

William Chan

✉ wchan212@gmail.com ☎ +1 650 450-9455 🌐 williamchan.ca

George Brown College Culinary and Baking Arts	2019 - current
Carnegie Mellon University PhD Electrical and Computer Engineering	2011 - 2016
University of Waterloo BASc Computer Engineering, Management Sciences Option	2006 - 2011
National University of Singapore Computer Engineering Exchange Student	2010

Professional Experiences

Ideogram AI - Cofounder/CTO	2022 - current
<ul style="list-style-type: none">• Mission: help people become more creative through generative AI	
Google - Google Brain Staff Research Scientist	2017 - 2022
<ul style="list-style-type: none">• I spend most of my time learning from my mentors and peers• Highlight: Fundamental research in the areas of iterative refinement generation (see Insertion Transformer, Imputer, WaveGrad, Palette, Imagen) with applications to speech recognition, machine translation, text-to-speech, text-to-image, and text-to-video.• Highlight: Fundamental research to improve speech recognition with generational leap in performance (see Listen Attend and Spell, SpecAugment and SpeechStew)• Mentored multiple AI residents, interns, and junior researchers• Research Areas: machine learning, deep learning, natural language processing, speech recognition, speech synthesis, computer vision	
Google - Google Brain Research Intern	2016
<ul style="list-style-type: none">• <i>Latent Sequence Decompositions</i>• <i>Very Deep Convolutional Networks for End-to-End Speech Recognition</i>• Research Areas: machine learning, deep learning, speech recognition	
Google - Google Brain Research Intern	2015
<ul style="list-style-type: none">• <i>Listen, Attend and Spell</i>• Research Areas: machine learning, deep learning, speech recognition	
Google - Google Brain Research Intern Part-Time	2015
<ul style="list-style-type: none">• 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 - Natural Language Processing Research Intern Part-Time	2014
<ul style="list-style-type: none">• 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 - Natural Language Processing Research Intern	2014

- SAFT Machine Intelligence
- Deep Learning, Natural Language Processing and Multilingual Coreference Resolution
- 7% 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, 22% increase in Click Through Rate and 22% increase in customer return on investment in certain scenarios (patent pending)
- Added various Map-Reduce analysis scripts
- Impact: 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
- Impact: 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
- Impact: 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
- Impact: 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
- Impact: 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
- Impact: 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)
- Impact: capital cost savings with custom in-house solutions vs commercial purchases

Publications

1. Su Wang*, Chitwan Saharia*, Ceslee Montgomery*, Jordi Pont-Tuset, Shai Noy, Stefano Pellegrini, Yasumasa Onoe, Sarah Laszlo, David Fleet, Radu Soricut, Jason Baldridge, Mohammad Norouzi†, Peter Anderson†, **William Chan**†, "Imagen Editor and EditBench: Advancing and Evaluating Text-Guided Image Inpainting," in arXiv 2022.
2. Rosanne Liu, Dan Garrette, Chitwan Saharia, **William Chan**, Adam Roberts, Sharan Narang, Irina Blok, RJ Mical, Mohammad Norouzi, Noah Constant, "Character-Aware Models Improve Visual Text Rendering," in arXiv 2022.
3. Jonathan Ho*, **William Chan***, Chitwan Saharia*, Jay Whang*, Ruiqi Gao, Alexey Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David Fleet, and Tim Salimans*, "Imagen Video: High Definition Video Generation with Diffusion Model," in arXiv 2022.
4. Daniel Watson, **William Chan**, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi, "Novel View Synthesis with Diffusion Models," in arXiv 2022.
5. Chitwan Saharia*, **William Chan***, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi*, "Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding," in NeurIPS 2022.
6. Jonathan Ho*, Tim Salimans* Alexey Gritsenko, **William Chan**, Mohammad Norouzi, and David Fleet, "Video Diffusion Models," in NeurIPS 2022.
7. Chitwan Saharia, **William Chan**, Huiwen Chang, Chris A. Lee, Jonathan Ho, David Tim Salimans, J. Fleet, and Mohammad Norouzi, "Palette: Image-to-Image Diffusion Models," in SIGGRAPH 2022.
8. Daniel Watson, **William Chan**, Jonathan Ho, and Mohammad Norouzi, "Learning Fast Samplers for Diffusion Models by Differentiating Through Sample Quality," in arXiv 2022.
9. Jonathan Ho*, Chitwan Saharia*, **William Chan**, David Fleet, Mohammad Norouzi, and Tim Salimans, "Cascaded Diffusion Models for High Fidelity Image Generation," in JMLR 2022.
10. Yu Zhang, Daniel S Park, Wei Han, James Qin, Anmol Gulati, Joel Shor, Aren Jansen, Yuanzhong Xu, Yanping Huang, Shibo Wang, Zongwei Zhou, Bo Li, Min Ma, **William Chan**, Jiahui Yu, Yongqiang Wang, Liangliang Cao, Khe Chai Sim, Bhuvana Ramabhadran, Tara N Sainath, Françoise Beaufays, Zhifeng Chen, Quoc V Le, Chung-Cheng Chiu, Ruoming Pang, and Yonghui Wu, "BigSSL: Exploring the Frontier of Large-Scale Semi-Supervised Learning for Automatic Speech Recognition," in IEEE Journal of Selected Topics in Signal Processing 2022.
11. Daniel Watson, Jonathan Ho, Mohammad Norouzi, and **William Chan**, "Learning to Efficiently Sample from DDPMs," in arXiv 2021.
12. Chitwan Saharia, Jonathan Ho, **William Chan**, Tim Salimans, David Fleet, and Mohammad Norouzi, "Image Super-Resolution via Iterative Refinement," in PAMI 2022.
13. Nanxin Chen, Yu Zhang, Heiga Zen, Ron J. Weiss, Mohammad Norouzi, Najim Dehak, and **William Chan**, "WaveGrad 2: Iterative Refinement for Text-to-Speech Synthesis," in INTERSPEECH 2021.

