AI @ NICTA

**Nick Barnes, Peter Baumgartner, Tiberio Caetano, Hugh Durrant-Whyte, Gerwin Klein,
Penelope Sanderson, Abdul Sattar, Peter Stuckey, Sylvie Thiebaux, Pascal Van Hentenryck, Toby Walsh**

Abstract

NICTA is Australia's Information and Communications Technology (ICT) Centre of Excellence. It is the largest organisation in Australia dedicated to ICT research. Whilst it has close links with local universities, it is in fact an independent but not-for-profit company in the business of doing research, commercialising that research and training PhD students to do that research. Much of the work taking place at NICTA involves various topics in artificial intelligence. In this article, we survey some of the AI work being undertaken at NICTA.

Introduction

NICTA is the largest ICT research centre in Australia, having been established ten years ago in 2002. It has five laboratories in four Australian capital cities: Sydney, Canberra, Melbourne and Brisbane. There are currently around 700 staff and PhD students working at NICTA. In June 2009, the 100th PhD student to study at NICTA graduated. At present and averaged over the year, one new PhD student studying at NICTA graduates every 10 days. NICTA has close links with its university members (Australian National University, the University of New South Wales and the University of Melbourne) as well as with its partner universities (University of Sydney, Griffith University, Queensland University of Technology, University of Queensland and most recently Monash University). Many of the researchers at NICTA are seconded from these universities. In addition, most of the other researchers at NICTA hold adjunct positions at one of these universities, enabling them to teach courses and supervise PhD students. NICTA also has close links with a number of other research organizations (including Australia's CSIRO, France's INRIA, Japan's NII, and Germany's Fraunhofer Institute) and major companies (including Microsoft, Google, SAP, and Ericsson).

Research vision

NICTA was established with two main objectives; to undertake leading fundamental research in ICT and to develop outcomes of commercial or national benefit from this research for Australia. In support of these objectives, NICTA is structured around six major research groups and four business teams. The research groups are in machine learning, networks, computer vision, software systems, optimization, and control and signal processing. Each group comprises between one and two hundred research staff and students. All of these groups are contributors in some way to AI research at NICTA. The business teams are in Broadband and the Digital Economy (BaDE), Infrastructure Transport and Logistics (ITL), Health, and Safety, Security and Environment (SSE). These business teams represent major applications of ICT especially in the Australian context. Each of these teams are major consumers of AI research through their engagement with the research Groups.

This mixture of both fundamental research and business outcomes provides a dynamic, productive and challenging environment for AI researchers of all persuasions. The research projects described here span the range from formal methods, planning and optimisation, to bioinformatics, computer vision and human-computer interaction. In the rest of this article, we look in more detail at some specific research areas and describe some of the research going on in the five NICTA laboratories.

Optimisation

One of the largest concentration of researchers in AI in NICTA works on optimisation. The research in this area has been driven by applications like routing vehicles, folding proteins, and scheduling traffic lights. The research explores the interface between several areas: constraint programming, operations research, satisfiability, search, automated reasoning, and machine learning. New projects in the optimisation area are addressing several topics of especial relevance to Australia including disaster management, smart grids and homes, supply chains and logistics, as well as the interface between optimisation, social choice and machine learning.

Constraint Programming

The optimisation group has considerable strength in both modelling and solving optimisation problems using constraint programming and related technologies. We have pioneered sophisticated modelling languages for optimisation like Zinc (Marriott *et al.* 2008) and MiniZinc (Nethercote *et al.* 2007) as part of the ambitious G12 project (Stuckey

Figure 1: NICTA's headquarter building on the Australian Technology Park in Sydney, Australia.

et al. 2005). The broader aims of the G12 project are to tackle the so called "modelling bottleneck", automating the process of taking the specification of an abstract optimisation problem and solving it. As part of this project, we have developed some ground breaking solving methods like lazy clause generation.

Whilst fundamental research questions like how to refine models automatically and deal with issues like symmetry (Walsh 2008) and computational complexity (Bessière *et al.* 2007) drive some of the research, there is also considerable input from practical real world problems. For instance, NICTA has a close relationship with the Road Traffic Authority (RTA) of New South Wales. The RTA develop and sell the SCATS traffic light control system. It is in one of the most widely used and successful traffic control systems, with installations in 142 cities across 25 countries. NICTA is currently trialing a new optimisation based signal control method at a major intersection south of Sydney. The system is predicted to improve the flow of traffic through the intersection in peak periods by 5%. Such savings will soon add up to considerable benefits. Traffic congestion is estimated to cost Australia over $10 billion annually, and this amount is set to double by 2020.

Satisfiability

NICTA has been undertaking fundamental research on various aspects of satisfiability (SAT) since its foundation. Research has ranged from SAT-encoded CSPs to encoding temporal and spatial reasoning problems, to exploiting problem structure for SAT local search, estimating the cost of SAT solving, parameter tuning and participating in the international SAT solver competitions. In each of these areas, we have produced a number of important results. In addition, we have solved several open challenges in the field.

A comprehensive study of the mappings between CSPs and SAT (Walsh 2000) and the development of algorithms that exploit the structure of SAT-encoding of CSPs (Pham *et al.* 2005) inspired a SAT encoding of qualitative temporal networks, resulting in an efficient solution to the well known temporal reasoning problem (Pham, Thornton, & Sattar 2008a). Later the SAT encoding approach was successfully applied to qualitative spatial reasoning problems (Li, Huang, & Renz 2009).

One of the recognised shortcomings of local search procedures for SAT is that they perform less well than complete algorithms on difficult structured problems, while generally doing much better on random problems. By taking some inspiration from the CSP structure exploiting approach (Pham *et al.* 2005), we developed a new approach that looked at discovering dependencies between variables and using this information to build a dependency lattice that guides a local search in such a way that only the independent variables in a problem are flipped. This resulted in significant improvements in the efficiency of local search on a number of widely recognized SAT challenge problems, including the well-known parity-32 problem, and won an IJCAI distinguished paper award (Pham, Thornton, & Sattar 2007). Further, an improved version for the first time outperformed a state of the art complete search solver on the parity-32 benchmarks (Pham, Thornton, & Sattar 2008b).

Other work on SAT includes both empirical and theoretical investigations into the power and efficiency of SAT algorithms, particularly concerning the use of restarts (Huang 2010a), the interplay between components of SAT algorithms (Huang 2007), and estimating cost of SAT solving in terms of the search tree size (Kilby *et al.* 2006) and runtime (Haim & Walsh 2008). Significant progress was made on one challenging problem in dynamic local search algorithms, namely parameter tunning (Thornton & Pham 2008). NICTA also played a key role in the preparation of the Handbook of Satisfiability, that provides a comprehensive account of theoretical and empirical studies of SAT algorithms, and applications (Biere *et al.* 2009).

