REF Outputs

- Most important part of REF assessment is research outputs (typically publications, but can be software, corpora, …)

- Expect a range of 1–6 publications per submitted person with average of 2?, 3?, 4?

- Quality not quantity – Originality, Significance, Rigour

- 100 word statements to support each submitted output

- Edinburgh REF-2014 Informatics submission: http://results.ref.ac.uk/Results/BySubmission/961

- REF-2021 output period: Jan 2014 - Dec 2020

- REF-2021 website: http://www.ref.ac.uk

Examples from REF–2014 submission

Schema-Based Independence Analysis for XML Updates

Originality: A new technique for statically determining when XQuery updates affect the results of queries, applied to view maintenance. Handles a larger set of XQuery features than any previous work on this problem.

Significance: This work (and a companion DBPL 2009 workshop paper XML update semantics) led to a follow-up PVLDB 2010 paper; influenced others (Bidoit et al., Junedi et al.) who have adopted our benchmark. Established large asymptotic savings for static view maintenance.

Rigour: Detailed experimental evaluation based on a new benchmark we developed to exercise features not handled by prior techniques.

Portable compiler optimisation across embedded programs and microarchitectures using machine learning

Originality: First paper to develop a machine-learning based compiler that offers performance portability across a microarchitecture design space.

Significance: This work represents a departure from traditional compilers that rely on fixed heuristics. This is especially important in the context of today's heterogeneous processors found in mobile phones. Industry has expressed a particular interest in this type of technology and we are currently working with ARM to integrate our approach within their compiler.

Rigour: An extensive evaluation of the micro-architectural design space of 35 real-life applications (MiBench embedded benchmarks) was performed. HiPEAC Paper Award recipient.

Quasi-Newton Methods for Markov Chain Monte Carlo

Originality: High-dimensional continuous problems are not handled effectively by existing methods for approximate inference. We introduce the use of quasi-Newton Hessian approximations directly within an approximate inference method.

Significance: We are unaware of other examples of quasi-Newton approximations outside of optimisation. This opens the door for many other new methods for approximate inference. There is very little literature on using second order information within MCMC due to its computational complexity; this work breaks this bottleneck.

Rigour: We outperform state-of-the-art MCMC methods empirically on several real-world data sets. Paper published in the top machine learning conference.

Incremental, Predictive Parsing with Psycholinguistically Motivated Tree-Adjoining Grammar

Originality: Proposed a computational model that accounts for prediction, a key component in human parsing that has not previously been modelled.

Significance: The model unifies two main theories of human parsing: Dependency Locality Theory and Surprisal. They make contradictory assumptions and explain complementary experimental data. Our model unifies them, explaining an unprecedented range of data. The ground-breaking work garnered major prizes for Demberg: Glushko dissertation prize (Cogsci Society), distinguished dissertation runner-up (BCS). Computational Linguistics is the top journal in the eponymous field.

Rigour: Unlike previous psycholinguistic models, our parser is rigorously evaluated on coverage, precision, recall using standard Parseval metrics.

Cooperative Localization for Autonomous Underwater Vehicles

Originality: Presented an algorithm for cooperative AUV navigation using acoustic modems which is robust to packet loss and maintains consistency by continuously enumerating the most likely solution set of the AUV trajectories.

Significance: This publication was the first to deal with the AUV navigation supported by active mobile acoustic beacons. Additionally this work also overviews issues of connectivity and consistency, including demonstrating that a naive application of Kalman filtering is unsuitable in this domain.

Rigour: Published in the most rigorous international robotics journal (IJRR) with support from significant multi-year field experimentation.

Animals Versus Animats: Or Why Not Model the Real Iguana?

Originality: Challenges the wide-spread approach of simulating invented 'animats', arguing that it fails to have the intended relevance to understanding biology.

Significance: Published as a target article with 15 commentaries from the leading researchers in the field of adaptive behaviour simulation. With author's reply, the entire issue of the journal was dedicated to discussion of the argument presented. It develops the author's previous discussions of modelling methodology which have significantly shaped this research field.

Rigour: The argument is grounded in the philosophy of science, and provides a logical justification and detailed critique of counter-arguments.

The arrow calculus

Originality: The first rigorous justification of a widely-used notation for programming with arrows. Published the first typing rules for the notation, and showed nine idiosyncratic rules used previously can be replaced by five rules in two well-known patterns.

Significance: Enabled our work on formlets, which underpin a commercial product (Intellifactory WebSharper) two open-source Haskell libraries, Digestive Functors and Yesod, and libraries for Common Lisp, F#, Javascript, Racket, and Scala, all implemented by third parties.

Rigour: Fifteen pages of proofs in Appendix A.2 of Yallop's PhD thesis. JFP is the leading journal in functional programming.