

Huan Wang

Education

PhD		2013
	Supervised by Prof. Daniel A. Spielman , Department of Computer Science, Yale University, New Haven, U.S.A.	
M.Phil		2007
	Supervised by Prof. Xiaoou Tang , and Prof. Shuicheng Yan , Multimedia Laboratory, Information Engineering (IE), The Chinese University of Hong Kong, Hong Kong, China	
B. Eng (with honors)		2004
	Mixed Honor Class, Chu Kechen Honors College , Information Science and Electrical Engineering (ISEE), Zhejiang University, Hang Zhou, P. R. China	

Working Experience

Senior Manager, Salesforce Research	2019-
Senior Researcher, Salesforce Research	2018-2019
Senior Applied Researcher, Microsoft	2015-2018
Research Scientist, Yahoo Labs, New York	2013-2015
Adjunct Professor (machine learning CS6923), Computer Science Department, Engineering School, New York University	2014
Adjunct Professor (algorithm design and data structure), MFE program, Baruch College, The City University of New York	2014
Intern at Microsoft Research, Redmond	2011
Intern at Microsoft Research Asian, Beijing	2010
Teaching Fellow of "Graph and Networks", Yale	2010
Teaching Fellow of "Design and Analysis of Algorithms", Yale	2010
Teaching Fellow of "Computer Science and Modern Intellectual Agenda", Yale	2009
Research Assistant [supervised by Prof. Jianzhuang Liu], Multimedia Laboratory, CUHK (Quality Migrant Admission Scheme)	2007
Teaching Assistant of 'Probability Models and Application', CUHK	2007
Teaching Assistant of 'Image and Video Processing', CUHK	2006
Teaching Assistant of 'Multimedia Coding and Processing', CUHK	2006

Fields of Interest

- **Machine Learning**
Dictionary Learning, Compressed Sensing, Graph Theory, Semi-supervised Learning, Subspace Learning, Manifold Analysis, Spectral Analysis, Clustering, Classification, Regression, Deep Learning
- **Big Data Analytics**
Stream Data Analytics, Distributed Learning on top of Apache Spark, Storm, and Hadoop. Big Data Compression and Representation.
- **Computer Vision and Cognition**
Object/Face Representation and Recognition, Image Registration, Segmentation, Age Estimation
- **Data Mining**
Search Engine on Web and Maps
- **Natural Language Processing**
NER, Language Modeling, Knowledge Graph Extraction, Conversational NLP

Research/Engineering Experience*

AI4Software	2021-
<ul style="list-style-type: none">• Massive language models for low-code software development• Risk estimation for code change• Regression forecast based on code difference	
AIOPS	2020-2021
<ul style="list-style-type: none">• Leverage AI to improve Salesforce's operational efficiency• Apply time series forecasting and anomaly detection on Salesforce's Product• Service monitoring and proactive anomaly detection• Traffic forecasting and early planning• Open source repository released on Github (1.9K stars in three weeks, #1 on paperswithcode for two weeks)	
Conversational AI	2020-2021
<ul style="list-style-type: none">• Leading a team to develop AI techniques for Conversational Bots• Develop a natural language understanding engine combining techniques such as NER, intent estimation, knowledge graph construction, sentiment analysis, negation detection, emotion detection, and coreference resolution.	
Multilingual Name Entity Tagging	2019-2020
<ul style="list-style-type: none">• Wrote the first multilingual NER model training/serving pipeline at Salesforce• As the service expands, now it takes a whole team of engineers to maintain and develop	
Language Modeling for ASR	2018-2019
<ul style="list-style-type: none">• Transformer / convolutional network comparison on FairSeq• Integration into the acoustic models	
Accelerated Neural Network Serving using Taylor Expansion and Memorization	2017-2018
<ul style="list-style-type: none">• Using Taylor series to approximate the neural network function• Flatten the deep network into a single "shallow" network	
Rademacher Regularization for Dynamic Dropout Training	2017-2018
<ul style="list-style-type: none">• Connect the dropout rate and Rademacher complexity• Automatic tuning of the dropout rate	
Offline Deep Learning Based Ranker on Bing	2017-2018
<ul style="list-style-type: none">• Expensive offline query-url ranker for higher accuracy• Memorization table for fast service	
AI Question Answer on Bing	2016-2017
<ul style="list-style-type: none">• Concept embedding for better answer accuracy• Bi-LSTM+Attention model for passage ranking	
Address Search Engine (GeoCoder) on Bing Maps	2016-2016
<ul style="list-style-type: none">• Model training for the new generation address search engine (GeoCoder) on Bing Maps• Model shipped with 100% traffic on Bing Maps.	
Category/Entity Recommendation for Bing Local	2015-2016
<ul style="list-style-type: none">• Model training and system design for category/entity recommendation based on user behaviors• On flight on Bing Maps and SERP	
Ads Click Prediction using Neural Network	2014-2015
<ul style="list-style-type: none">• Implemented a simple back-propagation algorithm for the neural network training• On the production data, the algorithm gives the best performance.	
Compromised Account Detection using Random Forest on top of Apache Spark and Storm	2014
<ul style="list-style-type: none">• Implemented a distributed random forest training algorithm on top of Apache Spark framework• Testing stage is implemented on top of apache storm• Encouraging performance on the real-world data.	
Approximate Matrix Multiplication	2013-2014
<ul style="list-style-type: none">• Approximate large matrix multiplication using randomized algorithms• Provable performance guarantee	
Online K-Nearest Neighbor Classifier on top of Apache Storm	2013
<ul style="list-style-type: none">• Designed and implemented an online k-nearest neighbor classifier on top of apache storm.• Locality sensitive hashing and sketching techniques are used to accelerate the prediction	
Dictionary Learning on Large Natural Image Data	2012
<ul style="list-style-type: none">• Designed a monotone dictionary learning algorithm with unbalanced coefficient assumptions• Evaluated the algorithm on natural image data sets	

