

Kathryn Mazaitis: Curriculum Vitae

kathryn.mazaitis@gmail.com

<http://www.cs.cmu.edu/~krivard>

<https://github.com/krivard>

Profile

I develop software systems for information retrieval and machine learning, on both frontend and backend, web and command line.

I transform piles of haphazard, deadline-oriented code into hard-wearing, scalable, re-usable workhorses.

I have a keen awareness of visual and functional design principles, and how they can serve both the programmer and end user.

I am a champion collaborator, mentor, and worker bee.

I am team-driven, not topic-driven; I'll learn anything to work on complex problems with great people, and I will cheerfully share my knowledge and know-how.

I have a terrible memory and take excellent notes.

I solve problems with communication and iteration: I make diagrams when I can't figure out how to organize code, I make data visualizations when I'm pinpointing a bug, I write Stack Overflow questions when I can't figure out a setting. These are all artifacts designed to communicate, and it's through these acts of communication that I most easily solve problems — often answering my own questions in the making of them.

Keywords

Languages	Development & Project Management	Web frameworks	Statistics & Data Visualization
Java	Eclipse	Tomcat	R
Python	ant	Solr+Jetty	d3
HTML/CSS	GNU make	Struts	OmniGraffle
JavaScript	git	Databases	Processing
Shell Scripting	Trello	postgres	Interface Prototyping
Perl	MediaWiki	MySQL	Paper/Physical
LaTeX	SVN	SQLite?	Revolution
MATLAB	Trac	IR Systems	Flex?
PHP, JSP	CVS	Lucene	CogTool
C	Testing & QA	Solr	Operating Systems
Lisp	JUnit	Indri	Ubuntu linux
Prolog	Nose?	neo4j?	Fedora Core linux
JS frameworks	Selenium		Mac OS X
JQuery			Windows?
Knockout			
dojo/dijit			
Prototype			

Employment

Research Programmer, Machine Learning Department, CMU 2009-present

Supervisor William Cohen

Projects RollMe, TensorLog, GuineaPig, GNAT, ProPPR, SEAL, NELL components, Secondstring, GHIRL, Querendipity

William's group researches machine learning and information retrieval, with a focus on graph-walking algorithms.

My role spans everything from development of research software, documentation, and tutorials, to running experiments, to system administration and student project management.

Since 2009 I have worked on design and implementation of web frontends (Java, Tomcat, JSP, Struts, PHP/MySQL, HTML/CSS/JavaScript) from scratch and building on existing code; Java development on projects varying in size from <10 to 1000+ classes; API design; Java ports of Python prototypes;

command-line interface design and implementation; dataset building with web scraping tools, shell scripts, and Makefiles; speed, memory, and accuracy performance analyses; Amazon Mechanical Turk experiment design, execution, and analysis; and system administration for Ubuntu, Amazon Web Services, GPUs, network storage, and the lab's team account on GitHub.

Research Programmer/Analyst, Institute for Software Research, CMU 2007-2009

Supervisor Anthony Tomasic

Projects Mixer, VIO, Forms, Workflow by Example, Honeydew

Anthony's group researches the use of machine learning in office tasks such as scheduling meetings, managing contacts, and repetitive data gathering.

My work included UI prototyping using paper, works-like, and Wizard-of-Oz (WOZ) techniques, a WOZ usability study and associated data analysis, JavaScript development, Firefox plugin development, and collaborative software design.

Software Technical Lead, Vision Agribotics/ROCONA 2006

Supervisors David Barrett, Bryan L. Aivazian

ROCONA is a high-tech agricultural startup company in California which sponsored an Olin Senior Capstone project (SCOPE). In SCOPE, teams of students work with representatives from industry to complete a year-long study of a real engineering problem. Our team was tasked with the conversion of a Kubota tractor into a prototype robotic vehicle to navigate orange orchards.

As the software tech lead, I was active in the design and assembly of the sensor and safety systems as well as being the primary programmer for the project. Over the summer, I stayed on to transition to the next year's group of students and continue software tasks, including the development of a SQL-based logging system for the robot.