14. Edwin Ng, Chung-Cheng Chiu, Yu Zhang*, and **William Chan***, "Pushing the Limits of Non-Autoregressive Speech Recognition," in INTERSPEECH 2021.
15. **William Chan**, Daniel S. Park, Chris A. Lee, Yu Zhang, Quoc V. Le, and Mohammad Norouzi, "SpeechStew: Simply Mix All Available Speech Recognition Data to Train One Large Neural Network," in arXiv 2021.
16. Nanxin Chen, Yu Zhang*, Heiga Zen, Ron J. Weiss, Mohammad Norouzi, and **William Chan***, "WaveGrad: Estimating Gradients for Waveform Generation," in ICLR 2021.
17. Harris Chan, Jamie Kiros, and **William Chan**, "Multichannel Generative Language Model: Learning All Possible Factorizations Within and Across Channels," in Findings of EMNLP 2020.
18. Chitwan Saharia*, **William Chan***, Saurabh Saxena and Mohammad Norouzi, "Non-Autoregressive Machine Translation with Latent Alignments," in EMNLP 2020.
19. **William Chan***, Mitchell Stern*, Jamie Kiros and Jakob Uszkoreit, "An Empirical Study of Generation Order for Machine Translation," in EMNLP 2020.
20. **William Chan**, Chitwan Saharia, Geoffrey Hinton, Mohammad Norouzi and Navdeep Jaitly, "Imputer: Sequence Modelling via Imputation and Dynamic Programming," in ICML 2020.
21. Daniel Park*, Yu Zhang*, Chung-Cheng Chiu, Youzheng Chen, Bo Li, **William Chan**, Quoc Le and Yonghui Wu, "SpecAugment on Large Scale Datasets," in ICASSP 2020.
22. Harris Chan, Jamie Kiros and **William Chan**, "Multilingual KERMIT: It's Not Easy Being Generative," in NeurIPS Workshop on Perception as Generative Reasoning 2019.
23. Laura Ruis, Mitchell Stern, Julia Proskurnia and **William Chan**, "Insertion-Deletion Transformer," in EMNLP WNGT 2019.
24. Lala Li and **William Chan**, "Big Bidirectional Insertion Representations for Documents," in EMNLP WNGT 2019.
25. **William Chan***, Nikita Kitaev*, Kelvin Guu*, Mitchell Stern* and Jakob Uszkoreit, "KERMIT: Generative Insertion-Based Modeling for Sequences," in arXiv 2019.
26. Daniel S. Park, **William Chan**, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk and Quoc V. Le, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," in INTERSPEECH 2019.
27. Jonathan Shen, et al., "Lingvo: a Modular and Scalable Framework for Sequence-to-Sequence Modeling," in arXiv 2019.
28. Mitchell Stern, **William Chan**, Jamie Kiros and Jakob Uszkoreit, "Insertion Transformer: Flexible Sequence Generation via Insertion Operations," in ICML 2019.
29. Bo Li, Yu Zhang, Tara Sainath, Yonghui Wu and **William Chan**, "Bytes are All You Need: End-to-End Multilingual Speech Recognition and Synthesis with Bytes," in ICASSP 2019.
30. Sara Sabour, **William Chan** and Mohammad Norouzi, "Optimal Completion Distillation for Sequence Learning," in ICLR 2019.
31. Jamie Ryan Kiros and **William Chan**, "InferLite: Simple Universal Sentence Representations from Natural Language Inference Data," in EMNLP 2018.
32. Jamie Ryan Kiros*, **William Chan*** and Geoffrey Hinton, "Illustrative Language Understanding: Large-Scale Visual Grounding with Image Search," in ACL 2018.
33. 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.
34. **William Chan**, Yu Zhang, Quoc Le and Navdeep Jaitly, "Latent Sequence Decompositions," in ICLR 2017.
35. Yu Zhang, **William Chan** and Navdeep Jaitly, "Very Deep Convolutional Networks for End-to-End Speech Recognition," in ICASSP 2017.
36. **William Chan** and Ian Lane, "On Online Attention-based Speech Recognition and Joint Mandarin Character-Pinyin Training," in INTERSPEECH 2016.
37. Guan-Lin Chao, **William Chan** and Ian Lane, "Speaker-Targeted Audio-Visual Models for Speech Recognition in Cocktail-Party Environments," in INTERSPEECH 2016.
38. **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.**

39. **William Chan** and Ian Lane, “*Deep Recurrent Neural Networks for Acoustic Modelling*,” in arXiv 2015.
40. **William Chan***, Nan Rosemary Ke* and Ian Lane, “*Transferring Knowledge from a RNN to a DNN*,” in INTERSPEECH 2015.
41. **William Chan** and Ian Lane, “*Deep Convolutional Neural Networks for Acoustic Modeling in Low Resource Languages*,” in ICASSP 2015.
42. **William Chan** and Ian Lane, “*Distributed Asynchronous Optimization of Convolutional Neural Networks*,” in INTERSPEECH 2014.
43. 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.
44. **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.**

Patents

- US10540962B1, US20200026765A1, US10510004B2, US20190258713A1, US20190354808A1, US20200151567A1, US20200234011A1, US20200372356A1
- Others under review.

Service

- Reviewer (past and/or present) for: AISTATS, ICASSP, ICLR, ICML, INTERSPEECH, NIPS/NeurIPS