Proceedings of the Oxford Computer Science Conference 2024

General Chairs: Benedict Bunting and Mathias Jackermeier
Conference Committee: Hunar Batra, Junayed Naushad, Siyi Sun, and Theo Wang
With assistance from: Sarah Retz-Jones and Jodie Mattioli

11th June 2024

Contents

Programme

09:45	Registration and Coffee		
10:00	Introduction and Deputy Head of Department (Research)		
10:30	Session I		The Complexity of Causal Discovery
10:50		Xingjian Bai	Sorting with Predictions
11:10	Coffee		
11:35	Lightning I	Jiaojiao Ye	PQD: Post-Training Quantization For Diffusion Models
		Fan Yang	AutoVCL: A Self-balancing Algorithm for Variational Continual Learning
		Siyi Sun	Missing Data Imputation
		Benjamin Rodatz	An elementary theory of machine learning tasks
12:00	Session II		Network communication and equilibrium selection
12:20		Lia Yeh	Researching and Teaching Quantum Picturalism
12:40	Lunch and Poster Session	Can Koz	Multiview Diffusion Models for Deformable Object Generation
		Matthew Morris	Orbit-Equivariant Graph Neural Networks
		Antonio Norelli	ASIF: Coupled Data Turns Unimodal Models to Multimodal Without Training
			Fast & Fair: A Collaborative Platform for Fair Division Applications
		Jingjie Yang	Orbit-finite-dimensional vector spaces
		Xiaxia Wang	Faithful Rule Extraction for Differentiable Rule Learning Models
			Understanding Challenges and Design Opportunities for Digital Mental Well-Being Support in Saudi Arabia
14:00	OxWoCS Keynote Speaker	Philippa Gardner	
15:00	Coffee		
15:20	Session III	Mathias Jackermeier	Generalising Complex Long-Horizon LTL Instructions via Deep Reinforcement Learning
15:40		Mihaela Stoian	Neuro-Symbolic AI for Synthesising Realistic Data
16:00			NIT: Neural Implicit Trajectories using Diffusion-Guided Scores
16:20	Brief Break		
16:25	Lightning II	Oishi Deb	Diffusion Models for 3D Reconstruction and Articulation of Deformable Objects
			Hairy Ground Truth Enhancement for Semantic Segmentation
		Alastair McCullough	Enterprise Information Governance, Domain Specificity, and the Data Mesh Case
		Jesse Wright	Conversational AI at a Web scale
16:50	Session IV	Tyler Reinmund	Transitioning towards a proactive practice: A longitudinal field study on the implementation of a ML system in adult social care
17:10		Edward Salkield	Security analysis of cutting-edge adaptive protocols in space systems
17:30	Drinks Reception and Prizes		
18:50	Arrival at St Anne's College for Dinner at 7pm		

OxWoCS Keynote Speaker

This year's Keynote Speaker was kindly organised in conjuction with Oxford Women in Computer Science.

Verified Software Specification at Scale

Philippa Gardner

Abstract

A long-term dream for myself and others, including many pioneers from Oxford, has been to use machine-supported reasoning to establish scientific and mathematical method to the specification and analysis of large-scale modern open software systems. This presentation will provide a whirlwind tour of some of my group's recent work on the specification and validation of language standards, libraries, and general programs, emphasising what makes a specification suitable for task in hand, rigorously evaluated and practical for real-world applications. It will cover:

Mechanised Language Standards. This includes the WasmCert mechanised specification project, which is actively used to correct and prove properties about the formal W3C WebAssembly (Wasm) standard, and the Wasm-SpecTec project, which offers a new approach to language specification and was provisionally accepted for adoption by the W3C Wasm committee in June 2024.

Compositional Analysis Techniques and Tools. This describes the foundations of compositional symbolic execution, a technique which underpins many tools for verification and true bug-finding in large-scale open industrial software, such as Meta's Infer platform, with compositionality achieved by using function specifications written in a modern Hoare logic called separation logic.

Program Logics for Concurrent Programs. This surveys recent work on verification techniques for complex shared-memory concurrent programs using two-sided abstract specifications of concurrent operations: implementations are proved correct with respect to an independent specification; and clients can be verified using this independent specification. This work has led to the development of the Iris framework, used by hundreds of specialists and underpinning the industrial Bedrock hypervisor.

Despite significant progress, we still have a long way to establish a scientific, mathematical method for the validation of modern software systems, akin to standard engineering practices. A recent publication by the Office of the White House calls for action from the technical community: "By reducing memory safety vulnerabilities at scale, creators of software and hardware can better secure the building blocks of cyberspace. This report focuses on programming languages as primary building blocks and explores hardware architecture and formal methods as complementary approaches to achieve similar outcomes (February 2024)". Our field just might be mature enough to begin to realise our dream.