Spring Research, Cognition Lab, Olin College 2006

Supervisor Lynn Andrea Stein

Project Fresnel

Fresnel is a display vocabulary for RDF, much like XSLT is for XML. The Cognition Lab frequently works with RDF for its research in the Semantic Web, and until this project wrote custom parser/display programs for each application.

I worked with another student to develop a Python implementation of Fresnel to allow RDF to be used with current web tools with the same interface as existing backend data stores.

Website <http://sourceforge.net/projects/tobacconist/>

Summer Research, Intelligent Vehicles Lab, Olin College 2005

Supervisor David Barrett

Project Mini-ATV Conversion

This project employed a multidisciplinary student team to design, fabricate, and manufacture the systems necessary to convert a child-sized ATV into a robotic vehicle platform. The software team started completely from scratch, and produced a system which used a stereo camera to target and follow the movements of nearby obstacles using a servo-controlled laser pointer.

My work focused on the central AI, which processes data from other software agents and supplies the motor controller with commands for moving the vehicle. I worked in MATLAB, C, and Python to develop tools for networking and manipulating new data structures as well as the AI code.

Summer Research, Cognition Lab, Olin College 2004

Supervisor Lynn Andrea Stein

Project Course Development, *Foundations of CS*

This new course combining traditional courses in Algorithms, Data Structures, Complexity, Programming Languages, and Automata was first taught in fall of 2004.

I was in charge of assignment prototyping, collecting and evaluating resources, and bringing up course infrastructure (webserver, website, wiki, code development space), in addition to providing input regarding course content and organization.

Projects

RollMe

With William Cohen, Vidhisha Balachandran

RollMe is an organization and planning scheme for machine learning pipelines and research groups. It takes the form of an ecosystem of abstract task definitions, each with a train/test set, with the aim of methodically decomposing tasks into subtasks, swapping out different implementations, tracking task completion and performance, and building pipelines given a contiguous sequence of task implementations.

My contributions: Participated in team design sessions and built a domain model. Adapted a COTS CMS (Bolt) to support the RollMe object type hierarchy (twig, PHP). Wrote an interactive webapp to visualize the active task space and the wishlist task space in the live system, with hooks for starting unfinished tasks (javaScript, d3, Knockout). Established an in-house instance of the CodaLab cloud computation and research tracking software (<https://worksheets.codalab.org>, Stanford), added support for authless SMTP (Python), and fixed errors in the configuration templates (NGINX). Submitted pull requests and bugs as appropriate. Wrote task definition copy for a small sample of past lab projects. Partially converted a past student project to a series of CodaLab worksheets.

Website <http://www.rollme.ml.cmu.edu>

Visualizer <http://www.rollme.ml.cmu.edu/browser/>

TensorLog

With William Cohen, Fan Yang

TensorLog is a Python package which uses factor graphs to extend the advances made by ProPPR into the deep-learning space.

My contributions: Standardized a template for experiments on new datasets, and maintained old datasets in working order. Added support for learning multiple predicates simultaneously. Wrote predictions serializer to be compatible with existing analysis and visualization tools. Wrote a cross-compiler for theano, including a polymorphic Learner rig to re-use existing experiment structures. Added support for `argmax`. Adapted a student research project in tensorflow to use TensorLog's database- and type system logic. Served as the git oracle.

GitHub <https://github.com/TeamCohen/TensorLog>

Papers William W. Cohen, Fan Yang, and Kathryn Mazaitis. Tensorlog: Deep Learning Meets Probabilistic Databases, arxiv.org 2017

More Info William Cohen. Tensorlog: A Differentiable Deductive Database, arxiv.org 2016

InMind Movie Recommendations

With Rose Catherine Kanjirathinkal and William Cohen

InMind is a dialogue-based digital assistant project funded by Yahoo!. The movie recommender system is a RESTful Tomcat webservice running ProPPR.

