The current focus is on interaction graphs, which are append-only, temporal graphs used for analytics in telecommunications, transportation, and social media.

RESTful conversation mining
Contact: Prof. Cesare Pautasso
Given a set of logs tracking the HTTP interactions of multiple clients with a REST API, the project will develop a tool to analyze the logs and build a representation of the conversation between each client as it uses the API. The representation should use the RESTalk visual notation (an extension of BPMN for modeling RESTful conversations). Additionally, the mining could be extended to detect recurring conversation fragments and patterns.
More information: http://design.inf.usi.ch/publications/2015/ecsa

SAT-based techniques for Approximate Circuit Design
Contact: Prof. Laura Pozzi

As energy efficiency becomes a crucial concern in every kind of digital application, a new design paradigm called Approximate Computing (AC) gains popularity as a potential answer to this ever-growing energy quest. AC provides a different view to the design of digital circuits, by adding _accuracy_ to the set of design metrics.

So, while traditionally one could sacrifice area for delay, for example, or energy for area, etc, now the idea is to play with accuracy also, and pay a small loss in accuracy for a large improvement in energy consumption. This is particularly suited for error-resilient applications, where such small losses in accuracy do not represent a significant reduction in the quality of the result. While Approximate Computing can be applied at different levels -- from software to hardware -- in our group we are particularly interested in the design of approximate boolean circuits. In particular, we are research Approximate Logic Synthesis, which is the process of automatically generating, given an exact circuit and a tolerated error threshold, an approximate circuit counterpart where the error is guaranteed to be below the given threshold. The resulting circuit will be a functional modification of the original one, where parts will be substituted, or even completely removed.

While various algorithms have been proposed -- in and out of our group -- for the design of approximate circuits, we are currently exploring new SAT-based solutions. The SAT (or boolean satisfiablity) problem states the following: given a formula containing binary variables connected by logical relations, such as OR and AND, SAT aims to establish whether there is a way to set these variables so that the formula evaluates to true. If there is, the formula is SAT; if there isn't, the formula is UNSAT.

An astonishing number of problems in computer science can be reduced to the SAT problem -- including our approximate circuit design question -- and, in addition to this, astonishingly fast SAT solvers exist.

Hence, in this project we aim at designing (and improving our existent) SAT-based formulations and algorithms for circuit design, in order to generate ever more efficient approximate circuits.

Useful links:

An introduction to SAT: https://www.borealisai.com/en/blog/tutorial-9-sat-solvers-i-introduction-and-applications/

An example of approximate circuit design technique:
https://ieeexplore.ieee.org/document/8342067

A survey of approximate circuit design techniques:
https://www.inf.usi.ch/phd/scarabottolo/papers/ALS_survey.pdf

https://www.inf.usi.ch/faculty/pozzi/

Scalable State Machine Replication
Contact: Prof. Fernando Pedone
State Machine Replication (SMR) is a well-established replication technique used by many production systems, including Apache Zookeeper, Google Chubby, Windows Azure storage, Google Spanner, and many others. Scalable State Machine Replication (S-SMR) is a recent extension of SMR developed at the distributed systems group at USI that promises unlimited performance in addition to configurable fault tolerance. Some initial efforts, for example, resulted in a prototype that outperforms Zookeeper by almost an order of magnitude. This project will look into various aspects of S-SMR and contribute to cutting-edge research with high prospects of applicable results within a team of highly motivated and talented students.

Smart Group Activity Journal – Creating an Automated Activity Feed for Outdoor Sports (SHA21.D)
Contact: Prof. Marc Langheinrich
Skiing and snowboarding are highly social activities. Winter enthusiasts capture and share vast amount of pictures and videos during outdoor vacations. To support information exchange among groups of skiers and snowboarders, this project seeks to create a semi-automated group activity journal. The journal would be automatically shared among group members with an option to grant access to external observers who want to follow a particular participant or the entire group in near real time. The student should implement an app that allows one to both post to and visualize a shared "event stream". Events posted should be supported with a simple plug-in system, e.g., one could imagine a "slope tracker" that posts to the stream whenever one has finished a ski run. Types of content added to the shared stream should include, but are not limited to (1) a pin on a map with geolocation information to setup a meeting point while on the slope; (2) captured media content (photos, videos, an optional live-stream) of the run; (3) reference information in a form of text necessary for

descent (e.g., the time left until sunset, the operational hours of a ski lift at a particular location, conditions of the slope with detailed information about potential hazards during descent). Following the recent trend in instant messaging services such as Snapchat or iMessage, where a message can expire after some time, optionally one would be able to associate an expiration tag to any events added to the feed. An additional smartwatch interface should support quick entry of items, e.g., one could post a "hazard" on the slopes by selecting a hazard type and the system automatically adding time and location. The shared activity journal will be hosted on a central server and will be accessible by group participants or observers through the app. An optional integration with an optical head-mounted display for augmented reality (RideOn, Recon Snow 2) is encouraged. Intermediate to strong iOS/Android programming skills required, strong Web programming skills an asset. Hardware such as a smartphone, smartwatch, and augmented reality gear will be provided.

Supervised vs reinforcement learning for the training of the driving agent of a self driving car operating in a real and in a simulated environment
Contact: Prof. Paolo Tonella
Co-Supervisor: Matteo Biagiola

Self-driving cars have been in development for several years and nowadays they are slowly being deployed. Recent tests show that autonomous cars can drive without human intervention for many miles under certain conditions.
The key technology that led to the deployment of autonomous vehicles is deep learning. In particular, in such domain, given the presence of large and diverse datasets of human-annotated data, the most prevalent paradigm is Supervised Learning (SL), where the car is trained to reproduce human behaviour. However, it can be argued that the disadvantage of the SL paradigm is that the agent (i.e. the car) learns to predict how to steer which does not necessarily mean that it learns how to drive. In particular, the SL approach does not take into account the effect that each decision has on the the environment the agent operates on, which in turns influences future decisions.
This drawback can be solved by another prominent learning approach, i.e. Reinforcement Learning (RL). RL, in fact, allows learning a policy, thereby creating agents that are able to make their own decisions, take actions, react and adapt based on the feedback they receive from the environment. Since the task of driving is sequential in nature, our hypothesis is that autonomous driving can be better addressed when formalized as a RL problem rather than a SL problem.
The interaction with the environment can be also seen as a drawback for RL because it might be time-consuming, especially in the real-world. However, it can be shown that, in a self-driving car simulated environment (i.e. the DonkeyCar simulator [1]), when the RL agent is appropriately pre-trained, the training time of a RL agent is comparable to the one of a SL agent. The goal of this project is to train a RL agent on a real car (i.e. the DonkeyCar [2]) to drive on a small scale track we have in our lab. Once a RL agent is successfully trained we can rigorously analyze the Pros and Cons of each learning approach (RL and SL) both in simulation and in the real world.

[1] https://docs.donkeycar.com/guide/simulator/
[2] https://www.donkeycar.com/.

Testing Deep Learning Systems with Generative Models
Contact: Prof. Paolo Tonella
Co-Supervisor: Andrea Stocco
Deep Learning (DL) systems are successfully applied in a range of applications, including safety-critical ones (e.g., autonomous driving). Unlike traditional software systems, in which developers explicitly program the systems' behaviour, DL entails techniques that mimic the human ability to learn how to perform tasks through training examples. For each input, a DL

system will behave depending on how it was trained. Hence, it is particularly interesting to discover how a DL system handles inputs beyond the ones used for training.
Most of the existing test input generation techniques directly modify raw input data (e.g., change values of an image's pixels). However, these techniques often generate inputs that do not represent the real world. Instead, model-based approaches manipulate a model of the input that is then used to derive the actual raw data. In this way, model-based approaches can generate more realistic inputs. However, model-based approaches are only applicable when a model can be obtained for the inputs in the given domain.
Deep Generative models based on latent variables, such as Generative Adversarial Networks (GANs) and Variational Auto-Encoders (VAEs), are a novel family of solutions that allow to obtain and leverage the probability distribution of the input space. Generative models allow generating new test inputs without needing to directly manipulate raw inputs or defining a model of the input space.
The goal of this Master Thesis is to implement novel test approaches that exploit generative models for generating test inputs. The student will face a fundamental problem in research and industry: generating meaningful test inputs for DL systems in complex domains.