Our publications can be found at `https://www.doc.ic.ac.uk/~pg/publications/`

Biography

Philippa Gardner is a Professor in the Department of Computing at Imperial College London and a Fellow of the Royal Academy of Engineering. She has developed formal methods and well-engineered tools that have led to mechanised international language standards and machine-checked correctness guarantees for industrial-scale open software. She completed her PhD thesis, supervised by Professor Gordon Plotkin FRS at Edinburgh in 1992, had fellowship positions in Edinburgh and then Cambridge hosted by Robin Milner, and then moved to Imperial as a permanent academic in 2001. Her work has been funded by twenty years of fellowships, including a UK Research and Innovation Established Fellowship from 2018–2023, two EPSRC programme grants, and gift funding from Meta, AWS and Microsoft.

Gardner was the director of the Research Institute on Verified Trustworthy Software Systems (VeTSS), funded by EPSRC and NCSC, 2017-2023. She was General Chair for POPL in January 2024, the 51st Annual ACM SIGPLAN Symposium on Principles of Programming Languages, and is an organiser of the Isaac Newton Institute (INI) programme on 'Big Specification' in October 2024, having previously organised the INI programme on 'Verified Software' in 2022.

Abstracts of Talks

Session I

The Complexity of Causal Discovery

I will present an overview of complexity results on one of the most general scientific inference problems - figuring out the causal relationships between a bunch of observed variables.

Sorting with Predictions

Xingjian Bai

We explore the fundamental problem of sorting through the lens of learning-augmented algorithms, where algorithms can leverage possibly erroneous predictions to improve their efficiency. We consider two different settings: In the first setting, each item is provided a prediction of its position in the sorted list. In the second setting, we assume there is a "quick-and-dirty" way of comparing items, in addition to slow-and-exact comparisons. For both settings, we design new and simple algorithms using only $O(\Sigma_i \log \eta_i)$ exact comparisons, where η_i is a suitably defined prediction error for the ith element. In particular, as the quality of predictions deteriorates, the number of comparisons degrades smoothly from $O(n)$ to $O(n \log n)$. We prove that the comparison complexity is theoretically optimal with respect to the examined error measures. An experimental evaluation against existing adaptive and non-adaptive sorting algorithms demonstrates the potential of applying learning-augmented algorithms in sorting tasks.

Session II

Network Communication and Equilibrium Selection

Coordinating the behaviour of self-interested agents, for example in the presence of multiple Nash equilibria, is a major research challenge for multi-agent systems. We introduce an extension of the strategic-form game which allows players to communicate with each other in accordance with a directed network structure, which we term a networked game. We assume that communication on the network is costless, potentially unbounded, takes place prior to game play, and may be

binding or non-binding. For the known setup of minority games (a class of congestion games), we demonstrate that, under certain regularity conditions, for each class of Nash equilibria, there exists a network that induces only strategies within such class. This approach provides two perspectives in the context of multi-agent systems. Descriptively, it offers an explanation for why specific outcomes that require correlated actions between some subsets of players (i.e., there exists some latent communication network between players) are observed; Prescriptively, it offers a mechanism for selecting for particular class of equilibria when there exists a multitude of equilibria, without resorting to repeated interaction or a trusted third-party signaller (i.e., correlated equilibrium).

Researching and Teaching Quantum Picturalism

Lia Yeh

Quantum theory is often considered an advanced undergraduate or graduate-level subject with many mathematical prerequisites, particularly linear algebra and symbolic manipulation within the Hilbert space formalism. This talk is on 1) ongoing research challenging this view through using quantum graphical calculi for quantum high school level education, and 2) how these calculi are being applied in academia and industry research across quantum computing. We posit that transitioning from symbolic presentations to pictorial ones will increase the appeal of quantum education and research, attracting more diverse audience.

Session III

Generalising Complex Long-Horizon LTL Instructions via Deep Reinforcement Learning

Mathias Jackermeier

We propose a novel approach to learning goal-conditional policies capable of satisfying multiple tasks specified in linear temporal logic (LTL). Unlike previous myopic methods that decompose tasks into subtasks, our approach considers the whole task, thereby avoiding suboptimal performance. We extend our method to handle full LTL, including infinite-horizon tasks, a capability not explored in prior works. Our focus lies on complex control problems with continuous state and action spaces, offering solutions to challenges posed by sparsity of rewards in such environments. By integrating these innovations, our approach not only achieves efficient task satisfaction but also demonstrates zero-shot generalization to novel tasks, a critical feature for real-world applications.

