curriculum vitæ

Wren Gayle Romano[†]

wren@cpan.org
1600 Amphitheatre Pkwy. 000-000-0000 (cell)
Mountain View, CA 94043 *https://wrengr.org*

EDUCATION

Dual-Ph.D. Cognitive Science & Computational Linguistics, Indiana University **ABD**

Thesis *Chiastic λ-calculi and free word order in Japanese*
Advisor Lawrence Moss

M.S.E. Computer Science, Johns Hopkins University **2008**

Project Heterogeneous Typed Unification in Dyna
Advisor Jason Eisner

Graduate study in Computer Science, Portland State University **2005–2007**

Advisor Melanie Mitchell

B.A. Linguistics with Anthropology focus, Reed College **2003**

Thesis *Performative Differences of Gender on the Prosody of Am. English Speech*
Advisor Matt Pearson

EMPLOYMENT

Google LLC & X the moonshot factory, Mountain View, California USA

Software Engineer, Google Research **2019–**
Software Engineer, Foundry, X **2016–2019**

Hardware logic design for a classified project using Haskell, BlueSpec, Rust, and TensorFlow. Despite remaining classified, our project was the subject of the HIW 2019 keynote (delivered by Satnam Singh and Lennart Augustsson), so that much is public knowledge. The dual employment listing is because the project started off at X, but then graduated and was acquired by Google in 2019.

I was one of the original team for the project, so I've worked on a bit of everything. Some particular achievements include: designing novel arithmetic/DSP units, and optimizing their RTL implementations; guiding decisions on the topology (both connectivity and bandwidth) of a network of SIMD-like units; developing algorithms for embedding high-dimensional tensor computations into said topologies (patent pending), and providing analysis of said algorithms, with explicit controls for adjusting the tradeoff between RAM bandwidth, communication bandwidth, and throughput/

[†]Prior to 2014 published under the surname Thornton.

latency; providing analyses of particular TensorFlow models, both to inform clients of the performance/efficiency of their models under the aforementioned embedding, and to guide compiler developers by discovering novel operations and explicating the specific difficulties those operations raise for the current compiler design; and, finally, designing quantization schemes for representing real numbers with integers rather than floating point.

Software Engineer, Bugs−−, Chrome Security, Google **2016–2017**

Project ClusterFuzz/OSS-Fuzz

Developing infrastructure for using fuzzing as part of a continuous integration process. My particular project was: given stacktrace information from a new crash report, and access to the repository's history, how can we determine which commit/patch introduced the regression so that we can assign the bug to the appropriate developer/team.

McMaster University, Hamilton, Ontario Canada

Doctoral Fellow **Summer 2011**

Topic Embedded Domain Specific Languages in Haskell

Mentor Christopher Anand

Following the *McMaster Workshop on Domain Specific Languages*, I stayed on to advise graduate students on writing domain specific languages in Haskell. The students' projects ranged from an interactive textbook with 3D diagrams for teaching electromagnetism, to using Haskell's type system to statically enforce dimensional analysis in the software for magnetic resonance imaging.

Center for Language & Speech Processing at Johns Hopkins University

Research Programmer **2008–2009**

Project Joshua

Mentor Chris Callison-Burch

Architectural and API design for the Joshua machine translation system to improve extensibility, modularity, scalability, and robustness. In addition to work on Joshua itself, I was one of the researchers for Johns Hopkins' participation in DARPA's GALE project for large-scale machine translation. My work on GALE included maintaining the end-to-end translation pipeline and running experiments to tune parameters and configuration of the pipeline.

Free Geek, Portland, Oregon USA

Chair of the Board **2006–2007**

Board of Directors **2005–2007**

Council Member **2004–2007**

Legal and fiscal responsibility for a 501(c)(3) non-profit community organization that recycles used technology to provide computers, education, internet access, and job skills training to those in need in exchange for community service.

In my time on the board, total income increased 57% to $748,111 (28% of the increase went to net profit), e-waste recycling increased 32% to 550 tons, and in 2007 26% of donated systems and CRTs were put into reuse. I also helped rewrite the bylaws and organized with sister organizations around the country as we transitioned to become a national organization.

Administrator of Systems and Security **2004–2007**

Build Instructor **2004–2007**

Teaching volunteers how to build computers from parts, install GNU/Linux, and use the commandline. Setting up testing stations for incoming hardware. System administration for a small network of Linux servers and a diskless terminal lab.