The medial axis of a simple polygon in linear time
Contact: Prof. Evanthia Papadopoulou
The medial axis transform (MAT) is a well-known shape descriptor that is used in diverse scientific areas. Given a simple polygon P, its medial axis is the part of the Voronoi diagram of the sides of P which lies inside P. If the polygon P is convex, its medial axis can be computed in linear time by a very simple randomized incremental algorithm, known since the late 80s. If the polygon P is simple, its medial axis can still be computed in linear time, however, the existing algorithms, which were developed in the late 90s, have been too complicated to be of practical use. Recently, we have generalized the simple randomized approach for convex polygons to general simple polygons. The randomized incremental algorithm remains surprisingly simple, almost as simple as the original technique for convex inputs, but it can still compute the medial axis of any simple polygon in linear time. We would like to bring this elegant, new technique to life through a MaRS project in computational geometry. The technique is applicable to all kinds of Voronoi diagrams with a tree structure, thus, several extensions are possible for the interested student. The ideal candidate would combine good algorithmic, programming, and analytical skills while having an interest in geometric problems. Ideally the implementation could be done within the CGAL library in C++.

Towards an Intelligent Post-training Mutation Tool for Deep Learning Systems
Contact: Prof. Paolo Tonella
Co-Supervisor(s): Gunel Jahangirova, Nargiz Humbatova
Deep Learning (DL) has become an integral part of many groundbreaking projects and products we use everyday. As quality and safety remains the main concern for the developers and users of modern products based on Artificial Intelligence, different techniques aimed at assessing their quality are of increasing interest for research community. Mutation testing is one of such techniques: it deliberately seeds faults in the form of small syntactic changes into the program under test to create a set of faulty programs called mutants. The general principle underlying this approach is the assumption that faults used by mutation testing represent the mistakes that programmers usually make. Mutation testing aims to assess the quality of a given test suite in terms of its capability to detect faults.
In traditional software systems the decision logic is often implemented by software developers in the form of source code. In contrast, the behaviour of a DL system is mostly determined by the training data set and the training program, i.e. these are the two major sources of defects for DL systems. There currently exist two tools that are designed specifically for performing mutation testing for DL systems. However, one of the tools is a pre-training one, which means it injects the faults into system prior to the training and thus is

computationally expensive, while the second one, a post-training mutation tool, injects faults that are random and not very likely to happen in real world. Such faults usually introduce slight noise or modifications to a randomly selected subset of weights or change a structure of an already trained DL model by adding/deleting its layers or replacing the activation function.
The goal of the project is to develop a new post-training mutation tool that would solve the limitations set by the previous approaches, i.e. to introduce smarter and fast DL-specific mutation operators that produce stable and reliable results and would facilitate the increased interest for mutation testing in DL community.

Understanding End-user Attitudes towards Location Sharing Services
Contact: Prof. Marc Langheinrich
So-called “geosocial” applications allow one to share one's current location with friends, families, or even the public. In order to better understand why people are (or are not) using such applications in their daily life, we drew up a technology acceptance model (TAM) – a model that helps predicts what factors influence a person’s use or non-use of such services. The initial model was based both on prior work in location sharing, as well as on actual service use by 36 participants of a 4-week study we conducted. In a subsequent step, we now plan to verify this model in a much broader survey study.
We are thus seeking a highly motivated student to join us in planning, conducting, and analyzing an international survey study. A good understanding of location sharing technology, as well as basic statistics is required. You will learn how to design and administer survey research using crowdworking services such as Amazon Mechanical Turk or Prolific, as well as understanding the results with the help of analysis software such as SPSS and AtlasTI. Ultimately, the student will join our international research team to analyze and write up the results for submission to a high-profile journal.

Understanding Practices and Motivations for Sharing Physical Resources through Digital Services (SHA21.B)
Contact: Prof. Marc Langheinrich
Today, vast amounts of user-generated and user-mediated content populates social networks. Current research has focused extensively on needs, practices, and concerns surrounding the sharing of photos and videos, textual information (e.g., status updates), and documents. However, in recent years, the scope of what is "shareable" has greatly increased, comprising not only audio-visual content but also preferences and tastes (e.g., playlists, food), physiological data (e.g., workouts), trips, and even information about and access to real-world artifacts (e.g., "couchsurfing"). A recent market trend is to share personal physical possessions, initially rooms and apartments (e.g., Airbnb), but more recently rides (Uber), cars (Getaround) and household items (Snapgoods). The goal of this project is to design, conduct and subsequently analyze an online survey that attempts to elicit current practices of usage of selected sharing economy services, as well as identify motivations to participate in such services. Additionally, the student should conduct contextual interviews involving different stakeholders of such service (e.g. users, non-users, owners, suppliers) to further understand the economic role of using such sharing-economy services. This projects requires strong analytical skills and a willingness to learn about a novel and emerging research field. Experience with performing survey research and/or interviews a plus.

Scalable State Machine Replication
Contact: Prof. Fernando Pedone
State Machine Replication (SMR) is a well-established replication technique used by many production systems, including Apache Zookeeper, Google Chubby, Windows Azure storage, Google Spanner, and many others. Scalable State Machine Replication (S-SMR) is a recent extension of SMR developed at the distributed systems group at USI that promises unlimited

performance in addition to configurable fault tolerance. Some initial efforts, for example, resulted in a prototype that outperforms Zookeeper by almost an order of magnitude. This project will look into various aspects of S-SMR and contribute to cutting-edge research with high prospects of applicable results within a team of highly motivated and talented students.

Blockchain, blockchain, blockchain
Contact: Prof. Fernando Pedone
Blockchain has gained much traction in recent years. From a topic restricted to specialized circles, it has made it to the general press with daily headlines, including many scandals. Beyond the hype, blockchain fostered the development of sophisticated distributed algorithms and models. And many interesting issues remain unaddressed. In this project, the student will team up with a group of talented researchers from USI and a leading blockchain enterprise to help advance the state of the art in the field.

Towards Efficient Dataset Reduction to Reduce the Costs of Mutation Testing for Deep Learning Systems
Contact: Prof. Paolo Tonella
Co-Supervisor: Dr. Nargiz Humbatova
Deep Learning (DL) has become an integral part of many groundbreaking projects and products that we use every day. As quality and safety remain the main concerns for developers and users of modern AI-based products, various techniques for assessing their quality are of increasing interest to the research community. Mutation testing is one such technique, in which errors in the form of small syntactic changes are deliberately introduced into the program under test to create a set of faulty programs, called mutants. The general principle underlying this approach is the assumption that faults used by mutation testing represent the mistakes that programmers usually make. Mutation testing aims to assess the quality of a given test suite in terms of its capability to detect faults.

In traditional software systems the decision logic is often implemented by software developers in the form of source code. In contrast, the behaviour of a DL system is mostly determined by the training data set and the training program, i.e., these are the two major sources of defects for DL systems. There exists a mutation testing tool called DeepCrime [1], which is designed to perform mutation testing on DL systems and is based on real DLspecific faults. However, it injects faults into a DL system prior to the training following a realistic fault injection scenario and is therefore computationally expensive. There is another tool DeepMutation++[2], which is computationally cheap because it injects faults into an already trained model. Such changes are random and not very likely to occur in the real world. These mutations usually introduce small noise or changes to a randomly selected subset of weights or change the structure of an already trained DL model by adding/deleting its layers or replacing the activation function.

The aim of the project is to explore methods for effectively selecting subsets of data from the existing dataset to enable (1) smarter and faster DL-specific post-training mutations, and (2) reducing the cost of pre-training mutations of DL systems. In particular, in the first use case, the student will explore ways to group inputs based on some common features so that perturbing the behaviour of an already trained model can produce meaningful mutants. In the second application, the student will investigate different dataset reduction techniques, with the aim of reducing the training time by reducing the amount of training data fed to a model, while maintaining the same performance of the trained model or the same mutation test results. It is worth noting that efficient dataset selection/reduction can be useful for various tasks beyond mutation testing.

