



REPORT ON BCTCS 2022

The 38th British Colloquium for Theoretical Computer Science

11–13 April 2022, Swansea University

Olga Petrovska and Monika Seisenberger

The British Colloquium for Theoretical Computer Science (BCTCS) is an annual forum in which researchers in Theoretical Computer Science can meet, present research findings, and discuss developments in the field. It also provides a welcoming environment for PhD students to gain experience in presenting their work to a broader audience, and to benefit from contact with established researchers.

BCTCS 2022 was hosted by Swansea University and held from 11th to 13th April 2022. The event attracted 67 registered participants, and featured an interesting and wide-ranging programme, including special sessions on explainable AI and theoretical aspects of security. A total of 28 contributed talks – predominantly by PhD students – were presented at the meeting alongside the six special session speakers and six keynote speakers.

A special theme for the meeting was **History of Theoretical Computer Science**. The four Keynote Speakers featured during the middle day of the colloquium were all chosen based on their lifetime achievements in developing theoretical computer science in the UK, and were invited to present personal histories of their fields of expertise. The meeting also featured a special discussion group on the pedagogy of theoretical computer science, led by Barnaby Martin and Eleni Akrida. Finally, there was a special session following the AGM devoted to celebrating two important birthdays: John Tucker, 70, who founded BCTCS in 1985; and Faron Moller, 60, who served as its President for 15 years (2004-2019).

BCTCS 2023 will be hosted by Glasgow University from 3rd–5th March 2023. Researchers and PhD students wishing to contribute talks concerning any aspect of Theoretical Computer Science are cordially invited to do so. Further details are available from the BCTCS website at www.bctcs.ac.uk.

Invited Talks

Cliff Jones (Newcastle University)

Formal Methods in the UK - A (Personal) History

I owe my real start in “formal methods” to two spells in Vienna but that was long enough ago that I can talk about many years of developing and promoting such approaches in the UK. I’d like to describe how difficult things were back in the 70s, how much better they are today, but also to address why it has taken so long and the situation is still far from ideal. I won’t have time to get very technical but will provide pointers to more technical material.

Alexander Knapp (Ausburg University)

Specifying Event/Data-based Systems

Event/data-based systems are controlled by events, their local data state may change in reaction to events. Numerous methods and notations for specifying such reactive systems have been designed, though with varying focus on the different development steps and their refinement relations. We first briefly review some of such methods, like temporal/modal logic, TLA, UML state machines, symbolic transition systems, CSP, synchronous languages, and Event-B with their support for parallel composition and refinement. We then present $E\downarrow$ -logic for covering a broad range of abstraction levels of event/data-based systems from abstract requirements to constructive specifications in a uniform foundation. $E\downarrow$ -logic uses diamond and box modalities over structured events adopted from dynamic logic, for recursive process specifications it offers (control) state variables and binders from hybrid logic. The semantic interpretation relies on event/data transition systems; specification refinement is defined by model class inclusion. Constructive operational specifications given by state transition graphs can be characterised by a single $E\downarrow$ -sentence. Also a variety of implementation constructors is available in $E\downarrow$ -logic to support, among others, event refinement and parallel composition. Thus the whole development process can rely on $E\downarrow$ -logic and its semantics as a common basis.

Mike Paterson (University of Warwick)

Algorithms and Complexity in the UK - A (Personal) History

One person’s story of how they found their way into this area during the 60’s and 70’s.

Rick Thomas (University of Leicester / University of St Andrews)

Some Interactions Between Automata Theory and Algebraic Structures in the UK - A (Personal) History

We will survey some of the impact that automata and formal language theory has

had on the theory of algebraic structures, particularly group theory, focussing on some work done in the UK. We will also point out some results and questions in automata theory arising from these interactions.

Francesca Toni (Imperial College London)

Argumentative Explainable AI

It is widely acknowledged that transparency of automated decision making is crucial for deployability of intelligent systems, and explaining the reasons why some outputs are computed is a way to achieve this transparency. The form that explanations should take, however, is much less clear. In this talk I will explore two classes of explanations, which I call 'lean' and 'mechanistic': the former focus on the inputs contributing to decisions given in output; the latter reflect instead the internal functioning of the automated decision making fed with the inputs and computing those outputs. I will show how both classes of explanations can be supported by forms of computational argumentation, and will describe forms of argumentative XAI in some settings (e.g. for multi-attribute decision making and machine learning).