Our SAT solvers based on ideas presented in (Anbulagan *et al.* 2005) and (Pham *et al.* 2008) entered into the biennial SAT solving competitions and won Gold medals for the random SAT category of the 2005 and 2007 rounds. Later, an improved version of (Pham *et al.* 2008) won the Silver medal for the random SAT category and the first place for the parallel track in the 2009 round.

In summary, NICTA has contributed to several areas of SAT research, and made significant progress on a number of the SAT challenge problems set out by (Selman, Kautz, & McAllester 1997). These include Challenge 2 on solving the parity-32 problem, Challenge 6 on developing a variable dependency approach for local search, and Challenge 8 on characterising problem encodings.

Vehicle routing

Another example of the "use inspiration" in research at NICTA is in the area of vehicle routing where we have built a flexible solver for a wide variety of logistics problems. This solver, called Indigo, is based on a combination of Operations Research (OR) and Artificial Intelligence (AI) techniques (Kilby & Verden 2002). Each company that has a logistics component to their daily activities has different business rules and business processes. These, in turn, give rise to different constraints on the solutions. Constraint Programming (CP) offers a natural way to express these constraints. A standard constraint programming solver can be used to propagate the effects of each of these constraints onto the emerging solution. The Indigo solver combines techniques from both the OR and AI literature. It uses a variety of OR construction methods to create an initial solution. An AI improvement method called Large Neighbourhood Search is then used to improve the routes. A bespoke CP system is used to formulate and solve a variety of side constraints not typically handled by traditional vehicle routing solvers, such as limited docks, mutual exclusion (service request A XOR request B) and precedence constraints (request A before request B). Propagators for these constraints can be written independently of any other constraint, making maintenance much easier under this paradigm. See Figure 2 for more details.

Another example of use inspiration in optimisation research can be seen in the Future Logistics Living Lab. This is a collaboration between NICTA, SAP and the Fraunhofer Institute to showcase the latest ICT technologies, and to provide a "sand pit" where major companies like Linfox and Hamburg Sud can come together to help transform the transport and logistics sector. NICTA has, for instance, been working with several major (over $1 billion revenue) fast moving manufacturing companies. Using the Indigo solver,

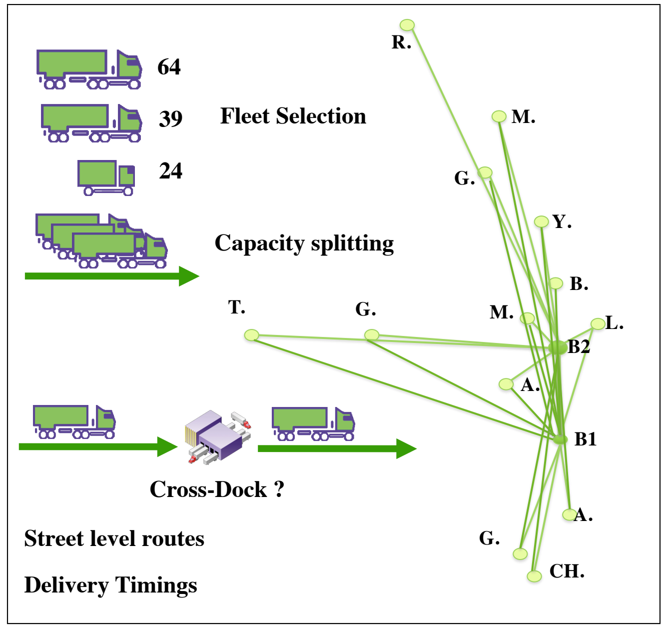

Figure 2: Example of a vehicle routing problem solved by the Indigo solver. The problem involves optimising the route choice based on street level mapping. The solver permits loads to be split as well as cross-docked. In addition, the best fleet mix is selected.

we have demonstrated how to produce significant savings in both local and regional distribution. Taking advantage of the solver's sophisticated modelling capabilities, we can answer complex strategic questions like how to optimise the fleet mix of trucks.

Planning and Diagnosis

NICTA does fundamental research on many aspects of automated planning and model-based diagnosis: from pathfinding to classical planning (satisficing and optimal) to probabilistic and temporal planning, and from diagnosis of circuits, to diagnosability analysis of discrete-event systems and networks of hybrid continuous and discrete systems.

In each of these areas, we have had a number of "world first" results. A prime example are methods for temporal *planning under uncertainty* that handle concurrency, time, actions with probabilistic outcomes and durations, continuous distributions and numeric state variables (Buffet & Aberdeen 2009; Little, Aberdeen, & Thiébaux 2005). RDDL, the new domain modelling language of the International Probabilistic Planning Competition, incorporates many of those features (Sanner 2010). More recently, we have designed the first exact methods for sequential decision processes with continuous non-linear stochastic dynamics (Sanner, Delgado, & de Barros 2011)

Influential contributions to *classical planning* via heuristic search, the approach that has dominated the last decade, include the h^m family of critical-paths heuristics, merge-and-shrink abstractions heuristics, and the landmark heuristic (Haslum 2006; Helmert, Haslum, & Hoffmann 2007; Richter & Westphal 2010). Our new results on planning via satisfiability, the rival approach, are moving SAT planning into untraditional territories such as cost-optimal and sub-optimal planning, where it is now competitive with heuristic search (Rintanen 2010; Robinson *et al.* 2010). Based on SAT, we have also designed some of the most efficient, generic methods for *diagnosis and diagnosability analysis* of discrete-event systems (Grastien *et al.* 2007).

Another driver is the construction of high-performance domain-independent planners and diagnosers based on those results. E.g., HSP*, LAMA, Madagascar, FPG, FOALP, NM-RDPP, and SATDiag are used by the community as benchmarks or received prizes at planning competitions. We are also exploring other aspects of the high-performance agenda, such as parallelising planning via heuristic search to benefit from the increasing availability of large-scale parallel clusters (Kishimoto, Fukunaga, & Botea 2009).

As in the other research groups, our research is increasingly informed by real-world problems. We have worked on military operations planning with the Australian Defense Science Organisation (Aberdeen, Thiébaux, & Zhang 2004). We also contribute to NICTA's effort in transport and logistics; e.g. path-finding and compression of all-pairs shortest path databases (Botea 2011).

Smart energy grids optimising generation, storage, transportation, distribution and consumption of energy, will offer formidable challenges at the intersection of planning, diagnosis, and control. Even conventional grids stretch the limit of existing technology. For instance, faults in such systems often lead to a range of secondary abnormalities, which in turn generate alarm cascades that overwhelm operators. For example, we are currently evaluating the use of DES diagnosis to intelligently process the alarm flows produced by an Australian power transmission utility (Haslum & Grastien 2011).