Nargiz Humbatova, Gunel Jahangirova, and Paolo Tonella. 2021. DeepCrime: mutation testing of deep learning systems based on real faults.In Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA ’21)

Hu, Qiang, Lei Ma, Xiaofei Xie, Bing Yu, Yang Liu, and Jianjun Zhao. Deepmutation++: A mutation testing framework for deep learning systems." In 2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE'19)

Enhancing Deep Learning Mutations via Weight-specific Fuzzing
Contact: Prof. Paolo Tonella
Co-supervisor: Dr. Jinhan Kim
Mutation testing in Deep Learning (DL) systems has been used to enhance the robustness and reliability of DL systems, by assessing their test sets using model variations, referred to as "mutants". One category of DL mutations is a post-training mutation, which perturbs the weights of trained models in order to subtly break the original model. This method provides a measure of the quality of a test set by identifying whether the test set can effectively discern the mutated model. The procedure of breaking the DL model employs several Mutation Operators (MOs) such as Gaussian Fuzzing (GF) -- an operator that targets a portion of weights to adjust their values based on a Gaussian distribution. However, this approach neglects the intrinsic properties of the weights. The impacts of MO should differ across the weights, neurons, layers, and model structure, taking into account local characteristics. Therefore, this project intends to develop a new weight-specific MO that accommodates the distributions of weights derived from multiple original models. The student will develop this new mutation strategy, which involves two main steps. The initial phase requires understanding the state-of-the-art post-training mutation technique, using its source code as a base to build the new weight-specific MO. The next step involves assessing the performance of this new MO across different dataset/models to determine the sensitivity and killability of its mutants.

Test Input Prioritization for Autonomous Driving Systems
Contact: Prof. Paolo Tonella
Co-Supervisor: Tahereh Zohdinasab
Autonomous driving systems have garnered significant attention in recent years, with rapid advancements in technology and the increasing demand for self-driving vehicles. Ensuring the safety and reliability of these systems is of paramount importance before their deployment on public roads. Thorough testing is crucial to identify potential failures and assess the system's behavior under various scenarios. However, exhaustive testing of autonomous driving systems is resource-intensive, time-consuming and difficult due to the huge test input space to cover and the high labeling cost. Consequently, prioritizing test inputs and labeling only 'high quality' ones, i.e., tests more likely to trigger a failure, is necessary for test cost reduction.

Multiple test prioritization techniques have been proposed by Software Engineering researchers. These techniques are mainly evaluated by measuring the number of failures exposed by the prioritized tests. Simply counting the number of failures produces a distorted view on the quality of the prioritized test inputs since they may be extremely similar (or even identical) inputs that expose the same problem of the deep learning component. Ideally, the prioritized test set should contain very diverse tests. Therefore, defining and measuring test uniqueness is extremely important for test prioritization when considering inputs for DL systems.

This master thesis proposal aims to address the challenge of test input prioritization for autonomous driving systems, with the goal of enhancing testing efficiency while preserving diversity. The proposed research seeks to develop a method that can effectively prioritize the test inputs, enabling the identification of critical scenarios that pose higher risks or require specific attention during testing and to use a novel metric that takes into account also test uniqueness to evaluate and compare the existing test prioritization techniques.

The research methodology will involve the following key steps:
Literature Review: Conduct a comprehensive review of existing literature on autonomous driving systems, and test input prioritization techniques. Identify the limitations and gaps in current approaches.
Test Input Prioritization Framework: Design and implement a prioritization framework that incorporates various factors, such as safety concerns and diversity. The framework should efficiently rank the test inputs based on their importance and potential impact on system performance.
Performance Metrics and Evaluation: Define appropriate performance metrics to evaluate the effectiveness of the proposed prioritization method. Compare the prioritized test inputs with state of the art approaches in terms of the number of diverse critical scenarios detected and testing efficiency.

The expected outcome of this research is a novel test input prioritization method that enhances test diversity for autonomous driving systems. The proposed framework will aid researchers and developers in allocating testing resources effectively by focusing on critical scenarios and reducing redundant testing efforts. Furthermore, this research contributes to the broader field of autonomous driving by providing insights into the challenges and opportunities in testing and validating these complex systems.

Fast updates of the Constraint Delaunay Triangulation and other tree Voronoi structures Contact: Prof. Evanthia Papadopoulou
The constrained Delaunay triangulation (CDT) is a variant of the well-known Delaunay triangulation in which specified edges, sometimes called segments, are constrained to appear. These constraints have many uses, such as representing boundaries of non-convex objects, supporting better interpolation of discontinuous functions, and aiding the enforcement of boundary conditions in finite element meshes [1]. The CDT of a point-set is as close to being a Delaunay triangulation as possible subject to those constraints. The project will develop randomized linear-time algorithms to update a constraint Delaunay triangulation when a new segment constraint is inserted. This will lead to a fast incremental construction of the constraint Delaunay triangulation. CDTs are used, among others, in Delaunay refinement methods for mesh generation. The project is a concrete example of a new approach to compute tree and forest Voronoi diagrams in linear expected time.

The proposed algorithms correspond to recent research results based on the following two papers.

Fast segment insertion and incremental construction of constrained Delaunay triangulations, Jonathan Richard Shewchuk and Brielin C. Brown, Computational Geometry: Theory and Applications, 2015

Abstract Voronoi-like Graphs: Extending Delaunay's Theorem and Applications, Evanthia Papadopoulou, SoCG 2022 https://drops.dagstuhl.de/storage/00lipics/lipics-vol258socg2023/LIPIcs.SoCG.2023.52/LIPIcs.SoCG.2023.52.pdf

Partially Homomorphic Encryption for Stream Processing Frameworks
Contact: Prof. Patrick Eugster
Co-Supervisor: Dr. Pavel Chuprikov
Due to the rapid spread of IoT, billions of devices are expected to continuously collect and process sensitive data. Because of limited computational and storage capacity available on IoT devices, the current de facto approach is to send the gathered data to the cloud for computation. Unfortunately, using public (untrusted) cloud infrastructures for processing continuous queries including on sensitive data leads to concerns over data confidentiality. An attractive approach to preserving the confidentiality of continuous query processing while utilizing public clouds is through the use of partially homomorphic encryption (PHE). PHE

allows computations over encrypted data, without revealing plaintext values. Recently, we developed a set of symmetric cryptosystems that retain the homomorphic expressiveness of previous asymmetric cryptosystems while being more performant (https://github.com/ssavvides/symmetria). The goal of this project is to apply these existing symmetric PHE schemes to allow confidentiality-preserving continuous queries, using the Apache Storm stream processing framework into which we previously already integrated several asymmetric PHE schemes. The student will take part in rewriting select streaming queries to use symmetric PHE, designing efficient serialization and deserialization strategies, and introducing ciphertext handling nodes into Storm's topologies. Microbenchmarking may be further used to guide the design of resource usage.

Cost-based Mechanism Selection for Secure Cloud Computing
Contact: Prof. Patrick Eugster
Co-Supervisor: Dr. Pavel Chuprikov
The value of what can be derived from customer data is being increasingly recognized by many industries, making information the new currency. With the amount of data generated growing exponentially, many organizations have turned to the cloud in their search for costeffective information processing. This leads to security concerns of processing sensitive data using third-party untrusted computational resources. While there are both software (homomorphic encryption) and hardware (secure enclaves) techniques with the potential to perform such processing without leaking information, they have their own constraints and overheads, so that there is no single universal solution. We have designed a system, called Hydra, that supports a multitude of security mechanisms nicely decoupling privacy policies from queries. To guarantee compliance with a chosen privacy policy, Hydra introduces a domain specific language (DSL). Hydra is integrated with the Apache Spark streaming processor to provide users with the familiar query abstraction. The current limitation of Hydra is that the choice of the security mechanism is hardcoded, while our DSL is able to check compliance of any such choice. The student will be working on query execution heuristics in Hydra (which uses standard Spark SQL extension points), based on empirical performance measurements for different security mechanisms combined with Spark execution metrics.

Language-based Policy Checking for Secure Computing
Contact: Prof. Patrick Eugster
Co-supervisor: Dr. Pavel Chuprikov
The value of what can be derived from customer data is being increasingly recognized by many industries, making information the new currency. With the amount of data generated growing exponentially, many organizations have turned to the cloud in their search for costeffective information processing. This leads to security concerns of processing sensitive data using third-party untrusted computational resources. While there are both software (homomorphic encryption) and hardware (secure enclaves) techniques with the potential to perform such processing without leaking information, they have their own constraints and overheads, so that there is no single universal solution. We have designed a system, called Hydra, that supports a multitude of security mechanisms nicely decoupling privacy policies from queries. A security policy checker in Hydra checks compliance of queries with the security policy of interest — in particular, ensuring absence of insecure information flows. Hydra is based on Apache Spark — one of the most active open source projects boasting contributions from over 1200 developers spread across 300 companies, making it a unique code base to learn and experience the length and breadth of system design principles. As part of the project, the student will integrate, enrich, and optimize the security policy checker with SparkSQL — Hydra’s query processing pipeline.

