

Experience

- 10/2016 – current **Senior Machine Learning engineer**, Facebook Inc., New York.
- *Team characteristics*: Place Visit Detection team built a system to predict real world visitation. It is a custom, large-scale ML system with limited labeled and plenty of unlabeled data. It requires both backend and modeling skills as the team owns its serving infrastructure to all Facebook users in real-time.
 - *Tech lead role*: Lead team's efforts on Ads Attribution and Optimization. Build roadmaps, coordinated with cross functional partners and provided direction to junior engineers.
 - *Custom ML solution*: Designed and developed an end-to-end solution from mobile client to Ads infrastructure to improve Measurement for high-value clients. It unlocked a new source of revenue and is currently running in production.
 - *Cross team ML dependency*: Architected and supervised design of two connected ML systems across two teams with training, evaluation, and versioning. Contained ML dependencies in the design to enable automated end-to-end deployment. Identified the discrepancy between the needs of consumer and Ads clients which led to targeted improvements in both use cases.
 - *Semi-supervised learning*: Designed a subcomponent of the system which serves as a backbone of our prediction model since it provides data about each place in the real world. The sheer size of the collected data from all the Facebook users made it particularly challenging for offline processing yet through a few simplifying assumptions we managed to do it reliably.
 - *Training data*: Developed 3 generations of collection systems to accommodate changes to a broader Facebook environment. This included negotiating with other teams to include features that were compatible with our needs as well as adapting our training process to fit the new data distribution.
 - *Safe feature engineering*: Modified team's training methodology to prevent engineers from making configuration mistakes when shipping live models. This was particularly important given that the team's doubled in size.
 - *Cross team direction*: Influenced goal setting to ensure the team was aligned with Ads requirements to make meaningful progress on final business objectives.
 - *Expanding team's scope*: Identified new ways of using data provided by users and worked with Privacy partners to approve it. This effort was crucial to improve a larger Ads optimization effort in London in the longer term.
 - *Team development*: Mentored junior engineers on the direct team and a remote team in London.
 - *Recruitment*: Contributed to the company through regular interviewing as well as helping new hires find teams
- 08/2015 – 08/2016 **Software engineer**, Flux Factory Inc, San Francisco.
- Developed multiple new user-facing features around users' data pipelines including new data auditing tools.
 - Improved the RESTful Flux's DataTable API by working closely with their SDK team and gathering feedback from developers. Also cleaned up the API to return informative errors and appropriate HTTP error codes.
 - Gained working knowledge of deployment systems including Docker and CoreOS by being part of the on-call rotation. Also provided the user support team with help to access data in Google Cloud Storage and MySQL databases.
 - Coordinated API/service changes with a remote engineering team and ensured backwards compatibility with old APIs.

Master Thesis

- Title Application of convolutional neural networks to RL control problems
Supervisors Z. Ghahramani, M. Hoffman
Grade **1st Class**

Education

- 2011 – 2015 **Trinity College, University of Cambridge**, MEng, BA (Hons.), **Distinction (MEng), 1st Class (BA)**.
- Achieved **1st Class** results in courses including Machine Learning, Computer Vision, Statistical Pattern Processing, Speech and Language Processing, Computational Neuroscience and Social and Technological Networks Analysis

Awards

- 2013 – 2015 **Senior Scholar**, Trinity College, University of Cambridge.
2012 **Junior Scholar**, Trinity College, University of Cambridge.
2012 – 2015 **Tripes Examination Prize**, Trinity College, University of Cambridge.

Skills

- Programming **Hack/PHP, C/C++, Java, Python, Go, SQL, Distributed Systems, Spark**
Machine Learning **ML Systems design, Evaluation, Feature Engineering, ML Versioning**
Deep Learning, Reinforcement Learning, Probabilistic Rankings, Convolutional Neural Networks
Non-parametric models, Neural Networks, MCMC methods, CUDA, SVM, Decision Trees
Gaussian Processes, Mixture Models, Hidden Markov Models, LDA, EM, RBM