My contributions: Wrote a data extraction, parsing, and pruning system (GNU make, GuineaPig <https://github.com/TeamCohen/GuineaPig>) to continually update our dataset, based on IMDB's (pre-Jan2017) static data exports and (Jan2017+) Amazon S3. Stepped in to make emergency major updates to the recommendations web service for a demo while the primary developer was unexpectedly unavailable. Made the necessary changes to handle JSON queries instead of HTML forms, and model data reusably (JSP, Java). Added support for human-readable movie names instead of numeric IDs. Added a new query

endpoint to include explanatory information for each recommendation, and extended the ProPPR program to provide the supporting data.

Paper Rose Catherine, Kathryn Mazaitis, Maxine Eskenazi, William W. Cohen. Explainable Entity-based Recommendations with Knowledge Graphs (poster paper), RecSys 2017.

QUASAR: Two datasets for question-answering and reading

With Bhuvan Dhingra and William Cohen

QUASAR presents two new large-scale datasets aimed at evaluating systems designed to comprehend a natural language query and extract its answer from a large corpus of text.

My contributions: Developed a prototype system to build simple relevant context documents for each question, forming the "(S)earch" phase of QUASAR. The prototype was composed primarily of existing information retrieval and natural language processing software, filling in the gaps with Java and GNU make. The rough pipeline:

1. Search ClueWeb (<http://lemurproject.org/cluwebo9.php/>) for the question text,
2. Parse the plain text out of the resulting HTML,
3. Optionally split the text into sentences (Stanford CoreNLP),
4. Load the data for each question into a lucene index (<https://lucene.apache.org/>),
5. Search the index for the question text,
6. Assemble the top scoring items and their noun phrases into a context document and candidate solutions list (Stanford CoreNLP),
7. Partition the resulting query records by complexity and ground truth (i.e. whether the context actually included the answer to the question after all that) (awk, perl).

The partitions were then handed off to a graduate student for analysis in several baseline reading systems.

For the trivia set, I cleaned a collection of 50,000 trivia questions originally designed for "pub quiz"-style tournaments. For the software-domain set, I converted 25,000 definitional sentences of software concepts (from StackOverflow; see GNAT project below) into 37,000 cloze-style questions, for which the search corpus was the top 50 StackOverflow threads for the head tag of the definition. Converted a sample of query records for each set into a webapp for a human labeling task (python wsgi, Solr) and provided a set of labels.

Paper Bhuvan Dhingra, Kathryn Mazaitis, William W. Cohen. Quasar: Datasets for Question Answering by Search and Reading, arXiv 2017.

DIEJOB: Inference and IE in the Biomedical Domain — a GNAT system

With Lidong Bing, Bhuvan Dhingra, Jong Hyuk-Park, and William Cohen

DIEJOB is a system for extracting biomedical category and relation facts from entity-centric corpora, targeted at large unstructured corpora in particular. For drugs, the target corpus is DailyMed, published by NLM. For diseases, the target corpus is WikiDisease.

My contributions: Wrote scrapers (scrapy) and parsers (xml) for several small well-structured corpora to use as distant labels. Advised on the tuning of ProPPR hyperparameters to produce stable results.

Website <http://curtis.ml.cmu.edu/gnat/biomed/>

Paper Lidong Bing, Bhuwan Dhingra, Kathryn Mazaitis, Jong Hyuk Park, William W. Cohen. Bootstrapping Distantly Supervised IE using Joint Learning and Small Well-structured Corpora, AAAI 2017

Software Development Knowledge Base — a GNAT dataset

With William Cohen, Eli Whitney, Lam Wing Chan, Joseph Gibli

Our software KB was developed as a summer research project, with a team of undergrads who had taken an introductory course in machine learning, but were not computer science majors. The task was to explore machine learning tasks in the context of software development entities such as form the set of tags on popular programming advice site StackOverflow (<https://stackoverflow.com>). The students generated entity labels over open-domain categories, built an ontology based on their labels, wrote a simple classifier in ProPPR, compared it to baselines, explored options for NLP features to improve performance, and ran an inter-annotator-agreement study.

My contributions: Ran daily stand-up meetings, and gave impromptu tutorials on machine learning and ML software concepts. Designed, generated, ran, and analyzed labeling activities. Guided the ontology-building process. Scraped, queried, parsed, and organized raw data from StackOverflow and DBPedia, including noncompliant CSV data (scrapy, GNU make). Tried six ideas that failed to align the StackOverflow KB with DBPedia using string-similarity metrics; the seventh yielded acceptable precision, if low recall.