Compile-time Verification of Fault-tolerant Distributed Systems
Contact: Prof. Patrick Eugster
Co-Supervisor: Dr. Pavel Chuprikov

Software defects cost our IT-centered society exorbitant amounts of money. To make matters worse, driven by the advent of paradigms such as cloud computing and blockchains, software has been becoming increasingly distributed, i.e., its execution spans many processes. Besides having to avoid “conventional” intrinsic defects in the actual software, programmers now have to cater for partial failures, e.g., the possibility that certain processes or hosts fail while others continue to operate. Catering for these requires complex protocols, making implementation error-prone. Traditional “full depth” verification of programs involve lengthy verification processes requiring much manual effort and expert knowledge and are thus easily left out of the loop. The goal of this project is to apply and improve a prototype fault-tolerant event-driven programming framework that allows verification of component interaction as part of compilation. Our target are distributed middleware systems following earlier experiences applying it to the cluster manager component of the Apache Spark system. Along the way, extensions and practical additions to the respective domainspecific language may be investigated together with runtime optimizations, and performance evaluation conducted to demonstrate low overheads.

Portable Programmer-agnostic use of Trusted Hardware
Contact: Prof. Patrick Eugster
Co-Supervisor: Dr. Pavel Chuprikov
Malware and other attempts of tampering with computer software remain a dominant challenge to computer security. While several trusted execution environments (TEEs) allow programs to be shielded from attacks (e.g., Intel SGX, ARM Trustzone, AWS Nitro), leveraging these requires expert knowledge in security and the respective TEE, in addition to deep understanding of the corresponding programs. Even without considering the performance characteristics of different TEEs, TEE-based programs are not portable across TEEs of different vendors due to different APIs and functionalities proposed. We thus propose to use TEEs in combination with program anomaly detection (AD). By creating models of programs and comparing these against executions at runtime, AD can be applied without modifications to software. By tracing appropriate features, AD can detect various attacks with low overhead. However, being implemented fully in software, existing AD solutions have a fundamental flaw – their own mechanisms for tracing executions and comparing them to models are not protected from tampering. The goal of this project is thus to implement an AD monitor in a secure manner leveraging a TEE and other hardware features. This generic implementation of the monitor is independent of any monitored program thus supporting portability across TEEs of different vendors.

Unbounded Model Checking for TLA+
Contact: Prof. Patrick Eugster
Co-Supervisor: Dr. Rodrigo Otoni
Correctness of distributed systems is of critical importance. One good way to establish guarantees about the behavior of such systems is through the writing and reasoning of TLA+ specifications, with this approach having been used in practice for a number of years. One important limitation of existing automated reasoning support for TLA+ is that it is based on bounded techniques, weakening the guarantees given to users and potentially hiding vulnerabilities. The goal of this project is to address this problem via a novel unbounded model checking approach for TLA+, based on the solving of constrained Horn clauses (CHC). CHC-based automated reasoning is a growing field, in which successful results have been achieved already say in the context of blockchain verification, but whose potential has not yet been fully tapped. The goal of this project is to extend the open-source symbolic model checker Apalache to improve the scalability of bounded TLA+ verification. Developing a CHC-based reasoning engine for Apalache and exploiting recent advances in CHC solving will enable efficient unbounded verification of TLA+ specifications.

Formal Modeling of Probabilistic Quantum Network Policies

Contact: Prof. Patrick Eugster
Co-Supervisors: Dr. Anita Buckley, Dr. Pavel Chuprikov
Quantum computing, communication, and sensing technologies offer fundamentally new ways for information processing. The objective of quantum communication is to transmit quantum states, which may be entangled, causing stronger correlations. The no-cloning theorem (i.e., qubits cannot be copied) makes quantum communication inherently secure, leading to several novel applications. The distribution of entangled qubits (Bell pairs) between distant end-nodes will be the main task of the quantum Internet of the future, and the main challenge will be scaling. We are developing QNetKAT (Quantum NetKAT), a language and logic for dealing with and reasoning about quantum networks. QNetKAT has primitives for creating and transmitting Bell pairs, together with parallel and sequential composition operators, and offers a simple way for expressing quantum network policies. In the course of this project the student will get familiar with the components of quantum networks and protocols for long distance entanglement distribution. Decoherence, losses, and noise-errors cause stochastic behavior of quantum operations. The goal of this project is to develop the QNetKAT language with probabilistic semantics. The main tasks will consist in extending the language with new primitives for expressing probabilistic behaviors, and implementing these in the NetSquid quantum network simulation platform (https://netsquid.org/) using Python.

Automatic Feedback for PyTamaro Web Activities
Contact: Prof. Matthias Hauswirth
Co-Supervisor: Luca Chiodini
Over the last two years, the LuCE research group has been developing PyTamaro, an educational Python library to learn programming using graphics.
To ease the adoption of the library, the group has also developed an online platform, PyTamaro Web (https://pytamaro.si.usi.ch/), enabling learners to program with PyTamaro directly in their browser, without requiring any installation. The platform now contains more than 100 activities, which interleave explanations and code cells in which to write Python programs (in the style of a Jupyter notebook). The platform also contains a dozen of curricula that guide the learners through those activities.
PyTamaro Web is being actively developed and is used by hundreds of users per day, both in Swiss high schools and worldwide. It also hosts a curriculum part of the international Hour of Code initiative.

While working through the programming activities, learners can submit their code and see the output it produces, but so far they receive no indication about whether their solution (i.e., the produced graphic) is correct. The goal of this project is to extend the capabilities of PyTamaro Web to automatically check the correctness of the student's solution and automatically produce feedback, to help in the learning process.

Experience with Python and web development (React / TypeScript) is beneficial. If you are interested in programming languages and in research on teaching and assessing the understanding of programming, and eager to see your efforts having an immediate impact on thousands of students learning to program, contact us to learn more!

Better Documentation for Python Libraries
Contact: Prof. Matthias Hauswirth
Co-Supervisor: Luca Chiodini
Over the last two years, the LuCE research group has been developing PyTamaro, an educational Python library to learn programming using graphics.
To ease the adoption of the library, the group has also developed an online platform, PyTamaro Web (https://pytamaro.si.usi.ch/), enabling learners to program with PyTamaro directly in their browser, without requiring any installation. The platform now contains more

than 100 activities, which interleave explanations and code cells in which to write Python programs (in the style of a Jupyter notebook). The platform also contains a dozen of curricula that guide the learners through those activities.
PyTamaro Web is being actively developed and is used by hundreds of users per day, both in Swiss high schools and worldwide. It also hosts a curriculum part of the international Hour of Code initiative.

A key aspect in working with any library, PyTamaro included, is being able to browse its documentation to quickly retrieve the piece of information needed (e.g., “which parameters does that function take? what is its return type?”). There are a number of standard tools to produce documentation (e.g., Sphinx). But despite the popularity of Python as a programming language, the documentation resulting from these tools can be extremely hard to navigate for beginner learners (and sometimes for proficient programmers as well!). The goal of this project is to develop a better format for the documentation of a library that takes into account the learners’ needs at the various stages of the learning process. We would like to integrate this into the PyTamaro Web platform for documenting PyTamaro, with possible extensions to other core parts of the Python standard library.

Experience with Python and web development (React / TypeScript) is beneficial. If you are interested in programming languages and in research on teaching and assessing the understanding of programming, and eager to see your efforts having an immediate impact on thousands of students learning to program, contact us to learn more!

Deep Learning of Diversification Processes: Simulation and Inference Across Disciplines
Contact: Prof. Ernst Wit
Co-supervisor: Dr. Francisco Richter Mendoza
Species diversification models, such as the birth-death model, provide crucial insights into the mechanisms driving species evolution and adaptation. These models are instrumental in explaining the emergence and extinction rates of species over time. Understanding these models is key to developing effective simulation frameworks, which can then be used to train deep neural networks. The training of neural networks with data generated from these models enables advanced statistical inference, uncovering patterns and relationships that are not immediately apparent. The potential of applying such models and inference techniques extends beyond biological systems. In social sciences, for instance, similar principles can be applied to study sociocultural evolution, economic dynamics, and population studies. This cross-disciplinary application showcases the versatility of the simulation framework combined with deep learning. The project will leverage the insights gained from species diversification models and other relevant literature, including key references such as [RJH+21] and [LLV+23], to develop a robust framework. This framework will not only model biological diversification but also simulate various social science scenarios. Subsequently, deep learning models, particularly neural networks, will be trained on these simulations for detailed statistical analysis and predictive modeling. The overarching goal is to bridge the gap between theoretical models of diversification and practical, data-driven applications in both natural and social sciences. By combining simulation, deep learning, and statistical inference, this project aims to provide a novel approach to understanding and predicting complex dynamic systems. This broad-spanning application of the proposed framework holds the potential to revolutionize how we approach problem-solving and decision-making in diverse fields.