Software systems

The software systems group looks at software systems across many different scales: from the low level of a microkernel operating system to the high level of cloud based systems. The AI research within this group is focused on automated reasoning and formal methods.

Automated Reasoning

Research on Automated Reasoning in NICTA is concerned with a variety of aspects of mechanizing logical reasoning. We develop push-button technology that can be used standalone or embedded in larger applications. The approach is grounded in basic research and is driven by application in NICTA projects and elsewhere.

The reasoning problems generated by real-world applications are usually non-trivial with respect to size and complexity. Moreover, different applications typically require different logics for domain modelling and different reasoning services. Correspondingly, we consider a variety of logics (propositional, first-order, higher-order) for classical and non-monotonic reasoning, theorem proving and model-finding, and interactive theorem proving, among others. In the following we highlight some of these developments. Naturally, there are overlaps with other areas.

Our work in the area of propositional logic includes the development of algorithms for more efficiently checking the satisfiability of formulas (Huang 2010b), particularly those that arise from real-world applications, and algorithms for compiling formulas into tractable forms, as well as exploiting the compiled structures for recurring similar reasoning tasks (Huang & Darwiche 2007).

In first-order logic, our focus is on instance-based methods, which have been established as viable alternatives to the more traditional resolution-based methods. One of the leading methods, the Model Evolution calculus (Baumgartner & Tinelli 2008), and its implementation have been co-developed at NICTA. Current research is concerned with extensions like including black-box reasoning for specialized background theories (Baumgartner & Tinelli 2011), to better support, e.g., application in software verification.

In higher-order logic, push-button tools are more limited in scope, so research aims to have machines act as *proof assistants*, helping humans prove difficult theorems. In this space, NICTA supports work on the HOL4 interactive theorem-proving system (Slind & Norrish 2008). This open-source system has a long history (starting in the 1980s), and is used around the world.

We develop non-monotonic reasoning techniques based on Defeasible Logics for normative reasoning. To accommodate the reasoning requirements in this domain, we consider extension by time, modal operators, and change management (Governatori & Rotolo 2010a). The framework and methodology currently proposed by NICTA was one of the first formal approaches to business compliance and allows for the most comprehensive and advanced conceptual model of the normative constraints a set of regulations can impose on business processes (Governatori & Rotolo 2010b).

Formal Methods

Formal Methods research in NICTA takes some of the techniques developed in other AI research areas such as static analysis, constraint solving, automated reasoning, satisfiability reasoning, and interactive theorem proving, and applies them to software development, in particular to software verification and quality assurance.

Key research areas are the semantics of programming languages, program verification and refinement calculi, the integration of various automated and interactive reasoning techniques into program analysis and verification frameworks, and scaling these methods to real-world complexity and code size.

Two projects that exemplify NICTA's work in Formal Methods are the Goanna static analysis tool for large industrial C/C++ code bases (Fehnker *et al.* 2007), and the L4.verified project (Klein *et al.* 2009) that provided the first implementation-level mathematical proof of functional correctness for an OS microkernel.

The Goanna tool, developed at NICTA, is now available as a commercial product from the spinout company Red Lizard Software[1]. It employs novel static analysis techniques, combined with model checking and constraint solv-

[1] `http://redlizards.com`

Figure 3: NICTA's stand at a recent CeBit exhibition.

ing to search for common predefined software defects such as buffer overflows or null-pointer dereferences with very low rates of false positives. The properties it searches for are easily customisable; they include memory corruption and leaks, code patterns that point to software quality issues, security vulnerabilities, API rule violations, and coding standards violations. Goanna fully automatically identifies over 100 types of serious defects.

In tune with NICTA's aim of employing basic research to solving real-world problems, the tool integrates tightly and easily into standard industrial development processes. It can be used as a drop-in replacement for the compiler in standard build processes. Integration with IDEs such as VisualStudio and Eclipse is available. Counterexample traces and error positions can easily be replayed within the IDE.

The second example project is the application of machine-checked, interactive proof in the Isabelle/HOL theorem prover to the seL4 microkernel (Klein *et al.* 2009). seL4 is a third-generation high-performance microkernel of the L4 kernel family. It's predecessor technology, the OKL4 kernel, is being marketed by NICTA spinout Open Kernel Labs[2] and at this time deployed in over 1.2 billion devices.

The proof the project developed shows that the C implementation of seL4 correctly implements its high-level functional specification. This is the first time that mathematical proof has successfully been applied to a real OS implementation on the scale of 10,000 lines of code. Microkernels provide fault isolation and security separation to application components. Formal verification provides ultimate assurance of correctness. Together, they enable a new way of building systems that has the potentially to fundamentally increase the assurance we can achieve of complex safety- and security-critical software.

Although it has been known in principle for more than 30 years, that formal proof can be applied to the implementation level, the complexity of real-world code has so far been prohibitive. To undertake this verification, the team has created a detailed formal semantics of the C programming language subset used in seL4, formal refinement and verification techniques in the theorem prover Isabelle/HOL that scale to large code and team size, and an innovative microkernel design and rapid prototyping technology that allowed the Formal Methods and OS teams to work together closely, interleaving kernel design, formal specification, implementation, and proof.

Current research in Formal Methods aims at achieving the same ultimate degree of assurance for systems on the scale of millions of lines of code. This is in principle made possible, not by linearly scaling the previous proof techniques, but by exploiting the verified kernel foundation and by microkernel-based security architectures. The research challenges that NICTA is addressing with defence and industry partners in this direction are in modelling and verifying concurrent applications, in formally, safely composing systems out of untrusted and trusted components, in integrating this method into the development process, and in making it feasible to apply in practice.

[2] `http://ok-labs.com`

Machine learning

The Machine Learning group at NICTA undertakes a wide range of activities, from theory building, modeling and algorithm development, to the use of machine learning in the solution of real-world problems. Much of the work is motivated by applications in domains such as health, document analysis, computer vision, social networking, natural language processing and preference elicitation. A large part of the core machine learning research is dedicated to learning theory, large-scale machine learning, graphical models, topic models, structured prediction and Gaussian processes.

One of the theoretical aims of the group is to better understand how learning problems can be represented and related. We are primarily problem- rather than technique-driven and are interested in, for example, questions of characterising when a problem admits an efficient solution or when one type of problem can be transformed into another. To date, this work has focused on classification, probability estimation, and divergence estimation problems. Many relationships between losses and divergences have been collected in (Reid & Williamson 2011) which have led to new surrogate regret bounds and tight generalisations of Pinsker's inequality. More recently, we have established a new geometric characterisation of which losses allow for quickly decaying regret in multiclass prediction with expert advice problems (van Erven, Reid, & Williamson 2011).

