

GREGORY DRUCK

gdruck@gmail.com
(443) 604-8603

SUMMARY

Former machine learning researcher who currently leads a team of scientists and engineers building AI products for the smart kitchen. Sees the big picture but is adept at diving into the details when needed. Enjoys working in fast-paced, collaborative environments that value creativity and leanness.

EXPERIENCE

Chief Data Scientist	June 2017 - Present
Head of Research	March 2014 - June 2017
Research Scientist / Engineer	May 2012 - March 2014
Yummly	Redwood City, CA

Tech lead and people manager for the Data Team, a ten person team that owns systems for content ingestion, search, personalization, computer vision, and data analysis, which together power many of the core features of the Yummly apps and website.

Key accomplishments:

Built ML systems to extract and parse online recipes, enabling Yummly to 1000X its publishers.

Developed ML algorithms that power Yummly's personalization engine, which synthesizes explicit preferences, implicit preferences, and context.

Created a tool to mine interaction data to surface insights into how users cook, which is heavily utilized by the product, content creation, and business development teams.

Built the Ingredient Recognition feature, which allows a user to scan (using object detection and text recognition) ingredients with a phone camera to track inventory and find recipes.

Diverse set of responsibilities also includes working closely with the product team to define strategy and serving as a member of the company leadership team.

Postdoctoral Scientist	September 2011 - May 2012
Yahoo! Research	Santa Clara, CA

Developed algorithms for mining user-generated refinements to online instructions, identifying entities and classifying intent in search queries.

Research Assistant	July 2005–August 2011
University of Massachusetts Amherst	Amherst, MA

Research in semi-supervised and active machine learning for information extraction and natural language processing with Dr. Andrew McCallum.

Intern, Research	June 2008–September 2008
Google	New York, NY

Developed an algorithm for semi-supervised dependency parsing that uses linguistic prior knowledge as a training signal. This work was published at ACL 2009.

Intern, Live Labs	June 2006–August 2006
--------------------------	-----------------------

Microsoft

Redmond, WA

Developed an algorithm for fast test time inference in graphical models using A* search with a learned heuristic function. This work was published at AISTATS 2007.

EDUCATION

University of Massachusetts Amherst

Amherst, MA

Ph.D., Computer Science, 2011

M.S., Computer Science, 2008

Adviser: Dr. Andrew McCallum

Focus: Machine Learning and Natural Language Processing

Johns Hopkins University

Baltimore, MD

B.S., Computer Science, 2005

Minor in Applied Mathematics and Statistics

RESEARCH EXPERIENCE

Thirteen refereed publications in top conferences (ACL, EMNLP, ICML, UAI, SIGIR, and KDD) with around 900 total citations in Google Scholar. See `gregorydruck.name` for a comprehensive list. Frequent program committee member and reviewer for conferences. Winner of Yahoo Key Scientific Challenges Award for Information Extraction, 2009. Speaker on food tech at conferences like the Computer Cooking Competition, Bon Appetech, eTail East, Customize, and SXSW.

TECHNICAL SKILLS

Machine Learning: Natural Language Processing, Information Extraction, Information Retrieval, Recommender Systems, Computer Vision, Data Analysis, Classification, Regression, Ranking, Clustering, Sequence Labeling, Parsing, Deep Learning, Semi-Supervised Learning, Active Learning

Languages: Scala, Python, Clojure, Java, SQL

ML Libraries: Tensorflow, scikit-learn, Stanford NLP, Mallet

Tools: Redshift, BigQuery, Spark, Hadoop, Solr, Elasticsearch, AWS, GCP, Docker, Jenkins, Github