References:
[LLV+23] I. Lajaaiti, S. Lambert, J. Voznica, H. Morlon, and F. Hartig. A comparison of deep learning architectures for inferring parameters of diversification models from extant phylogenies. bioRxiv, 03 2023.

[RJH+21] F. Richter, T. Janzen, H. Hildenbrandt, E. C. Wit, and R. S. Etienne. Detecting phylodiversity-dependent diversification with a general phylogenetic inference framework. bioRxiv, 07 2021.

XAI-Fed: Explainable AI for Federated Models in Wearable Sensing
Contact: Prof. Marc Langheinrich
Co-supervisors: Dr. Martin Gjoreski, Daniil Kirilenko
Federated learning and its combination with differential privacy is the latest technique for building privacy-aware machine-learning models. Its primary assumption – no data leaves the local data storage, has enabled its application in a variety of privacy-sensitive domains: mobile keyboard prediction, human mobility modeling based on GPS data, modeling from electronic health records, etc. Artificial Intelligence (AI) methods can bring significant and sustainable improvements to our lives. However, end-users must be able to understand those systems. Unfortunately, today's groundbreaking AI methods are black-boxed (i.e., the decision model and the process are not understandable). The increased complexity of AI algorithms has made previous eXplainable AI (XAI) tools unsuitable, including the fact that most of the XAI solutions are not designed to operate under privacy constraints. This project will investigate XAI techniques compatible with privacy-aware approaches (e.g., federated learning). The focus will be on counterfactual explainers [55] for wearable sensing data. Specific project tasks are: (i) Analyze XAI tools that can operate under privacy constraints, focusing on counterfactuals; (ii) Pre-process one dataset from wearable sensing systems. Example datasets include emotion recognition, activity recognition and energy expenditure estimation; (iii) Develop machine learning models for one of the datasets in step 2, and apply existing XAI tools on the models developed, including the method for generating counterfactual explanations, BayCon; (iv) Develop XAI tool for counterfactual explanations that can operate under privacy constraints.

Causal Attention for Concept Embedding Models
Contact: Prof. Marc Langheinrich
Co-supervisor: Dr. Pietro Barbiero
Concept Bottleneck Models (CBMs) represent a significant advancement in interpretable machine learning, functioning by constraining predictions to pass through an intermediate layer of human-understandable concepts. However, these models often face a trade-off between accuracy and interpretability. To address this, Concept Embedding Models (CEMs) have been developed, which use high-dimensional representations of concepts to maintain interpretability while enhancing model accuracy. Despite their progress, a key limitation in CEMs is the absence of causal mechanisms in task predictors, which currently rely on simple linear layers or basic neural models without considering the causal relevance of concepts to specific task labels.
The primary goal of this project is to innovate within the field of CEMs by designing a causal attention mechanism for task predictors. This mechanism aims to selectively emphasize the most relevant concepts for a given task, embedding a layer of causal understanding into the decision-making process. The project seeks to blend the interpretability of CBMs with the accuracy of high-dimensional concept representations, ultimately creating a model that not only performs well but also aligns with human reasoning and intervention strategies. This endeavor stands to significantly enhance both the efficacy and transparency of CEMs in practical applications.

Out-Of-Distribution Generalization in Concept Bottleneck Models via Latent Active Learning
Contact: Prof. Marc Langheinrich
Co-supervisor: Dr. Pietro Barbiero
Concept Bottleneck Models (CBMs) have become a cornerstone in interpretable machine learning, primarily due to their structure that forces predictions through an interpretable layer

of concepts. This design significantly aids in both interpretability and targeted interventions. On the other hand, Active Learning, a paradigm where the model actively queries specific data points to label, optimizes learning efficiency, particularly in scenarios with limited labeled data. It becomes especially critical in handling Out-of-Distribution (OOD) data, which refers to data that differ significantly from the training distribution and often pose a challenge for models trained on specific datasets. Integrating active learning with generative CBMs can be particularly fruitful, as it enables the exploration and better understanding of the latent concept space, especially in OOD contexts.
The project aims to leverage active learning strategies to improve OOD generalization in CBMs. To this aim, the project seeks to navigate the latent concept space of generative CBMs, sample OOD embeddings from this space, and provide concept supervisions using an active approach. This approach will not only facilitate the model's adaptation to OOD scenarios but also enhance its overall interpretability and effectiveness in real-world applications where data distributions can significantly vary.

Interpretable Concept-based Semi-factuals
Contact: Prof. Marc Langheinrich
Co-supervisor: Dr. Pietro Barbiero
Concept Bottleneck Models (CBMs) have redefined the landscape of interpretable machine learning by ensuring that predictions are made through an explicit layer of human-understandable concepts. This bottleneck structure allows for direct interventions at the concept level, making CBMs invaluable for applications requiring transparency and explicability. An emerging area of interest in this domain is the generation of semi-factuals, which are plausible changes in concept labels which do not alter the downstream task prediction of the model. Semi-factuals serve as powerful tools in understanding model decisions by illustrating how changing some concept labels may not alter the final task prediction.
The goal of this project is to design and implement generative CBMs that can learn a latent concept space capable of modeling and generating concept-based semi-factuals. By doing so, we aim to enrich the interpretability of CBMs, providing users with meaningful insights about the boundaries and implications of potential interventions. The project focuses on enabling these models to sample from a distribution of semi-factuals at inference time, thereby offering a nuanced understanding of the model's decision-making process and the realistic scope of influencing these decisions.

Tabular deep concept reasoning
Contact: Prof. Marc Langheinrich
Co-supervisor: Dr. Pietro Barbiero
Tabular data, characterized by their structured format in rows and columns, are ubiquitous across various industries and scientific fields. Their applicability ranges from finance and healthcare to retail and beyond, making them one of the most common data types in machine learning applications. The ability to derive meaningful insights from tabular data is crucial, yet challenging, due to the complexity and diversity of the data. In this context, Deep Concept Reasoner (DCR) emerges as a novel neural-symbolic approach that stands out for its ability to automatically generate interpretable probabilistic programs from predicted concepts. This approach bridges the gap between deep learning's predictive power and the interpretability of symbolic reasoning.
The aim of this project is to conduct a comprehensive evaluation of the Deep Concept Reasoner (DCR) in comparison to existing state-of-the-art machine learning models specialized in handling tabular data. The focus will be on two primary metrics: accuracy and interpretability. By assessing DCR's performance in these areas, the project seeks to establish its efficacy and potential as a tool for both accurate prediction and meaningful interpretation in applications dealing with tabular data. This evaluation will contribute to the understanding of neural-symbolic models' roles in practical machine learning tasks, particularly in scenarios where both high accuracy and clear interpretability are essential.

Development of a Spectral-Temporal Transformer for Enhanced Biomedical Signal Analysis
Contact: Prof. Marc Langheinrich
Co-supervisors: Dario Fenoglio, Dr. Martin Gjoreski
Transformers have become the leading machine learning models in various fields due to their singular ability to remember long-term dependencies and identify meaningful correlations for predictions. Surpassing traditional models like convolutional neural networks (CNNs), Long Short-Term Memory networks (LSTMs), and Recurrent Neural Networks (RNNs), they excel in sequence-based tasks, including natural language processing and time-series analysis [2]. This superiority is largely attributed to the attention mechanism, which allows the network to discern dependencies of each sequence element in relation to others.
Despite these advancements, the application of Transformers in biomedical signal analysis, such as in electrocardiograms (ECG) and electroencephalograms (EEG), is still under-explored. These signals conceal considerable diagnostic information within their spectral domain, essential for accurate medical assessments. Current deep learning networks, like STResNet, utilize Fourier transformation to capture critical frequency domain information from these signals.
This project proposes to develop a Spectral-Temporal Transformer model, inspired by STResNet's frequency analysis and transformers' long-term memory capabilities. The model aims to integrate spectral and temporal features in biomedical signals using an attention mechanism.