The research in graphical models and structured prediction focuses both on modeling and algorithmics, and applications such as rank estimation (Petterson *et al.* 2009), graph matching (Caetano *et al.* 2009) and multi-label classification (Petterson & Caetano 2010). A recent major achievement was the development of faster algorithms for maximum-a-posteriori inference in discrete graphical models (McAuley & Caetano 2011). Traditional belief-propagation algorithms are designed for the worst-case scenario and do not exploit the structure of the input data in order to make computations more efficient. In that work we presented exact max-product algorithms that have improved expected-case computational complexity under a reasonable assumption on the distribution of the input data. In practice we verify substantial speedups in carrying out tasks that are often modeled as inference in graphical models, such as text denoising, optical flow computation and protein design.

We also research into ways of scaling-up machine learning algorithms to deal with the data deluge arising from modern technologies. For instance, we recently investigated how the stochastic gradient descent algorithm can be implemented on a parallel architecture like a General Purpose Graphical Processing Unit (GPGPU), trading precision on one processing unit for delayed updates and resulting in more parallelism and overall speedup and precision gain (Xiao, McCreath, & Webers 2011). It was shown that the approach is limited by the memory bandwidth between the main processor and GPGPU, which may become less of an issue for future GPUs with more on-board memory.

One of the application focuses of the group is machine learning for structured text analysis and retrieval. This has motivated research in two areas: non parametric methods and topic models. Within topic models, we have developed techniques for structured documents, for instance documents with sections (Du, Buntine, & Jin 2010a) or sequential chapters (Du, Buntine, & Jin 2010b), and also methods for including semantic information like word relatedness, for instance to improve the understandability of topics. Within non-parametric modelling, the work on topic models has lead us to develop new techniques for machine learning with discrete hierarchical models using hierarchical Pitman-Yor distributions.

Another application focus is preference elicitation. This is the task of eliciting preferences from a user in order to make (approximately) optimal decisions or recommendations on behalf of that user. Because the number of potential preferences is very large, it is crucial to optimize preference elicitation queries and their sequence to obtain the best outcome for the user in the fewest queries. Bayesian inference and learning methods are ideal for this task since they provide the crucial probabilistic information required to compute the value of information of a potential query, i.e., the expected gain that would result from having an answer to the proposed query. Along these lines, (Guo & Sanner 2010) have looked at efficient and scalable methods for performing Bayesian preference elicitation and (Bonilla, Guo, & Sanner 2010) has examined more complex Gaussian Process and kernel models that account for shared preference information among multiple users. With this previous work and ongoing work in this area, NICTA researchers are developing advanced preference elicitation methods to build interactive recommendation systems that intelligently adapt to the needs of their users.

Computer Vision

NICTA's computer vision group draws strength from fundamental areas including geometry, recognition and detection, statistical pattern recognition and segmentation, and from approaches such as optimisation and machine learning. Computer Vision problems are often posed in the context of an application, which suits the nature of NICTA's use-inspired fundamental research approach. Some current driving applications are the bionic eye, hyperspectral imaging technologies, vision in road scenes, and visual surveillance.

Bionic Vision Australia (BVA) started in 2010, with the goal of developing a retinal implant to restore vision to people with visual impairment due to retinal degenerative conditions.[3] NICTA is a consortium member, and vision processing based on computer vision is one of its contributions. The consortium will be conducting human implanted trials of an electrode device in 2013, and is developing an implant with 1000 electrodes. In time devices may have higher resolution, however, visual prosthetic devices will always be limited by the residual damage from the cause of blindness. As such, the problem of vision processing to restore key functions of human vision with reduced resolution, and reduced dynamic range, using sets of wearable input cameras of comparatively high resolution. See Figure 4 for an example of the vision processing challenges tackled in this domain.

