

KAI SHENG TAI

kst@cs.stanford.edu
<http://kaishengtai.github.io>

LAST UPDATED
JUNE 1, 2021

RESEARCH INTERESTS Learning with limited labeled data, semi-supervised learning, incorporating prior knowledge in neural network architectures, sketching algorithms, representation learning

EDUCATION

2021 (expected) Ph.D. in Computer Science, Stanford University
Thesis: *Statistical Machine Learning Under Resource Constraints*
Advisors: Peter Bailis and Gregory Valiant

2015 M.S. in Computer Science, Stanford University

2013 A.B. in Physics, *magna cum laude*, Princeton University
Thesis: *Detecting Gravitational Waves from Highly Eccentric Compact Binaries*
Advisors: Frans Pretorius and Sean McWilliams

PROFESSIONAL EXPERIENCE

2016– Graduate Research Assistant, Stanford University

2015–2016 Senior Data Scientist, MetaMind (acquired by Salesforce in April 2016)

2014–2015 Research Assistant, Natural Language Processing Group, Stanford University

2014 Software Engineering Intern, Language Technology, Facebook

2013 Software Engineering Intern, Ads Integrity, Facebook

2012 Software Development Engineer Intern, Microsoft

PUBLICATIONS

Kai Sheng Tai, Peter Bailis, and Gregory Valiant. Sinkhorn Label Allocation: Semi-Supervised Classification via Annealed Self-Training. ICML, 2021.

Weiqiang Zhu*, **Kai Sheng Tai***, S. Mostafa Mousavi, Peter Bailis, and Gregory C. Beroza. An End-to-End Earthquake Monitoring Method for Joint Earthquake Detection and Association using Deep Learning. *NeurIPS Workshop on AI for Earth Sciences*, 2020. (*equal contribution)

Kai Sheng Tai, Peter Bailis, and Gregory Valiant. Equivariant Transformer Networks. ICML, 2019.

Vatsal Sharan*, **Kai Sheng Tai***, Peter Bailis, and Gregory Valiant. Compressed Factorization: Fast and Accurate Low-Rank Factorization of Compressively-Sensed Data. ICML, 2019. (*equal contribution)

Edward Gan, Jialin Ding, **Kai Sheng Tai**, Vatsal Sharan, and Peter Bailis. Moment-Based Quantile Sketches for Efficient High Cardinality Aggregation Queries. VLDB, 2018.

Kai Sheng Tai, Vatsal Sharan, Peter Bailis, and Gregory Valiant. Sketching Linear Classifiers over Data Streams. SIGMOD, 2018.

Kai Sheng Tai, Richard Socher, and Christopher D. Manning. Improved Semantic Representations from Tree-Structured Long Short-Term Memory Networks. ACL, 2015.

Kai Sheng Tai, Sean T. McWilliams, and Frans Pretorius. Detecting Gravitational Waves from Highly Eccentric Compact Binaries. *Physical Review D*, 2014.

SERVICE Reviewer for ICML (2019, 2020, 2021), NeurIPS (2019, 2020), JMLR (2020).

SKILLS Proficient in Python, C, C++, Java. Proficient with PyTorch.