Neuro-Symbolic AI for Synthesising Realistic Data

Mihaela Stoian

Deep learning models have shown their strengths in various application domains, however, they often struggle to meet requirements for their outputs. Here I will discuss how deep generative models for tabular data can be constrained such that their generated samples are guaranteed to be compliant with given requirements. This is achieved by automatically parsing the constraints and transforming them into a Constraint Layer seamlessly integrated with the models, following neuro-symbolic principles.The Constraint Layer not only ensures that the requirements are satisfied, but

also that the generated samples' distribution better matches the real data distribution compared to the samples from the baseline models.

NIT: Neural Implicit Trajectories using Diffusion-Guided Scores

Estimating precise camera pose trajectories is a longstanding problem in computer vision with downstream applications in robotics and augmented reality. Current methods characterize a camera trajectory as a sequence of discrete camera poses, but we argue that this does not accurately represent the continuous property of camera motion. Therefore, we propose NIT, a Neural Implicit Trajectory that uses a simple MLP to represent the sequence of poses. We train NITs with an SDS-style approach using pretrained diffusion pose models, and we show the properties and advantages of using our representation.

Session IV

Transitioning towards a proactive practice: A longitudinal field study on the implementation of a ML system in adult social care

Tyler Reinmund

Politicians and care associations advocate for the use of machine learning (ML) systems to improve the delivery of adult social services. Yet, guidance on how to implement ML systems remains limited and research indicates that future implementation efforts are likely to encounter difficulties. We aim to enhance the understanding of ML system implementations by conducting a longitudinal field study with a team responsible for deploying a ML system within an adult social services department. The ML system implementation represented a cross-organisational effort to facilitate the department's transition to a proactive practice. Throughout this process, stakeholders adapted to numerous challenges in real-time. This study makes three contributions. First, we provide a description of how ML systems are implemented and highlight practical challenges. Second, we illustrate the utility of HCI knowledge in designing workflows for ML-assisted preventative care programmes. Finally, we provide recommendations for future deployments of ML systems in social care.

Security analysis of cutting-edge adaptive protocols in space systems

Edward Salkield

Adaptive Coding and Modulation space protocols are currently being implemented and standardised to maximise data throughput on satellite links in the presence of signal fading and radio interference. We analysed these protocols, finding design vulnerabilities common to the different implementations. This talk will discuss the security model and implications of the research, which is under discussion with the European Space Agency, and will be published at ESA 3S next month.

Abstracts of Lightning Talks

Lightning I

PQD: Post-Training Quantization For Diffusion Models

Jiaojiao Ye

Diffusion models (DMs) have demonstrated remarkable achievements in synthesizing images of high fidelity and diversity. However, the extensive computational requirements and slow generative speed of diffusion models have limited their widespread adoption. In this paper, we propose a novel post-training quantization for diffusion models (PQD), which is a time-aware optimization framework for diffusion models based on post-training quantization. The proposed frame- work optimizes the inference process by conducting time- aware calibration. Experimental results show that our proposed method is able to directly quantize full-precision diffusion models into 8-bit or 4-bit models while maintaining comparable performance in a training-free manner, achieving a few FID change on ImageNet for unconditional image generation. Our approach demonstrates compatibility and can also be applied to 512x512 text-guided image generation for the first time.

AutoVCL: A Self-balancing Algorithm for Variational Continual Learning

Fan Yang

Variational continual learning is a turn-key learning algorithm that has state-of-the- art performance among the best continual learning models. In our work, we explore an extension of the generalized variational continual learning (GVCL) model, named AutoVCL, which combines task heuristics for informed learning and model optimization. We demonstrate that our model outperforms the standard GVCL with fixed hyperparameters, benefiting from the adjustment of the hyperparameter based on the difficulty and similarity of the incoming task compared to the previous tasks.

Missing Data Imputation

Siyi Sun

My presentation should be a review of current traditional machine learning and deep learning imputation methods for missing values, as well as my current research about incorprating transformer into data imputation process to elevate the imputation accuracy.