* Details available at: <http://cs.yale.edu/homes/wang-huan/project.html>

- Anomaly Detection in Bing Clusters** 2011
- Anomaly detection using the log data generated by the Bing cluster
 - Modeled the time series log data using multivariate Brownian Motion
- Matrix Decomposition and Dictionary Learning** 2010
- Proof on the uniqueness of dictionary learning
 - Fast algorithms on the approximate dictionary learning
- Curvature Analysis of Image Manifold** 2010
- Relation between the image derivatives and image manifold
 - An upper bound on the image manifold using image bandwidth
- Linear Regression by Localized Sample Selection** 2009
- A survey on local sample selection for regression
 - Proposed an l_1 based local sample selection algorithm for regression
- Graph Construction and Semi-supervised Learning** 2008
- A global way of graph construction
 - A sparse graph was derived without explicit sparse constraints.
- Correspondence Propagation for Image Registration** 2007
- Designed a transductive algorithm that utilizes prior knowledge to guide the bipartite matching process
 - Derived a closed-form solution that simultaneously preserves feature domain consistency and models geometric distribution.
- Factor Analysis for Image Ensembles** 2007
- Presented a statistical learning technique, the mode-kn factor analysis, to explore image ensembles.
 - Employed statistical Inference for the estimation of pose, illumination and identity
 - Enhanced the classification capability by interacting with the process of synthesizing data
- Misalignment Robust Face Recognition** 2007
- Proposed a misalignment robust framework for subspace learning algorithms to deal with the curse of correspondence problem in face recognition
 - Formulated the misalignment correction process as an L_1 norm optimization.
- Human Age Estimation from Facial Images** 2006—2007
- Took the nonnegative and uncertain properties of the human age into consideration
 - Formulated the age estimation problem as a two-phase semi-definite programming (SDP).
- Semi-supervised Regression on Multi-class/Multi-modality data** 2006—2007
- Derived a transductive procedure for the regression problem over multi-class/multi-modality data
 - Transduced labels across different class samples to pilot the regression.
- Tensor Subspace Analysis for Face Recognition** 2006
- Proposed the first convergent solution to Tensor Subspace Learning algorithms
 - Integrated Bayesian methods in model learning and inference for pose, illumination and identity estimation
 - Evaluated the recognition performance on face databases.
- Discriminant Analysis with Applications in Face Recognition and Data Classification** 2006
- Presented a novel solution that directly optimizes the trace quotient objective
 - Investigated the proposed algorithm systematically on face datasets and machine learning databases.
- Manifold Embedding and Clustering** 2006
- Proposed a spectral analysis algorithm for image clustering
 - Designed a new manifold embedding framework: Maximum Unfolded Embedding.
- Error Control Coding (Convolutional Code, Turbo Code and LDPC Code)** 2003—2004
- Simulated the encoding and decoding process using C++
 - Utilized maximum likelihood (ML/Viterbi) and maximum a posteriori (MAP/BCJR) decoding for the convolutional code.
- Embedded System Design and Implementation (a Digital Fiscal Register)** 2004
- Designed and implemented an embedded system, including hardware, software drivers and Graphic User Interface (GUI), to facilitate fiscal registration to local computer systems using an IC card.
- Micro-Control Unit Design** 2002—2003
- Used Schematic and Verilog-HDL in the design process
 - Four levels of pipelines
 - Implemented using an FPGA.