[3] www.bionicvision.org.au

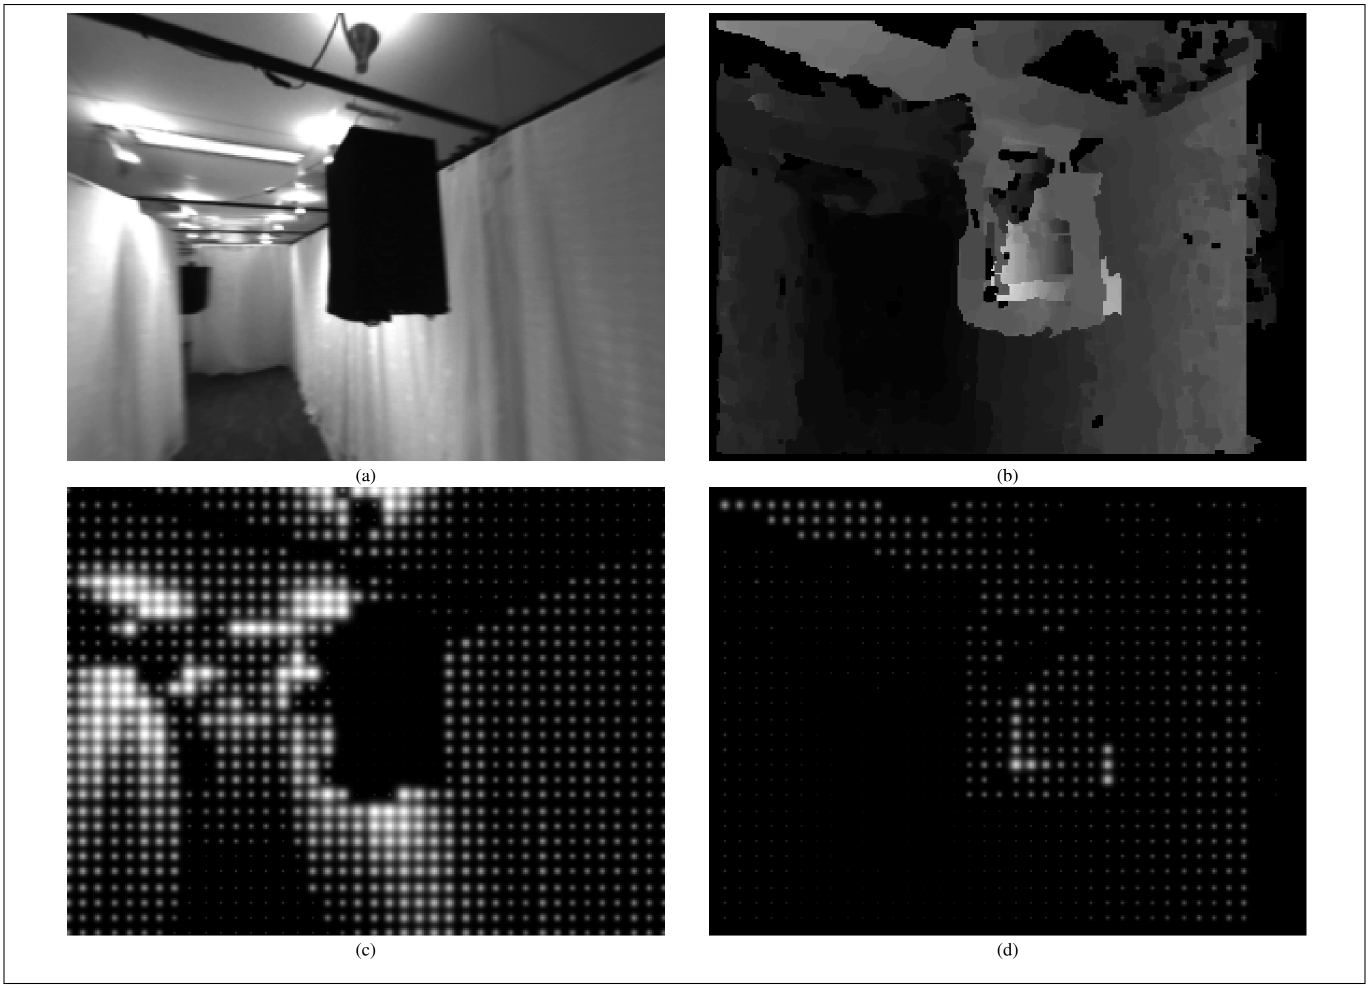

Figure 4: Image processing for the bionic eye. The first image (a) shows a navigation corridor, with an overhanging obstacle. The second image (b) shows the depth map corresponding to this. The third image (c) shows what this looks like if we render it with 30x35 phosphenes coding image intensity as brightness. Phosphenes are what people report seeing as the result of electrical stimulation of the visual system. The structure of the corridor is clear, and you can see the overhanging obstacle, but cannot judge its distance. The last image (d) uses phosphenes to render depth, the closer the brighter. Here you can see the nearby corridor wall fading away in depth, making structure visible, and the obstacle and its depth are apparent.

Key problems have been identified generally in the literature (e.g., focus groups (Keeffe *et al.* 2010)): for example orientation and mobility, and face recognition. Simulations of prosthetic vision can be used with normally sighted participants to refine approaches. The project has conducted orientation and mobility trials that demonstrate the value of providing representations including depth information, when negotiating overhanging obstacles (Barnes *et al.* 2011). New algorithms for robust and rapid detection of free-space and obstacles in disparity data (McCarthy & Barnes 2010) form a basis of new approaches. In low resolution, computer vision approaches can assist with fixation to facilitate high acuity recognition, such as for faces (He, Barnes, & Shen 2011). This approach is underpinned by fundamental research in face detection (Shen, Wang, & Li 2010).

In road scenes, AutoMap has technology to automatically find objects of interest in video. Geo-referenced road video images are searched, for example, for road signs that are important to navigation which are compiled into a map including their precise location. With commercial partners such as Sensis, a leading Australian map provider, NICTA's sign maps are already providing personal navigation information to drivers. Also, for the RTA, NICTA has conducted research around automated pedestrian detection (Shen, Paisitkriangkrai, & Zhang 2011).

The spectral imaging project conducts fundamental research to enable the next generation of hyperspectral cameras. The project has shown that one may simultaneously recover surface shape and photometric invariants from multi-spectral images, which is generally ill-posed (Huynh & Robles-Kelly 2009). In consumer cameras such multispectral approaches would allow, for example, the modification of lighting models from a single image without scene other information, or to recover properties of specific objects for surveillance applications.

Face recognition in low resolution images is important in applications like border control in international airports. Recent work has developed improved methods for image set recognition (Harandi *et al.* 2011), taking advantage of matching a carefully selected subset of images from video, rather than single images. Currently frame selection is based on a novel fast patch-based probabilistic face image quality measure (Wong *et al.* 2011).

NICTA is also conducting theoretically-driven fundamental research, such as camera motion estimation: recovering rotation across multiple images using L1 averaging (Hartley, Aftab, & Trumpf 2011); and, decoupling rotation and translation using antipodal points on hemispherical cameras (Lim, Barnes, & Li 2010). Also, in machine learning approaches, the dual formulation of boosting algorithms, which particularly improve detector performance (Shen & Li 2010).

Other areas

There are a number of other projects in NICTA which develop or exploit AI technologies. Two areas of especial note are bioinformatics, and human computer interaction.

Bioinformatics

There is almost universal agreement that future major advances is medicine and health will be strongly reliant on sophisticated information technology, hence the field of Bioinformatics is burgeoning. Unsurprisingly then NICTA has a significant number and wide variety of Bioinformatics projects underway. NICTA is fortunate in that its Victoria laboratory is situated at the heart of the fourth largest medical research precinct in the world, in the city of Melbourne, where over 10,000 medical researchers can be found with a 5 km radius of the laboratory. NICTA is partnered with many of the world leading medical research institutes that reside in this precinct, as well as other important institutes situated in other parts of Australia. There are a myriad bioinformatics research projects at NICTA, we highlight a few below.

One of the immediate challenges to computing generated by new DNA technology is the problem of processing the huge amounts of data that are generated by todays high throughput sequencing technology. The ”de novo” assembly problem looks how the sequence fragments generated overlap in order to reconstruct the original DNA sequence. The presence of measurement errors and genomic redundancy make this a computationally hard problem. Recent work at NICTA (Conway & Bromage 2011) attacks the problem from the point of view of resource usage – enabling researchers to perform this task with commodity computing rather than expensive supercomputing resources. The high throughput sequencing technology has made gathering the sequence fragments cheap; our technology makes the assembly cheap.

Determining 3D folding structure for proteins is one of the most challenging problems facing moleculare biology, with over 15 million proteins known but less than 100,000 with known structure. Determining proteins structure is a challenging task and NICTA is, with partners, developing technology towards answering this problem. MUSTANG (Konagurthu *et al.* 2006) is a leading tool for multiple structural alignment of proteins. Structural alignment of proteins is a key method for determining candidate phase information for new proteins whose structure is being determined by the molecular replacement method, and MUSTANG has been incorporated in the prediction pipeline toolset (Konagurthu *et al.* 2010) and helped determine the structure of a number of proteins including MACPF. On another front NICTA has helped develop new methods for *ab initio* structure prediction purely from sequence information (Hoque *et al.* 2011)

Current biomedical and genomic research is heavily dependent on biological knowledge, some of which has been human-curated into databases but much of which is only available in the scientific literature. Extracting such knowledge and converting it into a form over which data mining and pattern recognition techniques can be applied is a great challenge, but one with enormous potential benefit. Challenges include highly ambiguous terminology language that often compacts multiple relations, or biological ”events”, into a single phrase; and the use of tabular representations for relationships and other data. NICTA has developed genomic information retrieval methods (Stokes *et al.* 2009) that improve on standard approaches by using concept-based refinement.