The Computer Action Team at Portland State University

CLAW-1 (System administrator) **2006–2007**

Desk Cat (Tech support volunteer) **2005–2006**

System administration for a large complex network running Linux, Solaris, and Windows providing the computing resources for about 3,000 users for the Maseeh College of Engineering and Computer Science at Portland State University.

Work included converting the network from NIS to LDAP, maintaining tools for user account handling, teaching new volunteers basic and advanced Perl, and monitoring labs.

TEACHING EXPERIENCE

Graduate Associate Instructor, Indiana University

Graduate Courses

B522 Programming Language Foundations **Spring 2015**

Q520 Mathematics and Logic for Cognitive Science **Spring 2012**

Undergraduate Courses

Q320 Computation in Cognitive and Information Sciences **Spring 2013**

Q260 Programming for Cognitive and Information Sciences **Spring 2013**

Q101 Introduction to Cognitive Science **Fall 2012**

Graduate Instructor, Indiana University

Undergraduate Courses

Q250 Math and Logic for Cognitive and Information Sciences **Fall 2011**

PUBLICATIONS

P. Narayanan, J. Carette, W. Romano, C.-C. Shan, and R. Zinkov **Probabilistic inference by program transformation in Hakaru (system description).** In *Symposium on Functional and Logic Programming (FLOPS)*, March 2016, Kochi, Japan. pp. 62–79. http://wrengr.org/pubs/hakaru_flops2016.pdf

J. Carette, P. Narayanan, W. Romano, C.-C. Shan, and R. Zinkov **Building Blocks for Exact and Approximate Inference.** In *NIPS Workshop on Black Box Learning and Inference*, December 2015, Montréal, Canada. http://wrengr.org/pubs/blackbox_nips15.pdf

Z. Li et al. **Joshua 2.0: A Toolkit for Parsing-Based Machine Translation with Syntax, Semirings, Discriminative Training, and Other Goodies.** In *Proc. Joint 5th Workshop on Statistical Machine Translation (WMT) and Metrics MATR*, July 2010, Uppsala, Sweden. https://www.aclweb.org/anthology/W10-1718.pdf

A. Irvine, M. Kayser, Z. Li, W. Thornton, and C. Callison-Burch. **Integrating Output from Specialized Modules in Machine Translation: Transliterations in Joshua.** *Prague Bulletin of Mathematical Linguistics*, 93, 2010. pp. 107–116. http://wrengr.org/pubs/joshua_mt-marathon2010transliteration.pdf

Z. Li et al. **Joshua: An Open Source Toolkit for Parsing-based Machine Translation.** In *Proc. 4th Workshop on Statistical Machine Translation (WMT)*, March 2009, Athens, Greece. pp. 135–139. https://www.aclweb.org/anthology/W09-0424.pdf

Z. Li, C. Callison-Burch, S. Khudanpur, and W. Thornton. **Decoding in Joshua: Open Source Parsing-based Machine Translation.** *Prague Bulletin of Mathematical Linguistics*, 91, January 2009. pp. 47–56. http://wrengr.org/pubs/joshua_mt-marathon2009.pdf

W. Thornton. **Typed Unification in Dyna: An Exploration of the Design Space.** Masters Project Report, Johns Hopkins University, 2008. 39 pages.

W. Thornton. **Heterogeneous Strategies for Unification: Variable–Value Ordering and Optimized Structures.** Technical Report, Johns Hopkins University, 2008. 14 pages.

W. Thornton. **Performative Differences of Gender on the Prosody of American English Speech.** Undergraduate Thesis, Reed College, 2003. 59 pages.

PRESENTATIONS

W. Thornton. **Chiastic Lambda-Calculi.** At *Natural Language and Computer Science (NLCS) workshop*, 28 June 2013, New Orleans, Louisiana, USA. http://wrengr.org/pubs/chiastic_nlcs2013.pdf

W. Thornton. **Probability Smoothing for NLP: A case study for functional programming and little languages.** At *AMMCS minisymposium on Progress and Prospects in Model-Based Scientific Software Development (SS-SSD)*, 27 July 2011, Waterloo, Ontario, Canada. http://wrengr.org/pubs/smoothing_ammcs2011.pdf

W. Thornton. **Constrained scrambling in CCG: A case study in Japanese.** At *4th NASSLLI, student session*, 26 June 2010, Bloomington, Indiana, USA. http://wrengr.org/pubs/ccgjp_nasslli2010.pdf

