

William Chan

✉ williamchan@cmu.edu ☎ +1 650 450-9455 🌐 williamchan.ca

Education

Carnegie Mellon University

PhD Electrical and Computer Engineering 2016

MS Electrical and Computer Engineering 2013

Thesis: *End-to-End Speech Recognition Models*

Research Areas: *Machine Learning, Deep Learning, Speech Recognition*

University of Waterloo

BASc Computer Engineering, Management Sciences Option 2011

National University of Singapore

Exchange Student Computer Engineering 2010

Professional Experiences

Google Research - Google Brain Research Intern

2016

- Overdosing Google Brain with *Latent Sequence Decompositions*
- Architected Google Brain models with *Very Deep Convolutional Networks for End-to-End Speech Recognition*
- Research Areas: machine learning, deep learning, speech recognition

Google Research - Google Brain Research Intern

2015

- Teaching Google Brain how to *Listen, Attend and Spell*
- Research Areas: machine learning, deep learning, speech recognition

Google Research - Google Brain Research Intern Part-Time

2015

- Cooking secret sauces part-time while continuing PhD @CMU
- Various code optimizations, ??% efficiency in CPU and memory over baseline DistBelief implementation
- Research Areas: machine learning, deep learning

Google Research - Natural Language Processing Research Intern Part-Time

2014

- Continuation of summer internship but now working 1 day / week while continuing PhD @CMU
- ?x production inference code optimization under certain scenarios
- Research Areas: machine learning, deep learning, natural language processing

Google Research - Natural Language Processing Research Intern

2014

- SAFT Machine Intelligence
- Deep Learning, Natural Language Processing and Multilingual Coreference Resolution
- ?% increase in coref performance on the CoNLL dataset over baseline system
- Research Areas: machine learning, deep learning, natural language processing

Google - Machine Learning Ads Intern

2012

- Display Campaign Optimizer
- Replaced a heuristic ranking model with a machine learning ranking model, ??% increase in Click Through Rate and ??% increase in customer return on investment in certain scenarios (patent pending)
- Added various Map-Reduce analysis scripts
- Added Business Value: higher ROI for customers, higher revenue for Google and more satisfied customers

- TD Securities** - Quantitative Financial Engineer Intern 2010
- TD Bank's trade floor in the Quantitative Research Group
 - Optimized Monte Carlo Simulation through dividend events consolidation, 500%-1700% performance gain on dividend heavy deals
 - Added a new mixed Dividend Pricing Pricing Model, allows pricing of most exotic derivatives using an inter-blended discrete dollar and proportional dividend schedule
 - Architected new 2D Interpolation Infrastructure and added Thin Plate Spline Interpolation for Local Volatility Surfaces
 - Cholesky decomposition optimizations for FX/Equity correlation matrices by bootstrapping similar matrices used in Monte Carlo Simulations
 - Added Business Value: faster execution, more accurate pricing, happier clients, better risk management and hedging
- Amazon** - Software Engineer Intern 2009
- Amazon Prime AdAttribution
 - Re-architected database schemas (Oracle and MySQL)
 - Modified data mining infrastructure to more efficiently attribute Amazon Prime signups to specific ads for statistical analysis
 - Added Business Value: faster and more effective analysis of customer behaviour
- Google** - Disk Platform Engineer Intern 2009
- Research and analysis on disk latency and bandwidth
 - Data analysis on disk performance across entire Google's server fleet
 - Added a new Linux Kernel Power Capping Module to output detailed CPU usage to assist in power management across Google's server fleet
 - Added Business Value: maximize utility, lower capital and operating costs, negotiation leverage with disk vendors
- Intel** - Mobile Platform Architecture Engineer Intern 2008
- Adaptive Power Algorithms Research
 - Designed and prototyped several dynamic frequency algorithms for CPU/GPU to reduce dynamic power, +55% CPU power reduction with no performance penalty in certain apps/games
 - Multi-Chip Package Power Controller Integral Algorithm
 - New methodology to increase CPU and GPU performance while maintaining same thermal design power envelope
 - Statistical Data Analysis on Power Traces - analyze and isolate power events for simulation and forecasting next generation mobile platform architectures
 - Added Business Value: additional product value with zero capital and marginal unit cost
- NVIDIA** - Driver Software Engineer Intern 2007
- Driver Optimizations
 - Various driver optimizations including +20% performance on Call of Duty 2 on G80
 - Added Business Value: additional product utility without additional capital or marginal costs, product differentiation against competitor
- AMD** - Hardware Engineer Intern 2007
- High Bandwidth Video Playback Research and Development
 - Developed custom video system for 1920x1080 @ 120Hz playback with throughput of over 6Gb/s and designed to scale over 240Hz
 - Developed (Verilog) FPGA for a custom proprietary board for converting a DVI/HDMI signals to LVDS for a variety of LCD panels
 - Floating-point Optimizations (SSE2 and SSE3)
 - Added Business Value: capital cost savings (> 100k CAD) with custom in-house solutions vs commercial purchases