Self-Supervised Federated Learning for Sensor Data in The Wild
Contact: Prof. Marc Langheinrich
Co-supervisors: Dario Fenoglio, Mohan Li
While sensors and deep learning have achieved remarkable success in fields like human activity recognition (HAR), the Internet of Vehicles, or healthcare, one of the biggest challenges still facing the community is the exploration of unrefined, unlabeled data. Billions of sensor data have been generated daily from the edge devices, but labeling them for training models is an extremely laborious and knowledge-demanding task. As a result, much of this data remains unused, even though it could potentially improve models. Additionally, privacy concerns restrict access to user data, further limiting the size of datasets available for data-hunger tasks.
This project focuses on two innovative approaches to address these issues: Self-supervised learning (SL) and the recent Federated Learning (FL). The SL community has seen rapid growth with the introduction of encoders, as unlabeled data has proven to be a valuable source for representation learning. FL, recently introduced by Google, offers a distributed learning framework that protects privacy by keeping user data locally while maintaining strong model performance. This project aims to integrate these two approaches into Self-Supervised Federated Learning (SSFL) specifically for sensor data, to enhance representation capabilities for downstream tasks such as HAR.

Human Activity Recognition with Multi-modality and Multi-frequency Federated Learning
Contact: Prof. Marc Langheinrich
Co-supervisors: Mohan Li, Dario Fenoglio
The evolution of Human Activity Recognition (HAR) technologies, especially through wearable sensors like head-worn devices, has opened new avenues for advanced personal health and activity monitoring. Traditional HAR systems, often relying on centralized machine learning, face significant privacy and data security challenges. Federated Learning (FL) offers a promising solution by enabling decentralized, privacy-aware machine learning across multiple devices. This approach ensures no data leaves the local device, thus

addressing privacy concerns. In federated environments, however, there's a need for multi-sensor models that can collaboratively train across diverse user devices, overcoming the inherent heterogeneity.
While the HAR community is rapidly expanding, the scope of datasets explored remains limited. Current research predominantly focuses on mobile devices or smartwatches, which datasets may not adequately distinguish specific head-related activities, such as eating or talking. Instead, head-worn devices present a more suitable alternative for such activities. Inspired by this observation, our project proposes to approach HAR with a multi-modality and multi-frequency strategy. The multi-frequency aspect allows the framework to adapt to various environments, including low battery scenarios, potentially reducing computing costs while maintaining high accuracy.

An integrated platform for workplace human sensing
Contact: Prof. Marc Langheinrich
Co-supervisors: Mohan Li, Pietro Barbiero
The relationship between employee productivity and job satisfaction has been an intriguing topic for decades. The benefits of maintaining a high efficiency at work are multi-folded: boost the self-confidence of employees which further promotes their productivity; keep a good work-life balance; create a satisfying work condition and relieve stress from working, etc. Feeling and being more productive are way more important than extending work time and exhausting yourself out, which is the trigger to mental illness such as depression, sleep disorder, even suicides under high pressure. However, quantitatively evaluating the productivity of workers is not an easy task. Even though sensors industry has been growing fast in recent decades, sensing productivity and satisfaction at work, despite in urgent need from each one of us, remains unsolved with a handy approach.
This project aims to address the relationship between productivity, job satisfaction, and well-being in the modern workplace, we initiate our first endeavor to build an integrated platform to collect sensor data among employees in workplace.

Personalization for stress and mood recognition in diverse user group
Contact: Prof. Silvia Santini
Co-supervisor: Lidia Alecci
Mood has a significant impact on how people behave, think, and act. Psychological and social science research highlights that physiological aspects, including mood and stress, vary
between individuals. Consequently, a universal stress prediction model is ineffective due to these variations, as what works seamlessly for one person may yield inaccurate or inconsistent results for another. However, creating individual models for each person lacks scalability. To address this, the student will utilize an existing large dataset and apply clustering techniques to identify similarities among users. The aim is to develop a machine learning model optimized for the specific characteristics and behavior patterns within these user groups, providing a balanced and scalable solution for personalized stress prediction across a diverse population.

Protect privacy in wearable devices using data anonymization
Contact: Prof. Silvia Santini
Co-supervisor: Lidia Alecci
Wearables facilitate continuous data collection to monitor diverse human behaviors, covering activity, health, and stress. This personal data, including electrocardiogram, movements, and heart rate, is often accessible online for research. Despite the common practice of masking names with random identifiers, studies show that this is insufficient for user identity protection due to the subject-dependent nature of physiological data. This research utilizes existing data and models to implement and assess anonymization techniques such as noise addition and synthetic data generation. The objective is to

determine the extent of effective user identity protection while minimizing disruption to human behavior prediction.

A Novel Approach Using Bilateral Data Fusion for EDA Data Classification
Contact: Prof. Silvia Santini
Co-supervisor: Leonardo Alchieri
This thesis addresses the impact of lateralization on Electrodermal Activity (EDA) sensors in wearable devices. Lateralization, influenced by brain hemisphere activation, affects the accuracy of EDA readings based on the device's placement on a specific body side. Despite recent studies highlighting this issue, there is limited exploration of the potential benefits of using EDA devices on both sides simultaneously. The research aims to fill this gap by investigating how leveraging data from both sides concurrently can enhance classifier accuracy through machine learning. The focus is on datasets in the lab, with implications for medical-grade applications affected by lateralization. Success in demonstrating improved accuracy may revolutionize the field, particularly in sensitive medical tasks, offering more reliable predictions for tasks impacted by lateralization.

Uncertainty-aware Deep Learning in digital healthcare
Contact: Prof. Silvia Santini
Co-supervisor: Leonardo Alchieri
In this thesis, the objective is to investigate the use of Monte Carlo Dropout, a Bayesian Deep Learning technique, to make uncertainty prediction on Neural Network outputs when dealing with physiological data. In particular, the student will investigate the creation of a system using Early Exit to adjust the computational power based on the uncertainty prediction.
The student will leverage state-of-the-art Deep Learning models to make predictions on publicly available datasets for health applications, e.g., ECG data from healthy and unhealthy
individuals.

Empathetic Virtual Agents using Large Language Models
Contact: Prof. Silvia Santini
Co-supervisor: Nouran Abdalazim
The rapid advancement of pervasive systems and wearables has paved the way for personal informatics systems, e.g., personal assistants and chatbots. These systems can be deployed in different personal and professional settings. Such personal assistants aim at personalizing the user 's experience by taking into consideration the user's affection status and respond based on the user's emotions and mental state status.
In this project, we aim at developing an empathetic personal assistant tool that provides an affection-aware user experience. The proposed tool integrates users' affection from wearables with large language models, e.g., ChatGPT.

Embodied Large Language Models for Personalized Meeting Summarization
Contact: Prof. Silvia Santini
Co-supervisor: Nouran Abdalazim
The rapid advancement of pervasive systems and wearables has paved the way for personal informatics systems, e.g., personal assistants and chatbots. Such systems gear towards enhancing users' productivity in both work setting and personal life. In work settings, users leverage personal informatics systems in managing their tasks, tracking their progress and scheduling their meetings.
Meetings are crucial for communication, decision-making, and brainstorming. The effectiveness of meetings relies on participants' ability to remember key topics, often using meeting minutes or notes as memory cues. However, this approach is time-consuming, demands high attention levels, and may lead to varied perceptions among participants,

causing a lack of synchronization.
In this project, we aim at developing a personalized meeting summarization tool that aims at optimizing the meeting experience. The proposed tool integrates users' affection from wearables with large language models, e.g., ChatGPT.

Scalable State Machine Replication
Contact: Prof.Fernando Pedone
State Machine Replication (SMR) is a well-established replication technique used by many production systems, including Apache Zookeeper, Google Chubby, Windows Azure storage, Google Spanner, and many others. Scalable State Machine Replication (S-SMR) is a recent extension of SMR developed at the distributed systems group at USI that promises unlimited performance in addition to configurable fault tolerance. Some initial efforts, for example, resulted in a prototype that outperforms Zookeeper by almost an order of magnitude. This project will look into various aspects of S-SMR and contribute to cutting-edge research with high prospects of applicable results within a team of highly motivated and talented students.