John Tucker (Swansea University)

A (Personal) History of Logic in Computer Science

I will look at the emergence of logic as a central source of tools in computer science through the lens of my scientific development. I will concentrate on mathematical logic and its impact on programming and software engineering, starting in the 1960s.

Contributed Talks

Ruba Alassaf (University of Manchester)

Uniform interpolation in modal logic

A formula ϕ' is said to be a *uniform Σ -interpolant* of ϕ if, for any ψ over the signature Σ , ψ is implied by ϕ' if and only if ψ is implied by ϕ . Uniform interpolation has been the subject of many recent research papers due to its potential use in applications. In this talk, we introduce uniform interpolation and its variants, and present a saturation-based inference system that computes a *local* uniform interpolant for a formula and a “keep” signature.

Fahad Alhabardi (Swansea University)

Verification of bitcoin’s smart contracts in Agda using weakest preconditions for access control

We address the verification of bitcoin smart contracts using the interactive theorem prover Agda, focusing on two standard smart contracts that govern the distribution of bitcoins: Pay to Public Key Hash (P2PKH) and Pay to MultiSig (P2MS). Both are written in bitcoin’s low-level language script, and provide the security property of access control to the distribution of bitcoins. We introduce an operational semantics of the script commands used in P2PKH and P2MS, and formalise it in the Agda proof assistant using Hoare triples. We advocate weakest preconditions in the context of Hoare triples as the appropriate notion for verifying access control. Two methodologies for obtaining human-readable weakest preconditions are discussed in order to close the validation gap between user requirements and formal specification of smart contracts: (1) a step-by-step approach, which works backwards instruction by instruction through a script, sometimes stepping over several instructions in one go; (2) symbolic execution of the code and translation into a nested case distinction, which allows to read off weakest preconditions as the disjunction of accepting paths. A syntax for equational reasoning with Hoare Triples is defined in order to formalise those approaches in Agda.

Richard Allen (Swansea University)

Runtime verification for Android security

Users of computer systems face a constant threat of cyberattacks by malware designed to cause harm or disruption to services, steal information, or hold the user to ransom. Cyberattacks are becoming increasingly prevalent on mobile devices like Android. Attacks and countermeasures become more sophisticated in an ever-increasing arms race. A novel attack method is “collusion”, where the attack gets hidden by distributing the steps through many malicious software actors. We investigate the use of runtime verification to detect collusion attacks on the end-users device. We have developed a novel algorithm called *RRH* that is a variation

of an existing algorithm by Grigore Roşu and Klaus Havelund. Our approach is computationally efficient enough to detect collusion in realtime on the Android device and does not require prior knowledge of malware source code. Thus, it can detect future malware without modification to the detection system or the software under scrutiny.

Peace Ayegba (University of Glasgow)

Resource allocation problems in wireless communications

Wireless communications is transforming every sector of the economy as it enables billions of people to keep in touch with work and family via their smart devices. However, the current technology cannot meet the demands of the future (i.e., explosive growth in number of devices, better coverage and connection rate). For future wireless communications, the key technology that has the potential to enhance connectivity for billions of users is referred to as Cell Free massive Multiple-Input Multiple-Output (CF-mMIMO). One of the many challenges in CF-mMIMO is how to efficiently manage limited resources (spectrum, energy, and power) – the so called resource allocation problem in wireless communications. This talk will give a survey of some of the problems arising in this context, as well as mathematical techniques that have been explored to tackle them.

Martin Barrere (Imperial College London)

Trustworthy critical infrastructure systems

Critical infrastructure systems are systems that are considered vital for the proper functioning and development of a society. Some examples are power plants, smart grids, and water distribution networks. These systems are often composed of diverse interconnected subsystems, thus creating a complex network of highly interdependent software, processes, and hardware components. From a security perspective, cyber attacks on these cyber-physical systems can have very serious consequences such as flooding, blackouts, or even nuclear disasters. In this talk, I will cover some of the main concepts regarding critical infrastructure systems, their importance, and related security management challenges. Afterwards, I will present two of the research areas that we are currently working on, namely, identification of critical cyber-physical components and cyber-physical attack graphs. I will conclude my talk by describing a number of open problems and exciting security challenges that require further investigation.

Arnold Beckmann (Swansea University)

Blockchain-based cyber-physical trust systems and their application

Cyber Physical Trust Systems (CPTS) are cyber physical systems enriched with trust as an explicit, measurable, testable and verifiable system component. In this talk we describe our research in relation to CPTS that are driven by blockchain.