Website <http://curtis.ml.cmu.edu/gnat/software/>

GNAT: a Grounded NELL-like AKBC Toolkit

With William Cohen

GNAT aimed to collect the lab's accumulated knowledge and practice surrounding the analysis of entity-centric text corpora. Spin-off projects include the biomedical domain and software domain projects described above, as well as a smaller geo/venues project combining tweets and Yelp reviews. While GNAT never matured into a complete toolkit, it did give us the opportunity to distinguish common tasks from tasks which seemed similar in theory but in practice required individualized handling.

My contributions: Developed a build template for generating quick KB browsers using Solr (GNU make, json, xml). Wrote several small Java and Python programs for common dataset handling tasks, including soft joins, baseline entity linking, and normalization of hand-labeled data. Designed and wrote the website.

Website <http://curtis.ml.cmu.edu/gnat/>

GitHub <https://github.com/TeamCohen/GNAT>

ProPPR: Programming with Personalized Page Rank

With William Cohen, William Wang

ProPPR is a Java package which uses random-walk algorithms to make inferences over logic program proof trees, permitting higher recall without sacrificing accuracy. It was our first major software release on GitHub during my tenure with the lab.

My contributions: Ported the initial Python prototype to Java. Parallelized the stochastic gradient descent algorithm and the inference task. Generated modular architectures for learning and configuration. Developed a data format that was relatively memory- and disk-efficient while remaining human-readable. Generated multiple scaffolds for running experiments. Assisted grad students with writing research variations of learning and proving components. Ported a new proving paradigm based on the Warren Abstract Machine from Python to Java for version 2.0. Revised the entire system to better suit observed needs for variations, primarily in loss functions and regularization. Wrote an abstract scaffold on top of `java.util.concurrent` to permit anyone to multithread arbitrary code having distinct inputs and outputs. Developed a method and supporting code for hybrid models with some parameters trained by ProPPR and some trained by a (theoretically arbitrary) theano model; this would eventually inform the TensorLog cross-compiler system. Used profiling tools to optimize bottlenecks. Developed a set of sample evaluation datasets. Wrote API documentation, READMEs, diagrams, getting started guides, and walkthroughs. Ran scalability studies and reported results.

GitHub <https://github.com/TeamCohen/ProPPR>

Papers William Yang Wang, Kathryn Mazaitis, and William W. Cohen. Joint Information Extraction and Reasoning: A Scalable Statistical Relational Learning Approach, ACL 2015.

William Yang Wang, Kathryn Mazaitis, and William W. Cohen. A Soft Version of Predicate Invention Based on Structured Sparsity, IJCAI 2015.

William Yang Wang, Kathryn Mazaitis, Ni Lao, Tom Mitchell, and William W. Cohen. Efficient Inference and Learning in a Large Knowledge Base: Reasoning with Extracted Information using a Locally Groundable First-Order Probabilistic Logic, Machine Learning, 2015

William Yang Wang, Kathryn Mazaitis, and William W. Cohen. Structure Learning via Parameter Learning, CIKM 2014

William Yang Wang, Lingpeng Kong, Kathryn Mazaitis, and William W. Cohen. Dependency Parsing for Weibo: An Efficient Probabilistic Logic Programming Approach, EMNLP 2014

William Yang Wang, Kathryn Mazaitis, William W. Cohen. Programming with Personalized PageRank: A Locally Groundable First-Order Probabilistic Logic, CIKM 2013 (Honorable Mention)

Entity Linking with ProPPR

With William Cohen, Richard C. Wang, Frank Lin, Bhavana Dalvi, Jakob Bauer

Our entry for the NIST TAC Knowledge Base Population (KBP2014) Entity Linking Track uses ProPPR to identify which KB entity (identified by e.g. its Wikipedia URL) is being referenced in a particular mention in a document.

