Wei Dai (David)

(203) 300-9086

david@davidwd.org

davidwd.org

PROFILE

Skilled and detail-oriented software developer and engineering leader with extensive backgrounds in machine learning and distributed systems.

TECHNICAL SKILLS

Machine Learning: Parameter Server, Distributed Optimization, Deep Learning, Computer Vision, Medical Imaging, Graphical Models
Technologies: TensorFlow, Spark, Pulsar, GraphQL, Protocol Buffer, ZeroMQ, Docker, Kubernetes
Languages: Python, C/C++, Java

EXPERIENCES

Petuum Inc.

Senior Director of Engineering — Sept 2017 - Present
Director of Product Development — May 2017 - Sept 2017
Senior Manager — Jul 2016 - May 2017

- As part of the founding team, I help grow the engineering team from <10 in 2016 to over 60 full-time engineers in 2Q18. Both a trench worker and a manager, I influence product design as a domain expert in machine learning, and deliver quarterly internal platform product milestones from 3Q17 - 1Q18. Currently I lead the engineering division, marshaling the work of 8 teams and 4 front-line managers.
- I led the former medical imaging team and personally developed best in class deep learning methods around chest x-ray to assist radiologists in diagnosis. The project led to the first contract for the company. Petuum's healthcare solution assist doctor workflows by leveraging image, textual, and structured data such as chest x-rays, CT scan, admission note, lab test, and others. The team has grown to cover other non-healthcare industries.
- During a product pivot, I served as tech lead and co-designed and implemented a unified framework that flexibly combines all data and machine learning functions developed across platform and solution teams.
- Proposed and implemented org adjustment within platform teams with broad feedback and support, resulting in a simpler org structure that empowers engineering managers and tech leads.
- Fostering engineering culture by hosting bi-weekly tech talks, performing 1-on-1, planning and executing performance review and promotions for all engineering. Recommended engineering ladders and support long-term efforts in formalizing engineering and management titles.

Bosch Research
Research Intern — Feb 2016 - Aug 2016

- Used TensorFlow on a GPU cluster to train state-of-the-art convolutional neural networks for environmental sound analysis. Published two papers in ICASSP 2017.

Facebook
Software Engineer Intern — May 2015 - Aug 2015

- Developed a distributed machine learning backend for large-scale logistic regression using Petuum parameter server.
- Benchmarked Petuum against Facebook's internal system and open source Vowpal Wabbit; showed that Petuum achieves high system throughput and produces comparable to better models.

Google Inc. (Pittsburgh Office)
Software Engineer Intern — May 2013 - Aug 2013

- Contributed to the Ad Quality backend; developed a hyperparameter tuning framework to optimize SmartAds training system with convex and non-convex optimization algorithms; built a web frontend for other teams to interface with the framework. The system was later published as Google Vizier.

LinkedIn
Software Developer Intern — Jun 2012 - Aug 2012

- Implemented several background tasks in the payment backend using Java, Oracle SQL, Python, and Spring Framework.

EDUCATION

Carnegie Mellon University — 2018
Ph.D. in Machine Learning

California Institute of Technology — 2012
B.Sc with Honor in Computer Science

Wesleyan University — 2012
B.A. with High Honor in Physics and Mathematics

PUBLICATIONS

Over 20 publications in top machine learning, artificial intelligence, systems, and computer vision venues with more than 700 Google Scholar citations. Recipient of a best paper award. Two pending US patents.