We explain what blockchain is. We highlight some of the security properties of Blockchain-based CPTS and their formal proofs using the Tamarin Prover tool. We also describe the context of this research by highlighting the application of Blockchain-based CPTS in real world projects we are involved in.

Harry Bryant (Swansea University)

Exploring the IC3 algorithm to improve the Siemens-Swansea ladder logic verification tool

The last years have seen a radical transformation in railways. Among the new infrastructure, signalling software is a key technology to increase rail capacity and automation. However, verifying the safety of this software is a challenge for the railway industry. Siemens Mobility and Swansea University have developed an effective automated verifier for railway interlocking programs, which our project sets out to improve. The purpose of our research is to deal efficiently with the false-positive problem, where verification fails although the property under discussion holds. This problem arises due to the specific algorithmic approach taken. We are investigating if the false-positive problem can be addressed by Bradley's IC3 algorithm and its parallel and symbolic variations. The overall goal of our research is to offer a solution that scales to a challenge in rail industry.

Victor Cai (Swansea University)

Counterexample visualisation in the railway domain

In railway, the placement of equipment on a track plan is safety-critical and underlies many constraints (more than 300 are known in industry). For ETCS Level 2, our industrial partner has a design tool that lets engineers place equipment on track plans. The objective of our project is to visualise constraint violations in this tool, discovered through SMT solving. There is an approach under development that translates data from the design tool into a directed graph, which is then encoded in the SMTLib2 format, and finally checked against a set of design rules using the Z3 SMT solver. Our contribution is to take the unsatisfiable core of such a check and visualise it on the track plan in the design tool, by reversing the two model transformations that made the track data accessible to SMT solving. Our work provides a new instance of the challenge to translate from the language of a prover back into the language of the domain from which the problem originated.

Jamie Duell (Swansea University)

On explainable AI for medical diagnostics and its potential

Explainable AI (XAI) aims to provide intelligible explanations to users to support decision making. Current XAI methods, such as SHAP, LIME, and Scoped Rules, focus on calculating feature importance in predictions. Although XAI has attracted increasing attention, applying XAI techniques to healthcare to inform

clinical decision making is still challenging. Our experimental results show that XAI methods circumstantially generate different top features. The identified features, including shared top important features, demonstrate clinical significance in making diagnosis under different scenarios. Our work shows that medical diagnostics is a promising domain for applying XAI technologies. We aim to further this coalition of important features and the respective impact to surrounding features, by introducing peer-influence as a means of identifying underlying relations.

Hsuan Fu (Université Laval)

XAI in finance and economics

Abdul Ghani (Durham University)

Depth lower bounds in stabbing planes for combinatorial principles

This talk will introduce stabbing planes as a proof system, discuss its relevance, and indicate a few methods of finding lower bounds.

James Hinns (Swansea University)

Quantifying underspecification with feature attribution algorithms

For a given dataset and machine learning (ML) task, various predictors can be constructed with different reasoning that produce near-optimal test performance. However, due to this variance in reasoning, some predictors can generalise whilst some perform unexpectedly on unseen data. The existence of multiple equally performing predictors exhibits *underspecification* of the ML pipeline used for producing such predictors and presents challenges for credibility. In this talk, we propose identifying underspecification by estimating the variance of reasoning within a set of near-optimal predictors produced by a pipeline. We implement our approach by producing a set of well-performing predictors, computing their Shapley additive explanations and then measuring various distance and correlation metrics between them. We demonstrate our approach on multiple datasets drawn from the literature, and in a COVID-19 virus transmission case study.

Sahar Jahani (London School of Economics)

Automated equilibrium analysis of $2 \times 2 \times 2$ games

In this project, we study the Nash equilibria in three-player non-cooperative games in which each player has two strategies. We present a program that outputs a full description of these equilibria even for degenerate games, which has not been done before. These are the simplest games with more than two players. The players' best-response functions are non-linear and a new formulation for these is proposed. Finding Nash equilibria is divided into two main algorithms: one to compute the partially mixed equilibria, and one compute the completely mixed equilibria. Both algorithms are implemented in Python with a 3D representation of

best-response surfaces. The program will be added to the Game Theory Explorer, which is a web application for analyzing different types of games.

Amin Karamlou (University of Oxford)

Logical aspects of non-local games

A non-local game consists of a group of spatially separated players who are trying to convince an adversary (the referee) that they know the answers to a set of questions. Ever since the pioneering work of Bell in the 1960s this rather abstract framework has proven to be one of our most useful tools for understanding quantum advantage, that is, the ability of a quantum system to outperform a classical one. We focus on analysing non-local games by using techniques traditionally associated with the field of logic in computer science. I present an overview of the questions in this area, and the progress made thus far in solving them.