An elementary theory of machine learning tasks

Benjamin Rodatz

Idealised as universal approximators, learners such as neural networks can be viewed as "variable functions" that may become one of a range of concrete functions after training. In the same way that equations constrain the possible values of variables in algebra, we may view objective functions as constraints on the behaviour of learners. We extract the equivalences perfectly optimised objective functions impose, calling them "tasks". For these tasks, we develop a formal graphical language that allows us to: (1) separate the core tasks of a behaviour from its implementation details; (2) reason about and design behaviours model-agnostically; and (3) simply describe and unify approaches in machine learning across domains. As proof-of-concept, we design a novel task that enables converting classifiers into generative models we call "manipulators", which we implement by directly translating task specifications into code. The resulting models exhibit capabilities such as style transfer and interpretable latent-space editing, without the need for custom architectures, adversarial training or random sampling. We formally relate the behaviour of manipulators to GANs, and empirically demonstrate their competitive performance with VAEs. We report on experiments across vision and language domains aiming to characterise manipulators as approximate Bayesian inversions of discriminative classifiers.

Lightning II

Diffusion Models for 3D Reconstruction and Articulation of Deformable Objects

Oishi Deb

3D object articulation and reconstruction of deformable objects hold vast potential for applications in fields such as education, AR/VR, and medicine. However, creating and articulating such 3D models is a very cumbersome job and would need expertise in the handling of Articulated 3D graphics by highly skilled 3D Graphic Designers. This poses a challenging barrier to enabling a wide range of articulated 3D models and being able to create models for new object categories. Our goal is to leverage deep learning models to automatically generate these assets from a single input image. This task, while highly practical, is challenging as the model needs an inherent understanding of the potential shapes and appearances of the object. Gaining ground-truth data in 3D to develop this prior information is a very labour-intensive process. Recently, there has been increased interest in learning these priors from the vast array of images available on the internet. However, this approach has some drawbacks such as the absence of Multiview constraints, noisy data, occlusion, and a lack of different viewpoints. These challenges make current state-of-the-art methods fail in achieving accurate and high-fidelity results for novel articulated 3D objects. We address these limitations by integrating existing methods with advanced pre-trained 2D text-to-image Diffusion Models. These models introduce additional priors, yielding more accurate 3D reconstruction, and thus, it can generalize for unseen categories without requiring more training images.

Hairy Ground Truth Enhancement for Semantic Segmentation

Semantic segmentation is a key task within applications of machine learning for medical imaging, requiring large amounts of medical scans annotated by clinicians. The high cost of data annotation means that models need to make the most of all available ground truth masks; yet many models consider two false positive (or false negative) pixel predictions as 'equally wrong' regardless of the individual pixels' relative position to the ground truth mask. These methods also have no sense of whether a pixel is solitary or belongs to a contiguous group.

Enterprise Information Governance, Domain Specificity, and the Data Mesh Case

Alastair McCullough

This presentation offers a unique insight into research into the unfolding discipline of enterprise information governance. It considers what we may mean by such governance and looks at the new concept of domain specificity, using the globally-known data mesh sociotechnical method to help to understand such localisation in governance.

Conversational AI at a Web scale

Jesse Wright

In this talk, we propose a protocol for Trustworthy and Accountable Agentic Dialogues at a Web Scale. Interest in conversational AI is increasing amongst LLM researchers. However, the challenge of making Trustworthy and Reliable Web Agents remains open. We propose a near-term solution to this problem, by outlining a protocol which leverages existing Semantic Web technologies to provide a harness for conversational AI dialogues ensuring unambiguous, and logically sound, descriptions of (1) usage restrictions on exchanged data (2) data origins and provenance, and (3) transactional outcomes of dialogues. We implement personal AI agents which adhere to this protocol and demonstrate their abilities to perform scheduling tasks on the behalf of humans. Additionally, we use these agents to explore human-centric design patterns that enable balanced human-in-the-loop oversight according to user preferences

Abstracts of Posters

Multiview Diffusion Models for Deformable Object Generation

Can Koz

The advent of multiview diffusion models has opened new avenues in the generation of deformable objects, offering significant improvements in consistency, detail, and realism. This paper introduces a novel framework that leverages the capabilities of multiview diffusion models to generate high-fidelity representations of deformable objects. Our approach, extends MagicPony, synthesizes deformable object geometries from multiple viewpoints, ensuring that the generated shapes are not only visually plausible but also physically coherent when subjected to deformation.

Orbit-Equivariant Graph Neural Networks

Matthew Morris

Equivariance is an important structural property that is captured by architectures such as graph neural networks (GNNs). However, equivariant graph functions cannot produce different outputs for similar nodes, which may be undesirable when the function is trying to optimize some global graph property. In this work, we define orbit-equivariance, a relaxation of equivariance which allows for such functions whilst retaining important structural inductive biases. We situate the property in the hierarchy of graph functions, define a taxonomy of orbit-equivariant functions, and provide four different ways to achieve non-equivariant GNNs. For each, we analyze their expressivity with respect to orbit-equivariance and evaluate them on two novel datasets, one of which stems from a real-world use-case of designing optimal bioisosteres.