Human Computer Interaction

Finally, NICTA has several HCI projects which use and develop AI technologies. Organizations such as hospitals invest considerable amounts in software applications meant to provide easier access to medical records or to better organize the activity inside a hospital department. Previous experience indicates that the installation of new software often has negative, unintended consequences that can significantly diminish the usefulness of the product at hand (Ash, Berg, & Coiera 2004; Littlejohns, Wyatt, & Garvican 2003).

In Prospective ICT Evaluation (PICTE), we develop solutions for anticipating the impact that the installation of a new software application can have on a work environment such as a hospital or a department in a hospital. Being interested in the human factors aspect of the problem, we focus on how the planned changes would affect the staff working in that workplace. The earlier that undesired side effects can be identified, the easier and more cost-effective it is to take corrective action. The results of our research are intended to inform procurement and system acquisition decisions as much as to guide software development.

Previous PICTE research has relied almost exclusively on manual analysis (Sanderson *et al.* 2012) In our current work with Queensland Health, we are now introducing automated methods for prospective ICT evaluation (Botea & Sanderson 2011). We build models that represent work situations and work functions before the planned change (current model) and after the change (projected model). Models can be evaluated using methods developed in areas such as AI planning, model checking, workflows and business process modelling.

The evaluation focuses on the reachability of criteria that are relevant to the staff and their work routines. The main evaluation criteria include costs (for example, time taken, time uncertainty, information quality, mental workload, prospective memory load) and how well professional priorities and values are respected (for example, patient safety, patient throughput, infection control, quality of clinical notes, thoroughness of follow-through). Different professional groups such as doctors, nurses, allied health professionals, administrative officers may not be subject to the same costs or have the same priorities and values (Naikar *et al.* 2003). The differences observed between the evaluation of current and projected models as a result of the technical change let us evaluate the impact of the planned change on different professional groups. Technical changes that allow goals to be reached that are particularly costly for one or more professional groups are undesirable and point to the need for redesign or rearrangement of workplace roles and responsibilities in a shared, negotiated process.

We envisage that analysis will ultimately be performed as a mixed-initiative system, with the human identifying general nature of projected models that can then be tested by the automated reasoning.

Conclusions

AI forms a large part of NICTA's research portfolio. Indeed, AI impacts on almost every activity going on in NICTA. Computer vision algorithms are being developed to improve Australia's Bionic Eye. Optimisation methods are being used to reduce transport costs for logistical operations. Formal methods are being used to prove correct large systems like operating systems. Machine learning algorithms are being tuned to summarize documents. Automated Reasoning methods are being used to identify problems with business rules and violations of otherwise intangible aspects of work practice. The list could go on and on. Next time you are in our part of the world, you are encouraged to stop by and find out more.

Acknowledgements

NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program. Some of this research was also supported in part by Australian Research Council through its Special Research Initiative (SRI) in Bionic Vision Science and Technology grant to Bionic Vision Australia (BVA). A number of the AI projects also receive support from the Asian Office of Aerospace Research and Development.

References

Aberdeen, D.; Thiébaux, S.; and Zhang, L. 2004. Decision-Theoretic Military Operations Planning. In *Proc. ICAPS*, 402–412.

Anbulagan; Pham, D.; Slaney, J.; and Sattar, A. 2005. Old resolution meets modern SLS. In *AAAI*, 354–359.

Ash, J.; Berg, M.; and Coiera, E. 2004. Some unintended consequences of information technology in health care: The nature of patient care information system-related errors. *Journal of the Americal Medical Informatics Association* 11(2):104–112.

Barnes, N.; Lieby, P.; Dennet, H.; McCarthy, C.; Liu, N.; and Walker, J. 2011. Mobility experiments with simulated vision and sensory substitution of depth. In *ARVO*.

Baumgartner, P., and Tinelli, C. 2008. The Model Evolution Calculus as a First-Order DPLL Method. *Artificial Intelligence* 172(4-5):591–632.

Baumgartner, P., and Tinelli, C. 2011. Model evolution with equality modulo built-in theories. In Bjoerner, N., and Sofronie-Stokkermans, V., eds., *CADE-23 – The 23nd International Conference on Automated Deduction*, Lecture Notes in Artificial Intelligence. Springer. To appear.

Bessière, C.; Hebrard, E.; Hnich, B.; and Walsh, T. 2007. The complexity of global constraints. *Constraints* 12(2):239–259.

Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds. 2009. *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.

Bonilla, E.; Guo, S.; and Sanner, S. 2010. Gaussian process preference elicitation. In Lafferty, J.; Williams, C. K. I.; Shawe-Taylor, J.; Zemel, R.; and Culotta, A., eds., *Advances in Neural Information Processing Systems 23*. 262–270.

Botea, A., and Sanderson, P. 2011. Automated reasoning for prospective ICT evaluation. Technical Report NICTA-QRL-PICTE-2011-01, NICTA Queensland Research Laboratory.

Botea, A. 2011. Optimal Pathfinding without Runtime Search. In *Proc. AIIDE*.

Buffet, O., and Aberdeen, D. 2009. The Factored Policy-gradient Planner. *Artif. Intell.* 173(5-6):722–747.

Caetano, T.; McAuley, J. J.; Cheng, L.; Le, Q. V.; and Smola, A. J. 2009. Learning graph matching. *IEEE Trans. on PAMI* 31(6):1048–1058.

Conway, T., and Bromage, A. 2011. Succinct data structures for assembling large genomes. *Bioinformatics* 27(4):479–486.

Du, L.; Buntine, W.; and Jin, H. 2010a. A segmented topic model based on the two-parameter Poisson-Dirichlet process. *Machine Learning Journal* 81:5–19.

Du, L.; Buntine, W.; and Jin, H. 2010b. Sequential latent Dirichlet allocation: Discover underlying topic structures within a document. In *ICDM '10*, 148–157.

Fehnker, A.; Huuck, R.; Jayet, P.; Lussenburg, M.; and Rauch, F. 2007. Model checking software at compile time. In *Proceedings of the First Joint IEEE/IFIP Symposium on Theoretical Aspects of Software Engineering*, 45–56. IEEE Computer Society.

Governatori, G., and Rotolo, A. 2010a. Changing legal systems: legal abrogations and annulments in defeasible logic. *Logic Journal of IGPL* 18(1):157–194.

Governatori, G., and Rotolo, A. 2010b. A conceptually rich model of business process compliance. In Link, S., and Ghose, A., eds., *7th Asia-Pacific Conference on Conceptual Modelling (APCCM 2010)*, volume 110 of *CRPIT*, 3–12. ACS.

