David Ding

Website: davidding.me
Email: david@davidding.me

Experience

Google

Software Engineer — *New York, NY; 08/14 – 01/17, 05/17 – Present*

- My team takes business descriptions that clients want to add to Google Maps, and matches them with existing businesses in the canonical Maps database. These business descriptions are incredibly noisy and come from numerous sources like ad feeds and user edits, with ∼100,000 requests a day flowing through our automated system. My responsibilities:
  - Engineer high-quality features for our central ML clustering tool, subject to various memory, latency, and threading constraints. The tool retrieves all candidates, computes features, and scores everything within 2s 99% of the time, is thread-safe, and has a ∼100MB footprint.
  - Improve, retrain, and maintain the aforementioned ML model. Improvements to the input features and the model itself have led to a ∼30% relative gain in the tool's precision, to over 90%.
  - Develop and optimize online methods to retrieve matching business candidates for an input business from a database index and update the index, ensuring that production systems stay up while processing ∼1000 index updates per second.

ML Rotation Participant — *Mountain View, CA; 01/17 – 05/17*

- Researched interpretable ML models: specifically, extending the research group's highly interpretable lattice-based models to deeper and more flexible neural networks.
  - Conducted all the theoretical work and spearheaded the experimental efforts, showing that the proposed idea had real potential in practice.
  - Ultimately, achieved state-of-the-art performance for regression and classification with monotonicity guarantees, resulting in a NIPS 2017 publication (see below).

Intern (Site Reliability Engineering) — *Los Angeles, CA; 06/13 – 09/13*

- Designed, wrote, tested, and executed a tool for transferring live monitoring services from one cloud service to another, in preparation for a major infrastructure revision.

Publications

- Seungil You, **David Ding**, Kevin Canini, Jan Pfeifer, Maya Gupta. Deep Lattice Networks and Partial Monotonic Functions. In the proceedings of NIPS 2017. `arXiv:1709.06680`.

Education

Columbia University

M.S.; Machine Learning; GPA 4.0/4.0 — *New York, NY; 09/14 – 12/16*

California Institute of Technology

B.S.; Mathematics and Computer Science; GPA 4.0/4.0 — *Pasadena, CA; 09/10 – 06/14*

Technical Skills

- **Actively using**: C++, Python, Linux, Piper (version control), SQL
- **Significant experience with**: TensorFlow, C, Scheme, LaTeX, MATLAB

Some Hobbies and Interests

- **Music**: singing (New York City Gay Men's Chorus, tenor 2), piano, saxophone, guitar
- **Puzzles**: MIT Mystery Hunt (winning team 2012), Google Games (undefeated in LA area among all undergraduate and graduate teams, 2011-2014), and various other puzzle hunts.
- **Outdoors**: mountaineering (Aconcagua, Kilimanjaro, Mt. Whitney, and more), hiking (led week-long treks through the Sierra Nevada and the Pacific Northwest), skiing (type 3)

Last updated: August 2018