ASIF: Coupled Data Turns Unimodal Models to Multimodal Without Training

Antonio Norelli

I will show how multimodality can be achieved without any training, just by treating image representations **as if** they were text representations.

With ASIF, you can build a CLIP-like model with common unimodal pretrained encoders (like DINO and SentenceT) and a relatively small set of image-text pairs in minutes, without tuning a single neuron.

Furthermore, I will discuss how this construction comes with unique properties. Most notably, the possibility to explain the outcome of a classification in terms of the image-text pairs and iterate over them, since deploying a new version with updated samples takes a matter of seconds.

I will show experiments on standard zero-shot visual benchmarks that demonstrate the typical transfer ability of image-text models. Overall, ASIF represents a simple yet surprisingly strong baseline for foundation multimodal models, raising important questions on their data efficiency and on the role of retrieval in machine learning.

Fast & Fair: A Collaborative Platform for Fair Division Applications

Fair division, the study of how to fairly allocate resources among agents, has received substantial interest in the areas of artificial intelligence and multiagent systems. While there is an extensive theoretical literature on fair division by now, the developed algorithms are still mostly confined to research papers and inaccessible to the public. We attempt to bridge this gap by developing Fast & Fair, an open-source web application that hosts a number of fair allocation algorithms with user-friendly interfaces and explainable outcomes. In contrast to existing implementations, Fast & Fair is a collaborative platform that is open to community contributions and thereby facilitates the deployment of additional algorithms.

Orbit-finite-dimensional vector spaces

Jingjie Yang

Motivated by decidability issues of weighted register automata, we look at vector spaces that have an infinite but orbit-finite basis: for example, $\mathrm{Span}\{e_1, e_2, e_3, \ldots\}$ with a group action $(1\,2) * (e_1 - e_4) = e_2 - e_4$. I will present a classification result for the equivariant subspaces which says that counting local coefficients is all you need.

Faithful Rule Extraction for Differentiable Rule Learning Models

Xiaxia Wang

There is increasing interest in methods for extracting interpretable rules from ML models trained to solve a wide range of tasks over knowledge graphs (KGs), such as KG completion, node classification, question answering and recommendation. Many such approaches, however, lack formal guarantees establishing the precise relationship between the model and the extracted rules, and this lack of assurance becomes especially problematic when the extracted rules are applied in safety-critical contexts or to ensure compliance with legal requirements. Recent research has examined whether the rules derived from the influential NEURAL-LP model exhibit soundness (or completeness), which means that the results obtained by applying the model to any dataset always contain (or are contained in) the results obtained by applying the rules to the same dataset. In this work, we extend this analysis to the context of DRUM, an approach that has demonstrated superior practical performance. After observing that the rules currently extracted from a DRUM model can be unsound and/or incomplete, we propose a novel algorithm where the output rules, expressed in an extension of Datalog, ensure both soundness and completeness. This algorithm, however, can

be inefficient in practice and hence we propose additional constraints to DRUM models facilitating rule extraction, albeit at the expense of reduced expressive power.

Understanding Challenges and Design Opportunities for Digital Mental Well-Being Support in Saudi Arabia

Mental health and well-being difficulties are considered a growing and highly stigmatized concern in the Kingdom of Saudi Arabia (KSA), especially among young women. Saudi young women have a significantly high lifetime risk of stress, emotional distress, anxiety and depression. Despite the high interest in mobile health (mHealth) in the KSA and its potential to circumvent traditional barriers, research on its application in the Saudi mental well-being context is scarce. To fill this gap, this study aims to understand the main opportunities that mobile apps provide to support mental well-being and barriers to using existing mental health mobile apps from the perspective of young Saudi women. To achieve this goal, we have conducted semi-structured interviews with 20 young Saudi women. Interviews were designed to obtain deeper insights into local barriers and design opportunities, including how participants' cultural values affect their use of mental health apps. Interviews were transcribed and analyzed using Thematic Analysis. Our results suggest that while concerns about stigma and cultural barriers hinder seeking help, the introduction of technology providing privacy and anonymity has opened up new avenues for accessing mental support in the KSA. We discussed some inclusive design opportunities for digital mental health interventions that sensitively consider the identified cultural challenge and incorporate Saudi women's perspectives and design needs. This work further contributes to the wider HCI research on inclusive design targeting women's needs, where most of the existing research focused on the Western context to infer women's needs while the requirements of Muslim Arab women were yet under-researched.