My contributions: Wrote and tuned ProPPR program. Built associated dataset from multiple sources, including noncompliant XML. Wrote flexible (if gnarly, 400-line+) Makefile to support debugging and avoid recomputing data over files 1-12GB in size. Coordinated work with grad students to accomplish several variations on nil clustering. Coordinated with offsite grad alumni team to provide an additional submission based on a 2014 export of Wikipedia (Java, HttpClient). Wrote paper in collaboration with team.

Paper A Tale of Two Entity Linking and Discovery Systems, KBP 2014

SEAL: a Set Expander for Any Language

With William Cohen and Richard C. Wang

SEAL is a Java package which finds more examples of known categories and relations by identifying tabular information in web pages.

My contributions: Converted from delimited-string binary relations to architecturally supported binary relations. Added new searcher modules for ClueWeb, the Google Custom API, Bing, and an internal batch query relay. Performed accuracy and speed evaluation (R) with parameter sweeps for searcher-by-searcher tuning.

GitHub <https://github.com/TeamCohen/SEAL>

NELL: a Never-Ending Language Learner

With Tom Mitchell, William Cohen, Bryan Kisiel, Burr Settles, Partha Talukdar, Richard C. Wang, Jamie Callan, Dana Movshovitz, Tom Shen, Estevam Hruschka, Jr., and many others

NELL is a software system which has been running continuously since January 12, 2010, iteratively analyzing web text to build a knowledge base of categories and relations. It utilizes multiple independent sources of information to make highly confident predictions, achieving over 90% precision in most of the ontology.

My contributions: (1) SEAL component; inherited from Richard C. Wang. Rewrote the SEAL component to use the SEAL library directly instead of pinging a separate web service, including a custom stream sampler to make an accurate random selection of seed search terms from sources too big to fit in memory. (2) LatLong tinkertoy; with Tom Shen; portions inherited from Dana Movshovitz. Wrote a geolocation component to estimate latitude and longitude of city entities by interpolating known landmarks. Coordinated work with undergrad researcher to handle ambiguous placenames using a mixture-of-gaussians model. Added a component to the NELL web interface to display map locations where applicable (JSP, javaScript).

Website <http://rtw.ml.cmu.edu/rtw/>

- Map* <http://rtw.ml.cmu.edu/rtw/kbbrowser/map:building>
- Papers* T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, B. Yang, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. La, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, J. Welling. Never-Ending Learning, CACM 2017
- T. Mitchell, W. Cohen, E. Hruscha, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohammad, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, J. Welling. Never-Ending Learning, AAI 2015
- More info* Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. Toward an Architecture for Never-Ending Language Learning, AAI 2010

SCHNELL: Slot-fillers with Classifiers, Heuristics, and NELL

With William Cohen, William Wang, Malcolm Greaves, Guanyu Wang

Our entry for the NIST TAC Knowledge Base Population (KBP2013) Slot-Filling Track uses NELL data as distant labels to answer a limited set of questions (“fill slots”) about entities in a corpus. This was a heavily collaborative effort; everyone rotated though all parts of the code.

My contributions: Built a person-class relation dataset from dbpedia exports, built a location feature using exact matches into geonames, incorporated Stanford CoreNLP’s date/time extractor, revised serialization of concepts to better avoid recomputing data over large files, developed a Solr-based webapp for browsing results and performing error analysis, coordinated the team for several weeks in William’s absence, kept track of meetings and assignments.

Creativity Education

Independent study supervised by Haakon Faste

This work was inspired by an interaction design course for my Masters studies. I was frustrated by the way the rubrics and other artifacts seemed to work against the course goals, and started a conversation with the professor that turned into an independent study. I spent a semester working with the professor to investigate general principles of instruction for creative fields, with particular focus on self-awareness and formative failure. I designed and performed interviews and card-sorting activities, collected resources, iteratively synthesized our observations into a cohesive model, and wrote and published a paper on our findings.

- Paper* Kathryn Mazaitis and Haakon Faste. How learning works and design education: educating for creative awareness through formative reflexivity, DIS 2012 (Honorable mention)

Querendipity

With William Cohen, Ni Lao, Ramnath Balasubramanian, John Wolford, Jelena Jakovljevic; inherited as NIES from Andrew Arnold

Querendipity is a web application which uses graph search on a gene-aware corpus of biology publications to find serendipitous papers outside the normal reading list of a researcher.

