DEPARTMENT OF
COMPUTER
SCIENCE

Proceedings of the Oxford Computer Science Conference 2023

13 June 2023

Conference committee
Maximilian Doré, Mathias Jackermeier, Junayed Naushad, Leto Luana Riebel, Tein van der Lugt

With outstanding support from
Sarah Retz-Jones and Jodie Mattioli

Contents

Schedule

09:45 – 10:00	*Coffee*	*Atrium*
10:00 – 10:15	**Introduction by Jon Barrett**	*LTB*
10:15 – 10:50	**Overview of research themes**	*LTB*
10:50 – 11:30	**Session I** Julian DCosta: *On Becoming and Staying Positive* Jack Liell-Cock: *Algebraic Representations of Graphs*	*LTB*
11:30 – 11:50	*Coffee break*	*Atrium*
11:50 – 12:25	**Session II** Muhammad Abuzar Ghafari: *Continuous & Autonomous Security Compliance | Automating the Regulatory & Industry Standard Implementation for SMBs* Edd Salkield: *Satellite Spoofing from A to Z: On the Requirements of Satellite Downlink Overshadowing Attacks*	*LTB*
12:30 – 14:00	**Poster and Demo Session over lunch** Jonas Beuchert: *SnapperGPS: A small, low-power, low-cost location data logger* Wael Albayaydh: *Mobile App to Support Privacy Protection in Smart Homes* Benjamin Hardin: *Can interfaces make self-driving cars psychologically safe?* Jacob Beck: *Hypernetworks in Meta-RL* Valentin Bacher: *Volumetric reconstruction from 2D ultrasound images using NeRFs.* Ambre Bertrand: *Deep Learning-Based Emulation of Human Cardiac Activation Sequences* Sarah Aldaweesh: *Exploring Well-being Mobile Apps for Arabic Speakers* James Coleman: *Myocardial ischaemia and repolarisation impairment as causes of arrhythmias and ECG abnormalities in hypertrophic cardiomyopathy* Nihil Shah: *Introduction to Spoiler-Duplicator Game Comonads*	*Atrium*
14:00 – 15:00	**OxWoCS Keynote by Polina Golland:** *Learning to Read an X-ray: applications to heart failure monitoring*	*LTB*
15:00 – 15:20	*Coffee break*	*Atrium*
15:20 – 15:50	**Lightning talks** Wei-Chen Lee: *k-Prize Weighted Voting Game* Oishi Deb: *Recognising Signs in Continuous Signing Sequences* Albert Dasí: *In-silico trials for patient-specific selection of atrial fibrillation treatment* Oliver Sourbut: *Cooperation and Alignment in Delegation Games* Zev Shirazi: *Codensity monads and related constructions* Aleksandar Petrov: *Tokenizers Introduce Unfairness Between Languages*	*LTA*
15:50 – 16:20	**Session III** Benedict Bunting: *Operational Algorithmic Game Semantics* Paulina Smolarova**:** *Sampling from Random Cluster Model on random regular graphs*	*LTA*
16:20 – 16:40	*Coffee break*	*Atrium*
16:40 – 17:45	**Session IV** Aleksandar Petrov: *A General Certification Theory with S-Lipschitzness* Andrea Kainz: *A Decision Support System for Fairness Metrics in Binary Classification* Yiyuan Yang: *DCdetector: Dual Attention Contrastive Representation Learning for Time Series Anomaly Detection* Nicola Dinsdale: *SFHarmony: Source Free Domain Adaptation for Distributed Neuroimaging Analysis*	*LTA*
From 18:30	**Award ceremony with wine and canapés**	*Ashmolean*

OxWoCS Keynote

This year's keynote was organized by the Oxford Women in Computer Science Society (OxWoCS).

Learning to Read an X-ray: applications to heart failure monitoring

Polina Golland

We propose and demonstrate a novel approach to training image classification models based on large collections of images with limited labels. We take advantage of the availability of radiology reports to construct joint multimodal embedding that serves as a basis for classification. We demonstrate the advantages of this approach in
application to assessment of pulmonary edema severity in congestive heart failure that motivated the development of the method.