Nina Klobas (Durham University)

The complexity of temporal vertex cover in small-degree graphs

Temporal graphs model graphs whose underlying topology changes over time. Recently, the *temporal vertex cover (TVC)* and *sliding window temporal vertex cover (Δ -TVC for time-windows of a fixed-length Δ)* have been established as natural extensions of the classic vertex cover problem on static graphs with connections to areas such as surveillance in sensor networks. In this talk we initiate a systematic study of the complexity of TVC and Δ -TVC on sparse graphs. Our main result shows that for every $\Delta \geq 2$, Δ -TVC is NP-hard even when the underlying topology is described by a path or a cycle. This resolves an open problem and shows a surprising contrast between Δ -TVC and TVC for which we provide a polynomial-time algorithm. To circumvent this hardness, we present a number of exact and approximation algorithms for temporal graphs whose underlying topologies are given by a path, that have bounded vertex degree in every time step, or that admit a small-sized temporal vertex cover.

Veera Raghava Reddy Kovvuri (Swansea University)

Understanding the influence of controllable factors in feature attribution algorithms

Feature attribution algorithms enable their users to gain insight into the underlying patterns of large data sets through their feature importance calculation. Existing feature attribution algorithms treat all features in a data set homogeneously, which may lead to misinterpretation of consequences of changing feature values. In this work, we consider partitioning features into controllable and uncontrollable features, and propose the Controllable fActor Feature Attribution (CAFA) approach to compute the relative importance of controllable features. CAFA is different from existing feature attribution algorithms in the following aspects. First, as

already mentioned, we distinguish between controllable and uncontrollable features instead of treating them homogeneously. Secondly, we compute the relative importance of controllable features through selective perturbation and global-for-local interpretation. We applied CAFA to two medical datasets, a lung cancer dataset and a breast cancer dataset. The experimental results show that after using CAFA, though the model is built over all features, the explanations of controllable features do not interfere with the uncontrollable ones. We further studied CAFA on the COVID-19 non-pharmaceutical control measures dataset, which was collected by ourselves. This study also showed that the controllable and uncontrollable features are relatively independent.

Pardeep Kumar (Swansea University)

Security in cyber physical smart environments

Cyber physical systems (CPSs) are envisioned as one of the building blocks for the internet of things (IoT), critical infrastructures, automations, smart healthcare, smart environments, etc. In a typical CPS network, a sensor works as an eye and an actuator works as a leg to cater their respective applications via several low-powered heterogeneous communication technologies, protocols, and systems. However, such CPSs and their applications are largely distributed systems and subject to several attacks such as malicious attempts to hijack and denial of services. This talk will discuss some of the technical challenges and a lightweight security mechanism in cyber physical networks.

David Kutner (Durham University)

Payment scheduling in the interval debt model

Financial networks consist of banks (nodes) connected by debts (edges). We introduce the interval debt model, in which debts are due at integer time intervals on a temporal graph with lifetime T . We define valid payment schedules as ones consistent with the obligations of the nodes in the input, and study the problem of finding a schedule which minimizes (or maximizes) the number of nodes which go bankrupt. We also study the problem of bailout minimization, in which the aim is to minimize the bailout payments made by a financial authority to make possible a schedule with no bankruptcies. We find that bankruptcy minimization and maximization are NP-hard even on directed acyclic graphs of bounded degree; that bailout minimization is NP-hard on graphs of bounded degree; and that bankruptcy minimization remains weakly NP-hard on graphs with $O(1)$ vertices.

Laura Larios-Jones (University of Glasgow)

Reordering edges in temporal graphs to maximise reachability

Temporal graphs consist of an underlying graph (G, E) and an assignment t of timesteps to edges that specifies when each edge is active. This allows us to

model spread through a network which is time-sensitive. We will keep the set of timesteps the same and explore reordering them to optimise reachability. Previous work has mainly explored minimising spread for applications such as epidemiology. Here, we will be looking at the opposite problem of increasing movement through a graph. In particular, our goal is to reorder the timesteps assigned to the edges in our graph such that the minimum number of vertices reachable from any starting vertex is maximised. We will discuss which analogous structural and algorithmic results from the minimisation case hold in this case. Maximising spread can be useful in situations where we would like information or resources to be shared efficiently, such as advertising or even vaccine rollout.