Open Source Repos:

[Merlion: A Machine Learning Library for Time Series](#)

- 1.9K stars on Github in 3 weeks
- Ranked #1 by [PapersWithCode.com](#) for 2 consecutive weeks

[WarpDrive: Extremely Fast End-to-End Deep Multi-Agent Reinforcement Learning on a GPU](#)

[DataHardness: Evaluating SOTA Classification Models Against Bayes Optimality](#)

Publications*

Conference Papers:

- Aadyot Bhatnagar, Paul Kassianik, Chenghao Liu, Tian Lan, Wenzhuo Yang, Rowan Cassius, Doyen Sahoo, Devansh Arpit, Sri Subramanian, Gerald Woo, Amrita Saha, Arun Kumar Jagota, Gokulakrishnan Gopalakrishnan, Manpreet Singh, K C Krithika, Sukumar Maddineni, Daeki Cho, Bo Zong, Yingbo Zhou, Caiming Xiong, Silvio Savarese Steven Hoi and Huan Wang, 'Merlion: A Machine Learning Library for Time Series', <https://arxiv.org/abs/2109.09265>, 2021. Github Repo: <https://github.com/salesforce/merlion>
- Tian Lan, Sunil Srinivasa, Huan Wang, Stephan Zheng, WarpDrive: Extremely Fast End-to-End Deep Multi-Agent Reinforcement Learning on a GPU, <https://arxiv.org/abs/2108.13976>, 2021. Github Repo: <https://github.com/salesforce/warp-drive>
- Yu Bai, Song Mei, Huan Wang, Caiming Xiong, 'Understanding the Under-Coverage Bias in Uncertainty Estimation', **NeurIPS**, 2021.
- Ryan Theisen, Huan Wang, Lav R Varshney, Caiming Xiong, Richard Socher, 'Evaluating State-of-the-Art Classification Models Against Bayes Optimality', **NeurIPS**, 2021. Github Repo: <https://github.com/salesforce/DataHardness>
- Yu Bai, Chi Jin, Huan Wang, Caiming Xiong, 'Sample-Efficient Learning of Stackelberg Equilibria in General-Sum Games', **NeurIPS**, 2021.
- Tengyang Xie, Nan Jiang, Huan Wang, Caiming Xiong, Yu Bai, 'Policy Finetuning: Bridging Sample-Efficient Offline and Online Reinforcement Learning', **NeurIPS**, 2021.
- Rachel Luo, Aadyot Bhatnagar, Huan Wang, Caiming Xiong, Silvio Savarese, Yu Bai, Shengjia Zhao, Stefano Ermon, 'Localized Calibration: Metrics and Recalibration', <https://arxiv.org/abs/2102.10809>, 2021.
- Tong Niu, Semih Yavuz, Yingbo Zhou, Nitish Shirish Keskar, Huan Wang and Caiming Xiong, 'Unsupervised Paraphrasing with Pretrained Language Models', **EMNLP**, 2021.
- Wenpeng Yin, Huan Wang, Jin Qu, Caiming Xiong, 'BatchMixup: Improving Training by Interpolating Hidden States of the Entire Mini-batch', **ACL Findings**, 2021.
- Yu Bai, Song Mei, Huan Wang, Caiming Xiong, 'Don't Just Blame Over-parametrization for Over-confidence: Theoretical Analysis of Calibration in Binary Classification', **ICML**, 2021.
- Stanislaw Jastrzebski, Devansh Arpit, Oliver Astrand, Giancarlo Kerg, Huan Wang, Caiming Xiong, Richard Socher, Kyunghyun Cho, Krzysztof Geras, 'Catastrophic Fisher Explosion: Early Phase Fisher Matrix Impacts Generalization', **ICML**, 2021.
- Tong Niu, Semih Yavuz, Yingbo Zhou, Huan Wang, Nitish Shirish Keskar, Caiming Xiong, 'Unsupervised Paraphrase Generation via Dynamic Blocking', <https://arxiv.org/abs/2010.12885>, 2021.
- Yu Bai, Minshuo Chen, Pan Zhou, Tuo Zhao, Jason D Lee, Sham Kakade, Huan Wang, Caiming Xiong, 'How Important is the Train-Validation Split in Meta-Learning?', **ICML**, 2021.
- Minshuo Chen, Yu Bai, Jason D Lee, Tuo Zhao, Huan Wang, Caiming Xiong, Richard Socher, 'Towards understanding hierarchical learning: Benefits of neural representations', **NeurIPS**, 2020.
- Huan Wang, Nitish Shirish Keskar, Caiming Xiong, Richard Socher, 'Assessing Local Generalization Capability in Deep Models', International Conference on Artificial Intelligence and Statistics (**AISTATS**), 2020.
- Huan Wang, Stephan Zheng, Caiming Xiong, Richard Socher, 'On the Generalization Gap in Reparameterizable Reinforcement Learning', International Conference On Machine Learning (**ICML**), 2019.
- Peiliang Zhang, Huan Wang, Nikhil Naik, Caiming Xiong, Richard Socher. DIME: An Information-Theoretic Difficulty Measure for AI Datasets, Open Review, 2019. <https://openreview.net/forum?id=kvqPFy0hbF¬Id=5H2fhISQwz1>
- Huan Wang, Nitish Shirish Keskar, Caiming Xiong, Richard Socher, 'Identifying Generalization Properties in Neural Networks', NIPS workshop on Integration of Deep Learning Theories / Arxiv, 2018. [\[Blog\]](#)
- Ke Zhai, Huan Wang, 'Adaptive Dropout Training with Rademacher Complexity Regularization', **ICLR** 2018. [Authors contribute equally]
- Huan Wang, Christos Boutsidis, Edo Liberty, Daniel Hsu, 'Fast Matrix Multiplication via One-Side Element-wise Sparsification', preprint available upon request, 2015.