My contributions: Rewrote Andrew's wrapper-based prototype as a fully-fledged research web application in Tomcat, supporting live changes to the dataset and algorithm hyper-parameters, hot-swappable algorithm variations, user profiles, and relevance judgements. Wrote an additional component to permit visual gene browsing using an open-source tool. Designed and ran several relevance judgement experiments with our biologist collaborators, compiled and analyzed the data, and created a visualization to help interpret ranking differences between three variations.

Paper Ramnath Balasubramanian, Kathryn Rivard, William W. Cohen, Jelena Jakovljevic and John Woolford. Evaluating Joint Modeling of Yeast Biology Literature and Protein-Protein Interaction Networks, BioNLP 2012

Mixer

With Anthony Tomasic, John Zimmerman, Ian Hargraves, Ken Monkern, Steven Gardiner

Mixer uses programming-by-demonstration to turn repetitive information gathering requiring multi-step web surfing processes into something more like the fill-down behavior found in spreadsheets.

My contributions: Design and implementation of WOZ prototype in Flex. Operation of WOZ prototype in usability study. Analysis of study data. Collaborative design and implementation of genuine tool in JavaScript and Firefox plugin. Wrote the system architecture and WOZ study sections of publications.

Papers Steven Gardiner, Anthony Tomasic, John Zimmerman, Rafae Aziz, and Kathryn Rivard [Mazaitis]. Mixer: Mixed-initiative data retrieval and integration by example, Interact 2011 (Brian Shackel Award / Best Paper)

John Zimmerman, Kathryn Rivard [Mazaitis], Ian Hargraves, Anthony Tomasic, Ken Mohnkern. User-created forms as an effective method of human-agent communication, CHI 2009

VIO: Virtual Information Officer; Forms; WbE: Workflow by Example

With Anthony Tomasic, Issac Simmons, Marty McGuire, Jay Springfield

VIO/Forms/WbE were software components of the RADAR project which extracted common office tasks from email and generated tailored interfaces for completing each task.

My contributions: UI redesign of WbE including sketches, prototyping in Revolution (HyperCard), and paper pilot testing; wrote a column classifier for VIO (Java) to suggest field entries; wrote JavaScript for Forms to support a new functional-style API.

More info Zimmerman, J., Tomasic, A., Simmons, I., Hargraves, I., Mohnkern, K., Cornwell, J., and McGuire, R. M.. Vio: a mixed-initiative approach to learning and automating procedural update tasks, CHI '07.

Anthony Tomasic, R. Martin McGuire, Brad Myers, 2006. Workflow By Example: Automating Database Interactions via Induction. Carnegie Mellon University Technical Report Nr. CMU-ISRI-06-103, 2006.

Honeydew

With Anthony Tomasic

Honeydew was a brief-lived system for natural language meeting negotiation.

My contributions: UI sketches and storyboards for handling meeting scheduling in email.

Ballbot

With Ralph Hollis, inherited from Tom Lauwers; Microdynamics Systems Lab, Robotics Institute at CMU

Ballbot is a robot for human environments which balances on a single spherical wheel.

My contributions: When I inherited the project, the software was in its larval stages and based on bitwise communication with the hardware in C. I wrote an abstraction layer to provide for task scheduling and protect higher-level functions from the hardware details. This permitted the construction of behavior-based programming. I also worked on revisions of the controls system models and simulations. Upon departing for a leave of absence, I put together a comprehensive set of documentation describing the new software, tests in progress, pending tasks, and other crucial contextual knowledge acquired over my year with the project.

Competency-based grading

Competencies Working Group and Competencies Implementation Committee at Olin College

Olin College maintains a competency-based grading system in parallel with traditional letter grades. The Competencies Working Group designed the system to align with the Student Outcomes section of the ABET accreditation criteria for engineering programs, and the Competencies Implementation Committee carried out the plan. I served on both committees as a student.

My contributions: In addition to active participation in discussions, I performed many web-related tasks, including the development of an internal website to publish meeting minutes and related campus announcements, and designing a template-based web portfolio system deployed across the student body.