Dimitrios Los (University of Cambridge)

Balanced allocations with incomplete information

In balanced allocations, we have to allocate m tasks (balls) sequentially into n servers (bins) so as to minimise the gap, i.e. the difference of the maximum load to the mean m/n . The simple process of allocating each ball to a random bin achieves whp (with high probability) a $\Theta(\sqrt{m/n \log n})$ gap when $m \gg n$. A great improvement over this process is the two-choice process, where each ball queries the load of two randomly chosen bins and is then placed in the least loaded bin. It achieves whp a $\Theta(\log_2 \log n)$ gap for $m \geq n$. In the first part of this talk we give some background on balanced allocation processes. In the second part, we investigate variants of the two-choice process where allocations are made under incomplete information, giving insights into this dramatic improvement of two-choice over one-choice.

Michael McKay (University of Glasgow)

Envy-freeness in fixed size coalition formation games

In this talk we consider a variation of the three-dimensional stable roommates problem, which can be thought of as a coalition formation game. Agents specify preference lists over their peers, and the task is to partition the agents into sets of size 3 based on these preferences. Instead of seeking a “stable” partition, our aim is to construct a partition in which no agent has “envy” for another. There is a close relation to stability. In this talk we assume that agents have additive preferences. For three definitions of envy-freeness, we show that it is NP-hard in general to find envy-free partitions, but we present new results relating to restrictions of each problem in which an envy-free partition can be found in polynomial time. These results help us explore the boundary between NP-hardness and polynomial-time solvability in problems of coalition formation.

Steffan Mon (Swansea University)

Routing as a game

The internet can be viewed as a game played on a graph, where each vertex makes forwarding decisions based on selfish preferences over the path the data takes. Changes in forwarding decisions by one vertex may cause others to change their decisions, causing a ripple effect through the graph, resulting in degraded performance. Path formation games formalise and model this scenario along with other problems of interest. Although not every instance of path formation games has a solution, we prove how some special cases are guaranteed to have an efficiently computable solution. We also show how lacking a structure known as a dispute wheel is a sufficient condition for a game to have a solution.

Yoàv Montacute (University of Cambridge)

The path of hidden pebbles

Pebble games are a type of Spoiler-Duplicator game used in finite model theory to compare models and examine expressive powers of logical systems. Recently, these games were shown to provide an equivalent intermediary between fragments of different logical languages and categorical structures named comonads. This realises a unified language that can be used to study combinatorics categorically and to introduce alternative consolidated proofs for results with a common general form. This talk will begin with a current overview on this research project, originating from the seminal work of Abramsky, Dawar and Wang. We then consider a new type of pebble game which utilises the idea of a hidden pebble and yields a novel homomorphism-counting theorem via a comonad whose coalgebras characterise pathwidth.

Kiana Samadpour Motalebi (University of Manchester)

Tableau methods for modal logic

Logics of confluence, defined by generalisations of the axiom of confluence, cover a wide range of standard modal logics including modal logic K and single axiom extensions with the axioms T , B , S , BAN , $alt1$, $G0111$, G and D , which are simple instances of a generalised pattern of axioms of confluence. Two kinds of tableau calculi for this class of logics using respectively structural rules and propagation rules are studied. Structural rules ensure that the constructed models satisfy the characteristic frame conditions of the logic, whereas propagation rules only construct pre-models which are sufficient to discover unsatisfiability and can be completed or adapted to give concrete models.

Tamio-Vesa Nakajima (Oxford University)

Linearly ordered colouring of hypergraphs

Linearly ordered colouring is a type of special hypergraph colouring. In normal hypergraph colouring, one must colour each vertex with one of K different colours such that no hyperedge is monochromatic. In linearly ordered colouring,

one must colour each edge so that it has a unique maximum colour. In this talk we discuss several interesting results about such colourings – both tractability of certain cases, and hardness of others. In particular, we will discuss “promise” problems: rather than trying to find a linearly ordered colouring with K colours for arbitrary hypergraphs, we will assume that the hypergraph admits a linearly ordered colouring with fewer than K colours.

Hoang Nga Nguyen (Coventry University)

Towards systematic threat assessment and security testing for automotive OTA

Modern cars host numerous special-purpose computing and connectivity devices facilitating the correct functioning of various in-vehicle systems. These devices host complex software systems with over 100 million lines of code, requiring regular and timely updates for functional and security improvements. Addressing the shortcomings of the legacy update system, over-the-air (OTA) software update systems have emerged as an efficient, cost-effective, and convenient solution for delivering updates to automobiles remotely. While OTA offers several benefits, it introduces new security challenges requiring immediate attention, as attackers can abuse these update systems to undermine vehicle security and safety. This talk will present our model-based approaches to addressing the security assessment problem for automotive OTA systematically.

