David M. Chan

AI researcher in multimodal learning and human-computer interaction studying how AI and humans can work together to make the world a better place.

davidchan@berkeley.edu @
+1 303-475-5153 ☎
https://dchan.cc 🌐
DavidMChan ⌗
Berkeley, CA 📍

Education

2024
University of California, Berkeley, Ph.D. (EECS)
4.0 GPA, Thesis: Understanding, Building, and Evaluating Models for Context Aware Conditional Natural Language Generation, Adv. John Canny

2020
University of California, Berkeley, MS (EECS)
4.0 GPA, Thesis: TSNE-CUDA: GPU-Accelerated TSNE, Adv. John Canny

2017
University of Denver, BSc (Mathematics, Computer Science)
4.0 GPA, Summa Cum Laude, Honors with Distinction, Phi Beta Kappa

Experience

2021-2024
Amazon *Student Researcher*
- Developed and implemented large language models and foundational models for context-drive speech understanding leading to 13.8% relative WER improvements over baseline models.
- Researched model distillation and off-policy machine leading to algorithms reducing hot-fixing domain-alignment times by over 1K GPU hours, saving up to $3K per model training run.
- Designed multi-modal video-conditioned ASR model frameworks based on PyTorch and Tensorflow back-ends for pre-training, improving large model performance on low-resource devices by up to 45%.

Summer 2020
Google *Intern*
- Developed systems combining large language models (LLMs) and large multimodal models (LMMs) for Image captioning, leading to up to 84% MRR improvement over baseline approaches.
- Collaborated with a diverse team to design large-scale GPU training procedures with up to 1K GPUs across 128 Nodes for machine learning algorithms.

Summer 2019
Dropbox *Intern*
- Architected, developed, and deployed Dropbox's first ever deep learning based AI for multimodal content classification, achieving 0.96 F1 on content classification tasks.

2018
NASA Jet Propulsion Laboratory *Intern*
- Developed flight-ready provable AI systems for high-risk environments.
- Contributed core software to top-5 DARPA Subterranean Challenge Team.

2014-2017
Dreamface Technologies *Research Assistant*

Programming Languages

I have a fairly advanced knowledge of:

`Python`

I have written production code in:

`CUDA` `C++` `HTML` `TypeScript` `CSS`

I've written useful code in:

`R` `Matlab` `C` `C#` `Rust`

Technologies/Frameworks

AI/ML:

`Tensorflow` `PyTorch` `Jax` `Transformers` `SKLearn`

Distributed/Parallel:

`OpenMP` `MPI` `Spark` `Slurm`

Cloud:

`Kubernetes` `AWS` `GCP`

Development:

`Git/Svn` `Bash` `Bazel`

Misc:

`MS Office` `Photoshop` `In-Design` `Illustrator`

Selected Publications

Articles: 39   Citations: 1642   H-Index: 11   i10-index: 12

See, Say, and Segment: Teaching LMMs to Overcome False Premises, Computer Vision and Pattern Recognition (CVPR), 2024

CLAIR: Evaluating Image Captions with Large Language Models, Empirical Methods in Natural Language Processing (EMNLP), 2023

IC3: Image Captioning by Committee Consensus, Empirical Methods in Natural Language Processing (EMNLP), 2023

Multimodal Pre-Training for Automatic Speech Recognition, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2022

GPU Accelerated t-Distributed Stochastic Neighbor Embedding, Journal of Parallel and Distributed Computing (JPDC), 2019 (**Outstanding Paper Award @ HPML 2018, > 100 citations**)

Going deeper in facial expression recognition using deep neural networks, IEEE Winter conference on applications of computer vision (WACV), 2016 (> 1K citations)

Selected Awards/Grants

BAIR Commons Funded Project, Grant (2021-2023)

Berkeley CTSP, Fellowship (2020)

CITRIS Institute, Tech Innovation Grant (2019)

Herbert J. Greenberg Award for Excellence In Mathematics (2017)

Undergraduate Student Researcher of the Year (2017)

Departmental Service Award (2017)

Professional/Community Service

Berkeley EECS Faculty Search Committee (2024)

Berkeley EECS Graduate Admissions Committee (2020-2022)

Program Committee Member (WACV, CVPR, NeurIPS)

BAIR Graduate Peer Mentor (2021-Present)

BAIR Webmaster (2021-Present)