

pdgroves@gmail.com
http://petergroves.com
Chicago, Illinois

PETER GROVES

ABOUT

Software Engineer with a focus on Machine Learning.

Specialties: Machine Learning, DevOps, Data Modeling, Python, Java, Linux, Ocaml.

EXPERIENCE

NCSA — *Senior Software Engineer*

OCTOBER 2017 - PRESENT

As a Senior Software Engineer, my focus is on Backend and Machine Learning within the Visual Intelligence for Biology Group at the National Center for Supercomputing Applications.

- Lead Backend Developer of a platform for data science and web applications. This platform is based on common data types shared between frontend, backend, and data science jobs.
- Project Manager and Lead Developer of the DeepPhoton Project.
Our core focus is to build better tools for annotating complex medical images and simplify data management of the resulting training data for machine learning.
- Lead Backend Developer for Omix.
Omix is a web app that visualizes microbiome analytical results built in collaboration with Mayo Clinic's Center for Individualized Medicine.

ILLINOIS APPLIED RESEARCH INSTITUTE — *Senior Software Engineer*

SEPTEMBER 2015 - OCTOBER 2017

The Visual Intelligence for Biology Group moved from ARI to NCSA in 2015 due to a reorg - projects and team were not changed.

GROUPON — *Software Engineer*

SEPTEMBER 2013 - MARCH 2015

As part of the Automated Merchandising team, I was responsible for bundling deals (Groupons) into widgets to be displayed anywhere on the Groupon website or mobile apps. Personal areas of focus:

- Machine learning system for adjusting the relative rankings of widgets in different contexts.
- Integration test suite of the core system, written in Python.
- Build automation using Python, Capistrano, Maven, Jenkins, and cron.
- General software engineering of a Java webapp on an 8 person agile software team.
- International rollout of core service to 20 countries in Europe and Latin America.

INDEPENDENT CONTRACTOR — *Machine Learning Software Developer*

OCTOBER 2006 - SEPTEMBER 2013

As an independent contractor, I worked on the following projects:

- Nexlp (2013). An e-Discovery startup in Chicago. Their core product uses a natural language processing toolkit from the University of Illinois and a graph database (Neo4j) to analyze sets of millions of emails at a time. The primary deliverable was a pattern detection module that combined frequent item set analysis and anomaly detection to generate patterns of the form "Bob emailed Sally late at night about Chicago 12 times during the week of Dec 4, 2006, but normally this occurred 0.03 times/week."
- rVibe (2013). A boutique maker of training software for the pharmaceutical industry. Handled devOps for the company for nine months. Designed and implemented build tools and a performance benchmarking suite. Performed weekly deploys to production servers and adjusted agile release schedule and methodology as needed.
- Fuzzy economics project (2011/2012). Privately funded by a (stealth) organization, this was a bleeding edge project to create a system for creating expansive yet detailed ontologies of hypotheses and their supporting evidence. Built the full stack prototype using Java EE with JSP and the JavaBayes toolkit for Bayesian Networks.

DESIGNBYROBOTS — *Founder*

SEPTEMBER 2006 - SEPTEMBER 2013

Developed an algorithmic trading application written in OCaml using a genetic algorithm and statistical modeling techniques to find market price patterns on a time scale of less than one hour. Trading strategy optimization was the first application that uses DesignByRobot's data model for machine learning and automated design technology.

EDUCATION

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN — *MS Computer Science*

2001 - 2003

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN — *BA Agricultural Engineering*

1997 - 2001