Blockchain, blockchain, blockchain
Contact: Prof.Fernando Pedone
Blockchain has gained much traction in recent years. From a topic restricted to specialized circles, it has made it to the general press with daily headlines, including many scandals. Beyond the hype, blockchain fostered the development of sophisticated distributed algorithms and models. And many interesting issues remain unaddressed. In this project, the student will team up with a group of talented researchers from USI and a leading blockchain enterprise to help advance the state of the art in the field.

Towards Higher Order Mutation Testing for Deep Learning Systems
Contact: Prof.Paolo Tonella
Co-Supervisor: Dr. Nargiz Humbatova
Deep Learning (DL) has become an integral part of many ground-breaking projects and products we use every day. As quality and safety remains the main concern for the developers and users of modern products based on DL, different techniques aimed at assessing their quality are of increasing interest for research community. Traditional mutation testing deliberately seeds faults in the form of small syntactic changes into the program under test to create a set of faulty programs called mutants. The general principle underlying this approach is the assumption that faults used by mutation testing represent the mistakes that programmers usually make. Mutation testing aims to assess the quality of a given test suite in terms of its capability to detect faults.

In traditional software systems the decision logic is often implemented by software developers in the form of source code. In contrast, the behaviour of a DL system is mostly determined by the training data set and the training program, i.e., these are the two major sources of defects for DL systems. There exists a mutation testing tool called DeepCrime [1], which is designed to perform mutation testing on DL systems and is based on real DL-specific faults. It injects faults into a DL system prior to the training, following a realistic fault injection scenario, and thus it is computationally expensive.

The goal of this project is to explore the behaviour of higher order DL mutants: the idea is to generate mutants by simultaneously injecting multiple faults into the DL system under test and to analyse whether first order mutants are redundant w.r.t. the more complex ones. If this were confirmed, the adoption of higher order mutants would potentially reduce the cost associated with mutation testing and would create novel, interesting patterns of faulty DL model behaviour.

References:

[1] Nargiz Humbatova, Gunel Jahangirova, and Paolo Tonella. 2021. DeepCrime: mutation testing of deep learning systems based on real faults.In Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA '21)

Delaunay-like Graphs and the Constraint Delaunay Triangulation
Contact: Prof.Evanthia Papadopoulou
The constrained Delaunay triangulation (CDT) in the plane is a variant of the renowned Delaunay triangulation in which specified edges, sometimes called segments, are constrained to appear. These constraints have many uses, such as representing boundaries of non-convex objects, supporting better interpolation of discontinuous functions, and aiding the enforcement of boundary conditions in finite element meshes [1]. The CDT is as close to being the Delaunay triangulation as possible subject to those constraints. The project will implement a simple incremental construction to compute the constraint Delaunay triangulation (CDT) of points in the plane, which is based on a new technique to update the CDT in linear expected time, after inserting a new segment. The focus of the project will be on developing the latter technique.

The approach makes use of a new concept, termed the Delaunay-like graph, which extends the concept of a Delaunay triangulation, and is interesting in its own right. Updating a CDT efficiently is a concrete, simple, yet important, example, which illustrates this new concept. The MaRS project is a first step towards a longer term research project on Delaunay-like and Voronoi-like graphs.
The following papers are closely related.

[1] Fast segment insertion and incremental construction of constrained Delaunay triangulations, Jonathan Richard Shewchuk and Brielin C. Brown, Computational Geometry: Theory and Applications, 2015

[2] Abstract Voronoi-like Graphs: Extending Delaunay's Theorem and Applications, Evanthia Papadopoulou, SoCG 2023
https://drops.dagstuhl.de/storage/00lipics/lipics-vol258-socg2023/LIPIcs.SoCG.2023.52/LIPIcs.SoCG.2023.52.pdf

Platform for Learning Software Engineering
Contact: Prof. Matthias Hauswirth
While there is an abundance of online platforms for learning to program, few platforms go beyond writing small individual pieces of code. The Luce group at USI developed the PyTamaro Web learning environment (https://pytamaro.si.usi.ch/), which is used to teach Python programming using a motivating but theoretically well founded approach in several schools across Switzerland. In this project the student will contribute to an extension of the PyTamaro Web platform, with the goal of providing a novel educational collaborative software engineering experience.

Automated Assessment for Graphics-Based Programming Activities
Contact: Prof.Matthias Hauswirth
Unit tests are a commonly used way to automatically assess student code. While most educational code can be tested with assertions on simple values, code that produces graphics is harder to test. The student in this project will develop innovative ways to test the correctness of graphics-producing functions. This will be conducted in the context of the PyTamaro (https://pytamaro.si.usi.ch/) educational graphics library for Python, developed at the Luce group at USI, and in use in different schools across Switzerland.

Interactive Educational Game-Development Environment
Contact: Prof. Matthias Hauswirth

Game development is an attractive domain for introductory programming courses. However, many traditional game development frameworks overwhelm the students with accidental complexity that distracts from the essential programming concepts students need to learn. PyTamaro (https://pytamaro.si.usi.ch/) is a pure, immutable library for Python for creating the graphics that are at the core of many games. Under the guidance of the Luce group at USI, the student in this project will extend PyTamaro to provide a similarly clean and simple functionality for developing interactive graphical games.

Effect of Reflection Prompts on Learning in an Online Platform
Contact: Prof.Matthias Hauswirth
Learners studying with books are prone to gloss over material instead of reading it deeply. If material is presented electronically, such as in e-books and online learning platforms, the linear sequence of material can be interspersed with reflection prompts, where learners are expected to think about a piece of material before moving on. Such prompts may consist of multiple choice questions learners need to answer correctly, or they may simply consist of a question learners are expected to think through, without any need to enter an answer, and without the need of the answer being correct. In the context of the PyTamaro Web platform (https://pytamaro.si.usi.ch/) developed by the Luce group at USI, and used in multiple schools across Switzerland, the student in this project will study the effect of different kinds of reflection prompts on learning gains in an online platform for learning to program.

Collector's Drive and Intrinsic Motivation in Programming Tutorials
Contact: Prof.Matthias Hauswirth
Humans like to collect artifacts, and they like to show off their collections. The PyTamaro Web learning platform (https://pytamaro.si.usi.ch/) developed by the Luce group at USI, and used in multiple schools across Switzerland, allows students to collect the relevant code they write as they work through learning activities, so they can reuse their code in subsequent learning activities. In this project you will implement and study novel ways to manage and present this collection, and their effect on student motivation. The project focuses on intrinsic motivation, that is on turning the act of curating and presenting their code collection into an educational experience of its own.

Enhancing Privacy and Transparency in Federated Learning: Counterfactual Explainability for Healthcare Applications
Contact: Prof.Marc Langheinrich
Co-supervisors: Dr.Dario Fenoglio, Gabriele Dominici
The rapid advancement of Artificial Intelligence (AI) has revolutionized personal healthcare [1]. However, the use of such technologies often involves handling sensitive, personal, and confidential data. Traditional Machine Learning (ML) systems generally rely on centralized learning, where user data is collected and processed on a single machine, raising serious privacy and data security concerns. Federated Learning (FL) offers a promising alternative, enabling decentralized, privacy-aware ML across multiple devices [2]. FL ensures that data remains on local devices, thereby mitigating the risks associated with data centralization and facilitating the integration of new data from distributed sources.

Despite these privacy-preserving benefits, it is crucial that end-users understand how such systems operate. Many state-of-the-art AI models remain "black boxes" with opaque decision-making processes. As AI algorithms grow more complex, traditional eXplainable AI (XAI) tools are increasingly inadequate, particularly in privacy-conscious settings. Furthermore, most existing XAI approaches are not designed to function under the privacy constraints required by FL.

This project aims to explore XAI methods that are compatible with privacy-preserving frameworks like FL, focusing on the behavior of counterfactual explainers in this context.

Counterfactual explanations answer “what if” and “why not” questions (e.g., “Why is this patient classified as high risk for developing a disease rather than low risk?”). These models predict changes in the input (e.g., a patient’s situation) that would lead to the desired outcome (e.g., lower risk), helping users better understand the model’s behavior and identify actionable steps to alter outcomes. However, counterfactual explanations may pose privacy risks, particularly if the model has overfitted and the generated counterfactuals resemble training data points. This project will investigate when these privacy leaks occur and propose strategies to enhance both the quality and privacy guarantees of counterfactual explanations in FL. The specific objectives of this project include: (1) Investigating XAI tools that operate within privacy constraints, with a focus on counterfactual methods [3]; (2) Preprocessing a healthcare dataset containing patient data that tracks their health over time, leading to a prediction of disease risk. (3) Developing ML models using existing XAI tools in the FL setting, with an emphasis on generating high-quality counterfactual explanations [4][5]; (4) Designing a new XAI tool for generating privacy-compliant counterfactual explanations.