INVITED LECTURES

McMaster Workshop on Domain Specific Languages, McMaster University, 13–15 June 2011. http://www.cas.mcmaster.ca/~anand/DSL2011.html

DEPARTMENTAL LECTURES

Indiana University

"Chiastic Lambda-Calculi." *PL Wonks*, 13 March. **2015**
"Computational linguists should interact with other computer scientists or: what I did last summer." *CLingDing*, 13 February. **2015**
"Chiastic Lambda-Calculi." *Logic Seminar*, 4 November. **2014**
"A very brief introduction to Haskell." *CLingDing*, 11 February. **2013**
"Extending CCG for freer word order languages." *CL Lunch*, 30 March. **2010**
"Using Joshua." *CL Lunch*, 13 October. **2009**

RESEARCH EXPERIENCE

Hakaru (https://github.com/hakaru-dev/hakaru)

Graduate Research Assistant, Indiana University **2015–2016**

A probabilistic programming language implemented in Haskell. Unlike other PPLs, Hakaru presents inference procedures as source-to-source program transformations. These transformations include conditioning (by disintegration), simplification of integrals (via Maple), expectation, normalization, and sampling.

Built-To-Order BLAS

Graduate Research Assistant, Indiana University **2014**

A compiler for a domain-specific language for linear algebra programs, seeking to automatically generate efficient BLAS-like operators without relying on the BLAS API itself. Prior work focused on low-level optimizations

like loop fusion. My work enhances the compiler by adding high-level optimizations based on algebraic equalities.

Posta **2010–2014**

Qualifying project for Ph.D. Cognitive Science

Graduate Research Assistant, Indiana University **2010–2011, 2013–2014**

A Haskell library for HMM-based part-of-speech tagging and CCG supertagging. The current standard for parsing and machine translation involves a pipeline of many different systems, which makes it difficult or impossible for the systems to communicate beyond conveying their basic results. A better architecture would allow two-way communication between each of the systems so that higher-level processes could serve as oracles to guide the decisions of lover-level processes. The focus on producing command-line executable programs instead of reusable and robust libraries is partly to blame for the popularity of pipeline models. The goal of Posta is to provide a library which focuses on real-time, on-line, interactive, and any-time algorithms for tagging in order to facilitate the interactive communication we desire in complete parsing and translation systems.

HLTDI L^3 (http://code.google.com/p/hltdi-l3) **2009–2010**

A constraint-based solver written in Python for parsing, generating, and translating with Extensible Dependency Grammar (XDG). Within the current explosion in the quantity of information and in the means to access it, much of the world has been left behind because the information is not in a language that they understand. The L^3 project ("Learning Lots of Languages") has the long-term goal of developing a system to translate to and from many under-represented languages of the Global South and (less ambitiously) of creating tools to be used in information retrieval and computer-assisted language learning with these languages.

Joshua (https://cwiki.apache.org/confluence/display/JOSHUA/)

Research Programmer, Johns Hopkins University **2008–2009**

A scalable parsing-based statistical machine translation system written in Java whose goal is to create a coherent and configurable end-to-end machine translation system (as opposed to current state-of-the-art systems which connect diverse components with a morass of shell scripts, making it difficult to reproduce experimental results).

Dyna 2 (https://github.com/nwf/dyna)

Research project for M.S.E. Computer Science **2008–2009**

Weighted-logic programming language for natural language processing and dynamic programming. Major features include strong typing (including non-linear and recursive types), modules for separate proof universes, and

priority-based mixed forward-/backward-chaining inference (with memoization and truth maintenance). The first version of Dyna was a proof-of-concept for an untyped, single-module, forward-chaining fragment of the language. The second version is rewritten in Haskell and intends to cover all the major features.

FELLOWSHIPS AND AWARDS

Graduate Fellowship **2009–2010**

Tuition remission, health insurance, and stipend. ($30,000)

Scholarship for work on Dyna **2008**

Supported in part by the National Science Foundation under grant No. 0347822 to Jason Eisner. ($16,600)

Google Summer of Code **2005**

Project Paperboy WDM, an XML-based web-development framework.

Mentor Bart Massey, Portland State University

Designed, implemented, and documented a complete functional free/open-source project in under eleven weeks. ($5,000)

ACADEMIC SERVICE

Mentor for SIGPLAN-M, 2021.
Program Committee Member for HIW 2017.
Program Committee Member for OBT 2017.

December 10, 2021