Abstracts for Sessions I – IV

On Becoming and Staying Positive

Julian DCosta

I will describe our new decidability results on eventual non-negativity of matrices and checking whether a matrix subgroup contains a positive element. I will talk about the connections these problems have to classic theorems and conjectures in linear algebra and number theory.

Algebraic Representations of Graphs

Jack Liell-Cock

"The algebraic approach to modelling datatypes involves providing a set of construction primitives for injection into the type and higher-order compositions, along with equational laws to identify semantically equivalent constructions. These total and recursive definitions prevent usage errors from partial APIs and promote recursive computational patterns on the objects. We present an algebra for edge-indexed graphs using this approach. This first requires introducing a novel representation of multigraphs which puts precedence on the edge indices. The construction primitives and equational relations are then defined with respect to this structure, which gives rise to the algebra. The algebra translates directly into a Haskell implementation, which we further extend with functions for manipulating and transforming the edge graphs. We also introduce some homomorphisms in the graph algebra which compute the transpose, underlying set and shortest path algorithms. Finally, we present the beginnings of a generalisation of the algebra to hypergraphs which subsumes some node graph algebra constructions in literature."

Continuous & Autonomous Security Compliance | Automating the Regulatory & Industry Standard Implementation for SMBs

Muhammad Abuzar Ghafari

I wish to speak about(and present) a system i am developing which aims to help SMB to kickstart their Cybersecurity journey with minimum efforts to achieve their Security posture and compliance goals with continuous compliance and security automation in Public and Private Cloud Infrastructure.

Satellite Spoofing from A to Z: On the Requirements of Satellite Downlink Overshadowing Attacks

Edd Salkield

Spoofing attacks against satellite communications such as GPS have become increasingly common, as attackers increasingly have access to software-defined radio hardware. However,

existing research on GPS spoofing doesn't generalise well to other satellite systems, such as high data rate backhauls or satellite-to-customer connections, where the spoofing requirements are currently unknown.

We present the first systematic review of spoofing attacks in the general case, linking attack feasibility and impact to required budget through real-world experiments and channel simulations. We show that nearly all evaluated satellite systems were spoofable, even in the presence of a legitimate transmission, at a distance of 1km in the worst case and for a budget of ~2000 USD or less.

Operational Algorithmic Game Semantics

Benedict Bunting

The abstract from the paper is as follows:

We consider a simply-typed call-by-push-value calculus with state, and provide a fully abstract trace model via a labelled transition system (LTS) in the spirit of operational game semantics. By examining the shape of configurations and performing a series of natural optimisation steps based on name recycling, we identify a fragment for which the LTS can be recast as a deterministic visibly pushdown automaton. This implies decidability of contextual equivalence for the fragment identified and solvability in exponential time for terms in canonical form. We also identify a fragment for which these automata are finitestate machines. Further, we use the trace model to prove that translations of prototypical call-by-name (IA) and call-by-value (RML) languages into our call-by-push-value language are fully abstract. This allows our decidability results to be seen as subsuming several results from the literature for IA and RML. We regard our operational approach as a simpler and more intuitive way of deriving such results. The techniques we rely on draw upon simple intuitions from operational semantics and the resultant automata retain operational style, capturing the dynamics of the underlying language.

Sampling from Random Cluster Model on random regular graphs

Paulina Smolarova

The random cluster model is a random graph model originating from statistical physics, representing among others, percolation, Ising and Potts models. Sampling from this model is computationally hard in general, however, there are known classes of graphs and parameter ranges where we can find fast sampling algorithms. In this talk, I will present our new findings for sampling from the random cluster model on random regular graphs using Markov chains.

A General Certification Theory with S-Lipschitzness

Aleksandar Petrov

A Decision Support System for Fairness Metrics in Binary Classification

Andrea Kainz

In today’s world, machine learning systems are used in many sensitive areas of our lives such as credit scoring, disease prediction and job recruiting. In recent years, the concept of machine learning fairness has gained increasing attention as some prominent examples have demonstrated bias in computer systems. Within binary classification algorithms, there exist several definitions of the fairness of a machine learning model; however, it is impossible to fulfil all of them at the same time, and some of these metrics are interrelated or even incompatible in non-trivial scenarios. Publicly available fairness toolkits do not sufficiently support their users in selecting suitable fairness metrics. Therefore, the aim of my project was to create a user-friendly, interactive fairness tool to visualise pre-processing fairness metrics on the training data and post-processing fairness on a trained machine learning model with an integrated

decision support system that helps the user select appropriate post-processing fairness metrics according to the use case.