Literature:
Rieke, Nicola, et al. "The future of digital health with federated learning." NPJ digital medicine 3, no. 1 (2020): 119.
McMahan, Brendan, et al. "Communication-efficient learning of deep networks from decentralized data." In Artificial intelligence and statistics, pp. 1273-1282. PMLR, 2017.
Ezzeddine, F. et al. (2024). Differential Privacy for Anomaly Detection: Analyzing the Trade-Off Between Privacy and Explainability. In: Longo, L., Lapuschkin, S., Seifert, C. (eds) Explainable Artificial Intelligence. xAI 2024. Communications in Computer and Information Science, vol 2155. Springer, Cham. https://doi.org/10.1007/978-3-031-63800-8_15
Guyomard, Victor et al. “VCNet: A self-explaining model for realistic counterfactual generation.” ECML/PKDD (2022).
Dominici, G., Barbiero, P., Giannini, F., Gjoreski, M., Marra, G., & Langheinrich, M. (2024). Counterfactual Concept Bottleneck Models.

Self-Supervised Learning for Smart Glasses Data
Contact: Prof.Marc Langheinrich
Co-supervisor:Francesco Bombassei De Bona, Dr. Martin Gjoreski
Affective computing is an interdisciplinary field that develops systems capable of recognizing, interpreting, and simulating human affect. A fundamental assumption is that different mental states (e.g., emotions and stress) manifest through physiological and behavioral changes. These changes can be captured through various wearable technologies, including smart glasses. Smart glasses offer a unique, unobtrusive platform for monitoring physiological signals like facial expressions [1].

Early affect-recognition systems relied on traditional machine learning (ML) methods paired with hand-crafted features. However, modern systems increasingly utilize deep learning, which can be further improved with techniques like self-supervised and unsupervised learning [2, 3]. These approaches have shown promise in domains like image-based ML and video analysis and are beginning to show potential with smart glasses in semi-controlled environments. The utility of smart glasses for affective computing in real-world settings is an active area of research.

This project will explore personalization and domain-adaptation techniques to address important challenges in wearable computing: noisy data, limited data, and domain shifts in the labels and the sensor data due to subjectivity. Existing processing pipelines and the deep learning architectures will be augmented with the latest unsupervised and/or self-supervised learning techniques. These advanced techniques should produce more robust and data-efficient models (i.e., requiring fewer person-specific labels). Domain adaptation

will be explored within datasets (e.g., from one user to another) and across datasets and devices.

Literature:
Kiprijanovska, I., Stankoski, S., Broulidakis, M. J., Archer, J., Fatoorechi, M., Gjoreski, M., ... & Gjoreski, H. (2023). Towards smart glasses for facial expression recognition using OMG and machine learning. Scientific Reports, 13(1), 16043.
Meegahapola, L., Hassoune, H., & Gatica-Perez, D. (2024). M3BAT: Unsupervised Domain Adaptation for Multimodal Mobile Sensing with Multi-Branch Adversarial Training. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 8(2), 1-30.
Meegahapola, L., Hassoune, H., & Gatica-Perez, D. (2024). M3BAT: Unsupervised Domain Adaptation for Multimodal Mobile Sensing with Multi-Branch Adversarial Training. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 8(2), 1-30.

Privacy-Preserving Counterfactual Explainability in Federated Learning for Employee Well-being
Contact: Prof.Marc Langheinrich
Co-supervisor: Mohan Li, Daniil Kirilenko
The modern workplace increasingly relies on technology to monitor and enhance employee well-being and productivity. Collecting data on workers' mental states can provide valuable insights to improve job satisfaction and performance. However, this practice raises significant privacy concerns, especially when sensitive personal and psychological data are involved. Traditional machine learning models often depend on centralized data collection and processing, which can expose confidential employee information and lead to physical and mental exhaustion due to perceived surveillance.

Federated Learning (FL) offers a promising solution by enabling decentralized, privacy-aware machine learning across multiple devices or locations. In FL, models are trained collaboratively without exchanging raw data, thereby mitigating the risks associated with data centralization. This approach is particularly beneficial in settings where data privacy is paramount, such as employee mental health monitoring.

Despite the privacy advantages of FL, it's crucial that employees understand how these AI models predict their states. Many advanced AI models function as "black boxes," with decision-making processes that are opaque to end-users. This lack of transparency can hinder trust and acceptance of AI tools in the workplace. Furthermore, most existing Explainable AI (XAI) methods are not designed to operate under the privacy constraints of FL environments.
Counterfactual explanations offer a way to provide actionable insights by answering "what if" and "why not" questions (e.g., "Why is my stress level predicted as high rather than moderate?"). These explanations suggest minimal changes in input features that could lead to a desired outcome, helping employees understand how to improve their well-being. However, generating counterfactual explanations in FL settings poses unique challenges. If not carefully designed, these explanations can inadvertently reveal sensitive information about other employees, thus compromising privacy.

This project aims to explore and develop XAI methods compatible with privacy-preserving frameworks like FL, focusing on generating actionable counterfactual explanations for models predicting employees' mental states.

Efficient rendering for VR and AR devices
Contact: Prof.Piotr Didyk

Motivation
Generating images for novel virtual and augmented reality headsets is computationally expensive due to the spatial resolution, frame rate, and image quality requirements imposed by the hardware and the human visual system. At the same time, when rendering for these wide-field-of-view devices, a viewer cannot appreciate everything. In particular, we can achieve significant computational gains without substantial quality loss by rendering lower quality for peripheral vision, the so-called foveated rendering approach. However, it is still an open research question of what information has to be rendered to guarantee image quality matching the requirements imposed by the hardware and viewer.

Goal
The main goal of this project is to push the boundaries of foveated rendering, which is considered a key enabler of future VR and AR displays. This project aims to develop new image-based enhancement techniques for foveated rendering. More precisely, we postulate that we can generate images of significantly lower quality (low spatial and temporal resolution, inaccurate depth, etc.) and then enhance them using simple image-processing or image-based methods to make them indistinguishable from full-quality rendering. The main focus is on real-time techniques which can replace high-quality rendering. Therefore, also the techniques have to be simple. However, we can take advantage of any byproducts of the rendering procedure, e.g., texture information, depth map, optical flow, or any pre-computed information. We envision the final results of the project to be demonstrated in the end-to-end system consisting of low-cost rendering, enhancement, and display on one of the latest VR and AR devices, such as Varjo VR-3 (https://varjo.com/products/vr-3/).

Prerequisites:
• Good programming skills
• Background in computer graphics
• Basic knowledge of image and video processing
• Experience with graphics pipelines and rendering API such as OpenGL

Related literature:
• Taimoor Tariq, Cara Tursun, and Piotr Didyk. 2022. Noise-based enhancement for foveated rendering. ACM Trans. Graph. 41, 4. https://doi.org/10.1145/3528223.3530101
• Taimoor Tariq and Piotr Didyk. 2024. Towards Motion Metamers for Foveated Rendering. ACM Trans. Graph. 43, 4. https://doi.org/10.1145/3658141

Simulations of Dynamic Social Networks
Contact: Prof.Ernst-Jan Camiel Wit
This project explores dynamic social networks as marked counting processes, where the evolution of interactions between network nodes is modelled over time. The hazard rate governing these interactions can be described using a Cox model, forming the basis of a relational event model. The objectives of the project include developing methods for the efficient simulation of inhomogeneous and non-linear dynamic networks, capturing their complex temporal dependencies and structural characteristics. Additionally, the project involves designing and implementing efficient inference techniques to estimate the parameters of these models, enabling accurate representation and prediction of dynamic relational processes.

Implementing causal relational event models

Contact: Prof.Ernst-Jan Camiel Wit
Relational event models provide an efficient framework for describing the effect of drivers on the dynamics of temporal networks, capturing how interactions evolve over time. Data used to fit these models have always been purely observational. Although sociologists have always been interested in testing hypothesis-driven effects, the observational experimental designs have limited direct causal interpretation of the results. This MARS project aims to advance the field by exploring explicit causal implementations of relational event models, leveraging the concept of invariance causal prediction.