Grastien, A.; Anbulagan; Rintanen, J.; and Kelareva, E. 2007. Diagnosis of Discrete-Event Systems Using Satisfiability Algorithms. In *Proc. AAAI*, 305–310.

Guo, S., and Sanner, S. 2010. Real-time multiattribute Bayesian preference elicitation with pairwise comparison queries. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS-10)*, volume 9, 289–296.

Haim, S., and Walsh, T. 2008. Online estimation of SAT solving runtime. In *SAT*, 133–138.

Harandi, M.; Shirazi, S.; Sanderson, C.; and Lovell, B. 2011. Graph embedding discriminant analysis on grassmannian manifolds for improved image set matching. In *CVPR '11 Int Conf on Computer Vision and Pattern Recognition*. IEEE.

Hartley, R.; Aftab, K.; and Trumpf, J. 2011. L1 rotation averaging using the weiszfeld algorithm. In *Proceedings of 24th IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011)*, 3041–3048.

Haslum, P., and Grastien, A. 2011. Diagnosis As Planning: Two Case Studies. In *Proc. 5th Scheduling and Planning Applications Workshop (SPARK)*. To appear.

Haslum, P. 2006. Improving Heuristics Through Relaxed Search - An Analysis of TP4 and HSP*a in the 2004 Planning Competition. *J. Artif. Intell. Res. (JAIR)* 25:233–267.

He, X.; Barnes, N.; and Shen, C. 2011. Face detection and tracking in video to facilitate face recogntion with a visual prosthesis. In *ARVO*.

Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible Abstraction Heuristics for Optimal Sequential Planning. In *Proc. ICAPS*, 176–183.

Hoque, M.; Chetty, M.; Lewis, A.; and Sattar, A. 2011. Twin removal in genetic algorithms for protein structure prediction using low resolution model. *IEEE/ACM Transactions of Computational Biology and Bioinformatics* 8(1):234–245.

Huang, J., and Darwiche, A. 2007. The language of search. *Journal of Artifical Intelligence Research (JAIR)* 29:191–219.

Huang, J. 2007. A case for simple SAT solvers. In *CP*, 839–846.

Huang, J. 2010a. Extended clause learning. *Artif. Intell.* 174(15):1277–1284.

Huang, J. 2010b. Extended clause learning. *Artificial Intelligence* 174(15):1277–1284.

Huynh, C., and Robles-Kelly, A. 2009. Simultaneous photometric invariance and shape recovery. In *ICCV 09*. Springer-Verlag.

Keeffe, J.; Francis, K.; Luu, C.; Barnes, N.; Lamoureaux, E.; and Guymer, R. 2010. Expectations of a visual prosthesis: perspectives from people with impaired vision. In *ARVO*.

Kilby, P., and Verden, A. 2002. Flexible routing combing constraint programming, large neighbourhood search, and feature-based insertion. In *Proceedings of IJCAI-2011 Workshop on Artificial Intelligence and Logistics (AILog)*.

Kilby, P.; Slaney, J. K.; Thiébaux, S.; and Walsh, T. 2006. Estimating search tree size. In *AAAI*.

Kishimoto, A.; Fukunaga, A.; and Botea, A. 2009. Scalable, Parallel Best-First Search for Optimal Sequential Planning. In *Proc. ICAPS*, 201–208.

Klein, G.; Elphinstone, K.; Heiser, G.; Andronick, J.; Cock, D.; Derrin, P.; Elkaduwe, D.; Engelhardt, K.; Kolanski, R.; Norrish, M.; Sewell, T.; Tuch, H.; and Winwood, S. 2009. seL4: Formal verification of an OS kernel. In *Proceedings of the 22nd ACM Symposium on Operating Systems Principles*, 207–220. Big Sky, MT, USA: ACM.

Konagurthu, A.; Whisstock, J.; Stuckey, P.; and Lesk., A. 2006. MUSTANG: A multiple structural alignment algorithm. *Proteins: Structure, Function, and Bioinformatics* 64(3):559–574.

Konagurthu, A.; Reboul, C.; Schmidberger, J.; Irving, J.; Lesk, A.; Stuckey, P.; Whisstock, J.; and Buckle, A. 2010. MUSTANG-MR structural sieving server: Applications in protein structural analysis and crystallography. *Public Library of Science One*. http://www.plosone.org/.

Li, J.; Huang, J.; and Renz, J. 2009. A divide-and-conquer approach for solving interval algebra networks. In *IJCAI*, 572–577.

Lim, J.; Barnes, N.; and Li, H. 2010. Estimating relative camera motion from the antipodal-epipolar constraint. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 32(10):1907–1914.

Little, I.; Aberdeen, D.; and Thiébaux, S. 2005. Prottle: A Probabilistic Temporal Planner. In *Proc. AAAI*, 1181–1186.

Littlejohns, P.; Wyatt, J.; and Garvican, L. 2003. Evaluating computerised health information systems: hard lessons still to be learnt. *British Medical Journal* 326(7394):860–863.

Marriott, K.; Nethercote, N.; Rafeh, R.; Stuckey, P.; de la Banda, M. G.; and Wallace, M. 2008. The design of the zinc modelling language. *Constraints* 13(3):229–267.

McAuley, J. J., and Caetano, T. 2011. Faster algorithms for max-product message-passing. *Journal of Machine Learning Research* 12(4):1349–1388.

McCarthy, C., and Barnes, N. 2010. Surface extraction with iso-disparity contours. In *ACCV Proceedings of the Asian Conference on Computer Vision*.

Naikar, N.; Pearce, B.; Drumm, D.; and Sanderson, P. 2003. Designing teams for first-of-a-kind, complex systems using the initial phases of cognitive work analysis: Case study. *Human Factors* 45(2):202–217.

Nethercote, N.; Stuckey, P.; Becket, R.; Brand, S.; Duck, G.; and Tack, G. 2007. Minizinc: Towards a standard cp modelling language. In Bessiere, C., ed., *Proceedings of 13th International Conference on Principles and Practice of Constraint Programming (CP 2007)*, 529–543. Springer.

Petterson, J., and Caetano, T. 2010. Reverse multi-label learning. In Lafferty, J.; Williams, C. K. I.; Shawe-Taylor, J.; Zemel, R.; and Culotta, A., eds., *Advances in Neural Information Processing Systems 23*. 1912–1920.

Petterson, J.; Caetano, T.; McAuley, J.; and Yu, J. 2009. Exponential family graph matching and ranking. In Bengio, Y.; Schuurmans, D.; Lafferty, J.; Williams, C. K. I.; and Culotta, A., eds., *Advances in Neural Information Processing Systems 22*. 1455–1463.