DCdetector: Dual Attention Contrastive Representation Learning for Time Series Anomaly Detection

Yiyuan Yang

Time series anomaly detection is critical for a wide range of applications. It aims to identify deviant samples from the normal sample distribution in time series. The most fundamental challenge for this task is to learn a representation map that enables effective discrimination of anomalies. Reconstruction-based methods still dominate, but the representation learning with anomalies might hurt the performance with its large abnormal loss. On the other hand, contrastive learning aims to find a representation that can clearly distinguish any instance from the others, which can bring a more natural and promising representation for time series anomaly detection. In this paper, we propose DCdetector, a multi-scale dual attention contrastive representation learning model. DCdetector utilizes a novel dual attention asymmetric design to create the permutated environment and pure contrastive loss to guide the learning process, thus learning a permutation invariant representation with superior discrimination abilities. Extensive experiments show that DCdetector achieves state-of-the-art results on multiple time series anomaly detection benchmark datasets.

SFHarmony: Source Free Domain Adaptation for Distributed Neuroimaging Analysis

Nicola Dinsdale

To represent the biological variability of clinical neuroimaging populations, it is vital to be able to combine data across scanners and studies. However, different MRI scanners produce images with different characteristics, resulting in a domain shift known as the `harmonisation problem'. Additionally, neuroimaging data is inherently personal in nature, leading to data privacy concerns when sharing the data. To overcome these barriers, we propose an Unsupervised Source-Free Domain Adaptation (SFDA) method, SFHarmony. Through modelling the imaging features as a Gaussian Mixture Model and minimising an adapted Bhattacharyya distance between the source and target features, we can create a model that performs well for the target data whilst having a shared feature representation across the data domains, without needing access to the source data for adaptation or target labels. We demonstrate the performance of our method on simulated and real domain shifts, showing that the approach is applicable to classification, segmentation and regression tasks, requiring no changes to the algorithm. Our method outperforms existing SFDA approaches across a range of realistic data scenarios, demonstrating the potential utility of our approach for MRI harmonisation and general SFDA problems.

Abstracts for Poster and Demo Session

SnapperGPS: A small, low-power, low-cost location data logger

Jonas Beuchert

Biologists and wildlife conservationists use global navigation satellite systems (GNSS), e.g., the GPS, to track animals, study their behaviour, and make informed conservation decisions. Existing tracking devices are often expensive ($100-$10,000) and may require bulky batteries for long-term deployments. This prohibits studies with many animals and is our motivation for developing SnapperGPS. SnapperGPS aims to be a cheap, small, and low-power tracking solution. Its core idea is to make the hardware as simple and as energy-efficient as possible. We achieve this by doing as little signal acquisition and processing on the device as possible. Instead, SnapperGPS provides a web service that processes the signals in the cloud. This allows us to build a bare-bones receiver for <$30 that runs for >1 year. It employs the snapshot GNSS technology, which is particularly useful to track aquatic animals that surface only briefly, e.g., sea turtles.

Mobile App to Support Privacy Protection in Smart Homes

Wael Albayaydh

Smart home devices are collecting different types of data about users inside smart homes, which is raising privacy risks and concerns. We are working to design an app to support users protect their privacy while they are inside smart homes.

Can interfaces make self-driving cars psychologically safe?

Benjamin Hardin

As humans increasingly interact with AI and robotics and sometimes even engage in conversation with them, the question emerges of mental safety to the human. How does the user build a mental model of the system's performance? What are the consequences if a system manipulates or misleads the user? How can these harms be prevented?

The talk will introduce the idea of psychological safety in the context of human-vehicle interaction, how psychological safety differs from trust, and what techniques are currently proposed to create psychological safe vehicles.