Research Publications

1. Takaaki Hori, Shinji Watanabe, Yu Zhang and **William Chan**, “Advances in Joint CTC-Attention based End-to-End Speech Recognition with a Deep CNN Encoder and RNN-LM,” in INTERSPEECH 2017.
2. **William Chan**, Yu Zhang, Quoc Le and Navdeep Jaitly, “Latent Sequence Decompositions,” in ICLR 2017.
3. Yu Zhang, **William Chan** and Navdeep Jaitly, “Very Deep Convolutional Networks for End-to-End Speech Recognition,” in ICASSP 2017.
4. **William Chan** and Ian Lane, “On Online Attention-based Speech Recognition and Joint Mandarin Character-Pinyin Training,” in INTERSPEECH 2016.
5. Guan-Lin Chao, **William Chan** and Ian Lane, “Speaker-Targeted Audio-Visual Models for Speech Recognition in Cocktail-Party Environments,” in INTERSPEECH 2016.
6. **William Chan**, Navdeep Jaitly, Quoc Le and Oriol Vinyals, “Listen, Attend and Spell: A Neural Network for Large Vocabulary Conversational Speech Recognition,” in ICASSP 2016. **Best Student Paper Award.**
7. **William Chan** and Ian Lane, “Deep Recurrent Neural Networks for Acoustic Modelling,” in arXiv 2015.
8. **William Chan**, Nan Rosemary Ke and Ian Lane, “Transferring Knowledge from a RNN to a DNN,” in INTERSPEECH 2015.
9. **William Chan** and Ian Lane, “Deep Convolutional Neural Networks for Acoustic Modeling in Low Resource Languages,” in ICASSP 2015.
10. **William Chan** and Ian Lane, “Distributed Asynchronous Optimization of Convolutional Neural Networks,” in INTERSPEECH 2014.
11. Le Nguyen, Pang Wu, **William Chan**, Wei Peng and Joy Zhang, “Predicting Collective Sentiment Dynamics from Time-series Social Media,” in ACM SIGKDD WISDOM 2012.
12. **William Chan** and Jason Lohn, “Spike Timing Dependent Plasticity with Memristive Synapse in Neuromorphic Systems,” in IEEE IJCNN 2012. **Travel Fellowship for Spiking Neural Networks.**

Service

- Reviewer (past and/or present) for:
 1. IEEE Signal Processing Letters
 2. IEEE Transactions on Evolutionary Computation
 3. IEEE Transactions on Neural Networks and Learning Systems
 4. International Conference on Artificial Intelligence and Statistics (AISTATS)
 5. International Conference on Learning Representations (ICLR)
 6. International Conference on Machine Learning (ICML)
 7. INTERSPEECH
 8. Neural Information Processing Systems (NIPS)

References

- Ian Lane (Carnegie Mellon University)
- Bhiksha Raj (Carnegie Mellon University)
- Florian Metze (Carnegie Mellon University)
- Chris Dyer (Carnegie Mellon University and DeepMind)
- Quoc Le (Google Brain)
- Oriol Vinyals (DeepMind)
- Navdeep Jaitly (NVIDIA Research)