Pham, D.; Thornton, J.; Sattar, A.; and Ishtaiwi, A. 2005. SAT-Based versus CSP-Baced constraint weighting for satisfiability. In *AAAI*, 455–460.

Pham, D.; Thornton, J.; Gretton, C.; and Sattar, A. 2008. Combining adaptive and dynamic local search for satisfiability. *JSAT* 4(2-4):149–172.

Pham, D.; Thornton, J.; and Sattar, A. 2007. Building structure into local search for SAT. In *IJCAI*, 2359–2364.

Pham, D. N.; Thornton, J.; and Sattar, A. 2008a. Modelling and solving temporal reasoning as propositional satisfiability. *Artif. Intell.* 172(15):1752–1782.

Pham, D.; Thornton, J.; and Sattar, A. 2008b. Efficiently exploiting dependencies in local search for SAT. In *AAAI*, 1476–1478.

Reid, M. D., and Williamson, R. C. 2011. Information, divergence and risk for binary experiments. *Journal of Machine Learning Research* 12:731–817.

Richter, S., and Westphal, M. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *J. Artif. Intell. Res. (JAIR)* 39:127–177.

Rintanen, J. 2010. Heuristics for Planning with SAT. In *Proc. CP*, 414–428.

Robinson, N.; Gretton, C.; Pham, D.; and Sattar, A. 2010. Partial Weighted MaxSAT for Optimal Planning. In *Proc. PRICAI*.

Sanderson, P.; Xiao, T.; Freeman, C.; and Broxham, W. 2012. Envisioning healthcare work: Models for prospective evaluation of new systems. In *Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium (IHI 2012)*.

Sanner, S.; Delgado, K. V.; and de Barros, L. N. 2011. Symbolic Dynamic Programming for Discrete and Continuous State MDPs. In *Proc. UAI*.

Sanner, S. 2010. Relational Dynamic Influence Diagram Language (RDDL): Language Description. `http://users.cecs.anu.edu.au/~ssanner/IPPC\_\2011/RDDL.pdf`.

Selman, B.; Kautz, H.; and McAllester, D. 1997. Ten challenges in propositional reasoning and search. In *IJCAI*, 50–54.

Shen, C., and Li, H. 2010. On the dual formulation of boosting algorithms. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 32(12):1907–1914.

Shen, C.; Paisitkriangkrai, S.; and Zhang, J. 2011. Efficiently learning a detection cascade with sparse eigenvectors. *IEEE Transactions in Image Processing* 20(1):22–35.

Shen, C.; Wang, P.; and Li, H. 2010. Lacboost and fisherboost: Optimally building cascade classifiers. In *European Conference on Computer Vision 2010*, volume 2, 608–621.

Slind, K., and Norrish, M. 2008. A brief overview of HOL4. In Mohamed, O. A.; Muñoz, C.; and Tahar, S., eds., *Theorem Proving in Higher Order Logics, 21st International Conference*, volume 5170 of *Lecture Notes in Computer Science*, 28–32. Springer.

Stokes, N.; Li, Y.; Cavedon, L.; and Zobel, J. 2009. Exploring criteria for successful query expansion in the genomic domain. *Information Retrieval* 12(1):17–50.

Stuckey, P.; de la Banda, M. J. G.; Maher, M.; Marriott, K.; Slaney, J.; Somogyi, Z.; Wallace, M.; and Walsh, T. 2005. The g12 project: Mapping solver independent models to efficient solutions. In *Proceedings of 11th International Conference on Principles and Practice of Constraint Programming (CP 2005)*, 13–16. Springer.

Thornton, J., and Pham, D. 2008. Using cost distributions to guide weight decay in local search for SAT. In *PRICAI*, 405–416.

van Erven, T.; Reid, M. D.; and Williamson, R. C. 2011. Mixability is bayes risk curvature relative to log loss. In *Proceedings of the 24th Annual Conference on Learning Theory*.

Walsh, T. 2000. SAT v CSP. In Dechter, R., ed., *6th International Conference on Principles and Practices of Constraint Programming (CP-2000)*, 441–456. Springer-Verlag.

Walsh, T. 2008. Breaking value symmetry. In Fox, D., and Gomes, C., eds., *Proceedings of the 23rd National Conference on AI*, 1585–1588. Association for Advancement of Artificial Intelligence.

Wong, Y.; Chen, S.; Mau, S.; Sanderson, C.; and Lovell, B. 2011. Patch-based probabilistic image quality assessment for face selection and improved video based face recognition. In *Biometrics Workshop at CVPR*.

Xiao, F.; McCreath, E.; and Webers, C. 2011. Fast on-line statistical learning on a GPGPU. In *AusPDC2011*.

Biographical information

Nick Barnes is Research Group Leader in Computer Vision at NICTA. He leads VIBE project exploring vision processing strategies and techniques within the Bionic Vision Australia partnership.

Peter Baumgartner is a Research Leader in the Software Systems Research Group at NICTA. He works on automated deduction, in particular first-order logic theorem proving, developing applications for software verification, knowledge representation, and business rules/process analysis.

Tiberio Caetano is a Senior Researcher in the Machine Learning Research Group at NICTA. He works on graphical models, structured estimation and matching. He has also published work on computational biology, complex networks and text analysis.

Hugh Durrant-Whyte is the CEO of NICTA. Previously, he led the Australian Centre for Field Robotics (ACFR) and the ARC Centre of Excellence for Autonomous Systems (CAS).

Gerwin Klein leads the Trustworthy Systems project at NICTA. This plans to change the game of software design and implementation, aiming for unprecedented security, safety, reliability and efficiency.

Penelope Sanderson is the Professor of Cognitive Engineering and Human Factors at the University of Queensland. She is a Fellow of the Academy of the Social Sciences in Australia.

Abdul Sattar is the founding Director of the Institute for Integrated and Intelligent Systems, Professor of Computer Science and Artificial Intelligence at Griffith University, and a Research Leader at NICTA in the Optimisation research group.

Peter Stuckey is a Professor in the Department of Computing and Information Systems at The University of Melbourne and Research Leader in the NICTA Victoria Research Laboratory where he leads the G12 project within the Optimisation research group.

Sylvie Thiebaux is Professor in the School of Computer Science at The Australian National University, and Research Leader within the Optimisation research group, working in the area of planning and diagnosis.

Pascal Van Hentenryck is a professor in the Melbourne School of Engineering and the overall leader of the Optimisation research group at NICTA. He is a Fellow of AAAI.

Toby Walsh is a reseach leader at NICTA in the Optimisation research group, adjunct Professor at the Department of Computer Science and Engineering at the University of New South Wales, and external Professor of the Department of Information Science at Uppsala University. He is a Fellow of both AAAI and ECCAI.