Hypernetworks in Meta-RL

Jacob Beck

Deep reinforcement learning (RL) is notoriously impractical to deploy due to sample inefficiency. Meta-RL directly addresses this sample inefficiency by learning to perform few-shot learning when a distribution of related tasks is available for meta-training. While many specialized meta-RL methods have been proposed, recent work suggests that end-to-end learning in conjunction with an off-the-shelf sequential model, such as a recurrent network, is a surprisingly strong baseline. However, such claims have been controversial due to limited supporting evidence, particularly in the face of prior work establishing precisely the opposite. In this paper, we conduct an extensive empirical investigation and suggest a method that works without the need for additional tuning. While we likewise find that a recurrent network can achieve strong performance, we demonstrate that the use of hypernetworks is crucial to maximizing their potential. We also show that hypernetwork initialization is a critical factor in meta-RL, and that naive initializations yield poor performance. Surprisingly, when combined with properly initialized hypernetworks, the recurrent baselines that are far simpler than existing specialized methods actually achieve not just strong performance, but are state-of-the-art (SOTA) on standard meta-RL benchmarks.

Volumetric reconstruction from 2D ultrasound images using NeRFs.

Valentin Bacher

NeRFs and Kilo-NeRFs became increasingly popular in recent years. We are applying this technique to reconstruct volumetric data from 2D ultrasound images. To reduce the number of parameters and make the approach more explainable, we are using a physics informed rendering technique simulating the ultrasound acquisition.

Deep Learning-Based Emulation of Human Cardiac Activation Sequences

Ambre Bertrand

We have developed a machine learning approach to emulate algorithmic methods currently used in the context of cardiac electrophysiological simulations. Having learnt from a large number of cardiac activations sequence maps, our model is able to estimate the activation pattern across the heart's ventricles, given a starting location and conduction velocity. Our model is 500 times faster

than the original simulation algorithm at point of deployment, and reproduces results within 4.7ms RMSE (approx 5% accuracy).

Exploring Well-being Mobile Apps for Arabic Speakers

Sarah Aldaweesh

Myocardial ischaemia and repolarisation impairment as causes of arrhythmias and ECG abnormalities in hypertrophic cardiomyopathy

James Coleman

Hypertrophic cardiomyopathy (HCM) is a common inherited heart disease and a significant contributor to sudden cardiac deaths in the young. Acute myocardial ischaemia (a sudden shortage of blood supplied to the heart muscle) and chronic repolarisation impairment (where cardiac muscle cells are diseased such that cellular repolarisation is slowed) are hypothesised to cause lethal arrhythmias in HCM, but the mechanisms by which these disease factors interact are unknown.

Computational models of human cardiac electrophysiology were used to simulate acute myocardial ischaemia in tissue and in HCM ventricles, where ventricular simulations were informed by common distributions of ischaemia as analysed in perfusion CMR scans (N=28 patients). Using S1-S2 pacing protocols, arrhythmic risk was measured for scenarios in which the hypertrophied ventricular septum was affected by (i) ischaemia only, and (ii) ischaemia and repolarisation impairment. ECGs were derived from ventricular simulations and compared between scenario (ii) and a further scenario affected by (iii) repolarisation impairment only.

Simulations identified that repolarisation impairment can enable arrhythmias to be induced at a lesser severity of ischaemia than otherwise possible for ventricles with normal baseline repolarisation, and that arrhythmias can be increasingly sustained for ischaemia affecting the full width of the ventricular wall. Moreover, simulations identified that due to an abnormal electrophysiological response to ischaemia of myocardium affected by repolarisation impairment, ischaemia in HCM can manifest on the ECG as uprighting of the T wave (i.e. the ECG can transiently appear more normal during episodes of ischaemia).

Overall, our findings suggest that deleterious interactions between ischaemia and repolarisation impairment contribute to arrhythmic risk in HCM, and that consideration of transient uprighting of the T wave on exercise ECGs could improve the diagnosis of ischaemia in HCM.