One of the primary components of Olin's competencies system is a conference-style (posters and presentations) exposition event to evaluate projects from each student at the end of every semester. I designed and implemented the student-facing frontend of a software system automating the creation of the expo schedule. The interface took the form of a web application which used JavaScript and MySQL to collect project, poster,

and presentation data from students, as well as allowing them to handle schedule conflicts relatively painlessly by trading timeslots with another student. A revision of this system remained in use for five years after I graduated.

SIS: Student Information System

Course project: Human Factors and Interface Design, Olin College

SIS is a web application giving faculty access to the grades and transcripts of their advisees. The interface did not support common advising needs. Our team used interviews, task scenarios, paper prototypes, and extensive user testing to come up with a way to present advisee information to complement the way faculty actually performed their advising tasks. The system was implemented in MySQL and JavaScript.

Behavior-based Robotics

Intelligent Vehicles Lab at Olin College

David Barrett started this lab in fall of 2004 to research autonomous ground vehicles.

My contributions: Design of a behaviors and arbiter system to attain basic navigation of a robot using sensory data from GPS, stereo vision, and LIDAR. I wrote custom simulations in Python and MATLAB, and also explored a Python package for robotics developed at Bryn Mawr (<http://pyrobotics.org>).

Education

Coursework: MS Human-Computer Interaction, Carnegie Mellon Univ. 2009-2015

Part-time while employed at CMU.

Went on leave of absence after spring 2015 for health reasons; never returned because the capstone requires a full-time commitment. I had two courses and the capstone project remaining to successfully complete the degree.

Coursework spanned interaction and user experience design processes, software design for user interfaces, and the psychology of human attention.

Coursework: PhD Robotics, Carnegie Mellon Univ. 2006-2007

Left the program after summer 2007. I enjoyed research; I had trouble being a graduate student.

Awarded full funding from the Robotics Institute as well as a Barbara Lazarus Women@IT fellowship from CMU.

Research focused on modeling and control of Ballbot, a balancing robot.

Coursework spanned foundational robotics topics including mathematics, sensing, action, and cognition.

BS Engineering, Franklin W. Olin College of Engineering

2002-2006

Member of inaugural class and recipient of full tuition and room scholarship.

Coursework included traditional engineering topics, computer science, design, music, and entrepreneurship.

Most coursework was project-based and relied on teamwork and peer teaching, integral parts of the Olin College curriculum.

Activities

Fiber Arts: Knit, Dye, Spin

2014-present

I own two antique spinning wheels (a German upright and an unsigned Canadian Production Wheel) and about a dozen hand spindles which transform fluff into string, which I then knit into hats, gloves, and shawls. I grow dye-producing plants and kettle-dye wool yarn and fiber. I process raw fleece, fresh from the sheep, into a format suitable for spinning. I travel to fiber festivals in the midwest and mid-atlantic.

Social Dance & Dance Instruction

1997-2002, Woodinville, WA; 2002-2006, Boston, MA; 2006-2014, Pittsburgh, PA

I started dancing Western Squares after school in 1997, and have continued to study and enjoy social dance ever since. In 2004 I learned a swing dance called Lindy Hop and began to compete, and eventually taught weekly classes for student dance clubs at Olin, MIT, and CMU. I added West Coast Swing to my repertoire in 2005, and was brought on to teach for PghWCS in 2009 when national champion instructor Stephanie McHenry went on parental leave. My role at PghWCS grew to include co-leadership of a team of six instructors teaching four levels of dance classes each week, plus monthly one-on-one practicas. I taught private lessons to students of intermediate level and below, and ran periodic workshops for our instructor team to work on pedagogical and andragogical techniques and build rapport among instructors. I served as club treasurer for 2013 and 2014, collecting weekly club income from dances, paying teachers and DJs, managing reimbursements, paying club taxes, and generating quarterly finance reports. I also served on the Visitors committee, handling registration, attendance, and payment for weekend workshops from out-of-town instructors.

Other Interests

Cooking • Sewing • Gardening • Soapmaking
Sketching • Herbal medicine • Bookbinding • Hand embroidery