Uchenna Nnawuchi (Middlesex University)

Investigating the right to explainable AI in the context of the GDPR

The growth of Artificial Intelligence (AI) in various aspects of human society has affected the landscape of human cognition, social order and political power. Machine Learning (ML) algorithms are increasingly becoming responsible for decision-making in many industry sectors owing to their efficiency in performing tasks. However, the growing worry is that ML algorithms are fraught with opacity and bias. Some legal scholars and data scientists have claimed that the right to explanation exists and is a persuasively appealing approach for challenging and providing redress to ML opacity. However, the existence or otherwise of such a right is at the centre of global debate. This paper investigates whether there is a legal right to explanation of ML algorithms in the context of the GDPR. The paper analyses the argument from various scholars on both sides of the divide with a view to ascertaining whether the right actually exists.

Alpay Ozkeskin (University of Liverpool)

Scheduling with precedence constraints for electricity cost in smart grid

We study a scheduling problem arising in demand response management in smart grid. Consumers send in power requests with a flexible feasible time interval during which their requests can be served. The grid controller, upon receiving power

requests, schedules each request within the specified interval. The electricity cost is measured by a convex function of the load in each time slot. The objective is to schedule all requests with the minimum total electricity cost. Previous work has studied a special case where the power requirement and the duration a request demands are both unit-size. We extend this case by allowing jobs to have precedence constraints. We present a polynomial time offline algorithm that gives an optimal solution when the feasible time intervals for the jobs in order of precedence are the same.

Filippos Pantekis (Swansea University)

Towards massively parallel SAT

Parallelisation of SAT is an area that faces many challenges, including the sharing of information between threads causing restrictions to performance. In this work, we present steps in massively parallelising SAT checking using new GPGPU architectures in a manner that is both loosely coupled and fully scalable. We demonstrate techniques for successful parallelisation of SAT in these special environments and present early results in favour of their use.

Francis Southern (Swansea University)

Singular DP-reduction

It is useful to be able to reduce one set of boolean constraints to another, which is simpler relative to some complexity measure, in such a way that solutions can be translated back to the original problem. One operation which achieves this and has been utilised in both theoretical and practical work on the propositional satisfiability problem is Davis-Putnam reduction, which eliminates a variable v from a clause-set F by replacing the clauses containing v with their resolvents (on v). In general, this may increase the number of clauses. However, the number of clauses also decreases in the special case of singular DP-reduction, in which the eliminated variable occurs in one of its polarities only once. My talk will cover some structural properties of singular DP-reduction on minimally unsatisfiable clause-sets.

Theofilos Triommatas (University of Liverpool)

On maximising the visibility area with a rotating field of view

Imagine we have a polygon-shaped platform P and only one static spotlight outside of P . Which direction should the spotlight face to light most of the platform? More formally, we define and provide an algorithm for the problem of finding the maximum intersection between a polygon and a rotating field of view. If we were to guess the direction that yields the maximum intersection, we would have to go through an infinite number of guesses. We provide an FPTAS that approximates the rotation that yields the maximum intersection.

Galen Wilkerson (Imperial College London / University of Surrey)

Spontaneous emergence of computation in network cascades

Neuronal network computation and computation by avalanche supporting networks are of interest to the fields of physics, computer science (computation theory as well as statistical or machine learning) and neuroscience. Here we show that computation of complex Boolean functions on inputs arise spontaneously in threshold networks as a function of connectivity and antagonism (inhibition), computed by logic motifs. We also obtain lower bounds on function probabilities, and show that these results agree with those obtained by others.

Adam Wyner (Swansea University)

Explainable AI in law

Tansholpan Zhanabekova (University of Liverpool)

Deciding what is good-for-MDPs

Nondeterministic Good-for-MDP (GFM) automata are for MDP model checking and reinforcement learning what good-for-games automata are for synthesis: a more compact alternative to deterministic automata that displays nondeterminism, but only so much that it can be resolved locally, such that a syntactic product can be analysed. GFM has recently been introduced as a property for reinforcement learning, where the simpler Büchi acceptance conditions it allows to use is key. However, while there are classic and novel techniques to obtain automata that are GFM, there has not been a decision procedure for checking whether or not an automaton is GFM. We show that GFM-ness is decidable and provide an EXPTIME decision procedure as well as a PSPACE-hardness proof.