Introduction to Spoiler-Duplicator Game Comonads

Nihil Shah

I will present on a recent research programme drawing connections between categorical semantics and finite model theory. This connection uses comonads (a central object in categorical semantics) to internalise model comparison games (a tool in finite model theory to show inexpressiblity in a logic).

Abstracts for Lightning talks

k-Prize Weighted Voting Game

Wei-Chen Lee

We introduce a natural variant of weighted voting games, which we refer to as k-Prize Weighted Voting Games. Such games consist of n players with weights, and k prizes, of possibly differing values. The players form coalitions, and the i-th largest coalition (by the sum of weights of its members) wins the i-th largest prize, which is then shared among its members.

Recognising Signs in Continuous Signing Sequences

Oishi Deb, Andrew Zisserman

This work focuses on efficiently and accurately recognizing signs in continuous signing videos of British Sign Language. This is a very challenging task due to the speed of the signing, the large vocabulary, the within-class variation of a sign, and co-articulation where the sign is affected by those preceding and following it.
We investigate whether using human pose as an intermediate representation alleviates some of the problems of human appearance variability, and also leads to efficient training and inference. To begin, the project leverages access to large datasets of videos featuring British Sign Language, necessitating the extraction of key points from these videos frame-by-frame. By training a transformer with the extracted key points and their corresponding labels, we seek to develop a robust and accurate system for sign recognition. The new architecture, with its ability to capture long-range dependencies and effectively model spatial relationships, holds promise in capturing the intricate movements and gestures involved in the signing. By employing this model architecture and training it on the extracted keypoints, the project aims to overcome the challenges posed by the dynamic nature of signing-sequences and the diverse variations within signs.

In-silico trials for patient-specific selection of atrial fibrillation treatment

Albert Dasí

Atrial fibrillation (AF) is the most common arrhythmia, affecting 1 in 45 people in the UK. After first procedure treatment, 60% of patients will newly present in the hospital with recurrent AF.

During the presentation I will provide an overview of a computational framework that we have developed to identify optimal therapies for AF patients. The framework is based on in silico clinical trials, a concept that refers to the use of human-based modelling and simulation to virtually test medical treatments. Widely accepted in engineering, in silico trials are now emerging in medicine with a strong regulatory backing. Here I will cover the technical challenges involved in (1) constructing large populations of virtual patients that present variability in important clinical features; (2) conducting in-silico trials to evaluate state-of-the-art therapies for AF; (3) developing a decision algorithm that, based on virtual patient features, selects an optimal treatment.

The validation and adaptation of the new decision algorithm in clinical practice is expected to boost the efficacy of AF treatment.

Cooperation and Alignment in Delegation Games

Oliver Sourbut

In multi-principal-multi-agent 'delegation games', principals' welfare is a function of alignment, coordination ability, and 'welfare calibration'

Codensity monads and related constructions

Zev Shirazi

Tokenizers Introduce Unfairness Between Languages

Aleksandar Petrov

Recent language models have shown impressive multilingual performance, even when not explicitly trained for it. Despite this, there are concerns about the quality of their outputs across different languages. In this paper, we show how disparity in the treatment of different languages

arises at the tokenization stage, well before a model is even invoked. The same text translated into different languages can have drastically different tokenization lengths, with differences up to 15 times in some cases. These disparities persist across the 17 tokenizers we evaluate, even if they are intentionally trained for multilingual support. Character-level and byte-level models also exhibit over 4 times the difference in the encoding length for some language pairs. This induces unfair treatment for some language communities in regard to the cost of accessing commercial language services, the processing time and latency, as well as the amount of content that can be provided as context to the models. Therefore, we make the case that we should train future language models using multilingually fair subword tokenizers.

Awards

The conference committee awarded the following prices:

Price for best presentation

Nicola Dinsdale for their outstanding talk on *SFHarmony: Source Free Domain Adaptation for Distributed Neuroimaging Analysis*

Price for best poster and demo

Jonas Beuchert for their outstanding poster and demo of *SnapperGPS: A small, low-power, low-cost location data logger*

Price for best lightning talk

Oliver Sourbut for their outstanding talk on *Cooperation and Alignment in Delegation Games*

We thank the Computer Science department for the generous support of the conference.