- Huan Wang, John Wright, and Daniel Spielman, 'A Batchwise Monotone Algorithm for Dictionary Learning', arxiv, 2015.
- Daniel Spielman, Huan Wang, and John Wright, 'Exact Recovery of Sparse-Used Dictionaries', **Best paper award** of the 25th Conference on Learning Theory (COLT), Jun.2012. [authors are ranked in alphabetical order in COLT]
- Shuicheng Yan, and Huan Wang, 'Semi-supervised Learning by Sparse Representation', SIAM International Conference on Data Mining (SDM) Apr. 2009.
- Huan Wang, Shuicheng Yan, Thomas Huang and Xiaoou Tang, 'A Convergent Solution to Tensor Subspace Learning', International Joint Conferences on Artificial Intelligence (IJCAI 07 Oral presentation), Jan. 2007.
- Huan Wang, Shuicheng Yan, Jianzhuang Liu, Thomas Huang and Xiaoou Tang, 'Misalignment Robust Face Recognition', IEEE Conference on Computer Vision and Pattern Recognition (CVPR 08), Jun. 2008.
- Huan Wang, Shuicheng Yan, Thomas Huang and Xiaoou Tang, 'Trace Ratio vs. Ratio Trace for Dimensionality Reduction', IEEE Conference on Computer Vision and Pattern Recognition (CVPR 07), Jun. 2007.
- Huan Wang, Shuicheng Yan, Thomas Huang, Jianzhuang Liu and Xiaoou Tang, 'Transductive Regression Piloted by Inter-Manifold Relations', International Conference on Machine Learning (ICML07 Oral presentation), Jun. 2007.
- Shuicheng Yan, Huan Wang, Thomas Huang and Xiaoou Tang, 'Learning Auto-Structured Regressor from Uncertain Labels', International Conference on Computer Vision (ICCV 07), Oct. 2007.
- Huan Wang, Shuicheng Yan, Thomas Huang and Xiaoou Tang, 'Maximum Unfolded Embedding: Formulation, Solution, and Application for Image Clustering', ACM International Conference on Multimedia (ACM SIGMM06), Oct. 2006.
- Shuicheng Yan, Huan Wang, Thomas Huang and Xiaoou Tang, 'Ranking with Uncertain Labels', IEEE International Conference on Multimedia & Expo (ICME07 Oral presentation), May. 2007.
- Shuicheng Yan, Huan Wang, Xiaoou Tang and Thomas Huang, 'Exploring Feature Descriptors for Face Recognition', IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP07 Oral presentation, ~15%), Apr. 2007.

Journals:

- Daniel Spielman, Huan Wang, and John Wright, 'Exact Recovery of Sparse-Used Dictionaries', (Arxiv), 2012.
- Shuicheng Yan, Huan Wang, Jianzhuang Liu, Xiaoou Tang, and Thomas S. Huang, 'Ranking with Uncertain Labels and Its Applications', Frontiers of Computer Science in China (Journal), 2007.
- Shuicheng Yan, Huan Wang, Xiaoou Tang, and Thomas S. Huang, 'Learning Auto-Structured Regressor from Uncertain Labels', IEEE Transactions on Information Forensics and Security (TIFS), 2008
- Huan Wang, Shuicheng Yan, Jianzhuang Liu, Thomas Huang and Xiaoou Tang, 'Correspondence Propagation with Weak Priors', IEEE Transaction on Image Processing (TIP), 2008
- Shuicheng Yan, Huan Wang, Xiaoou Tang, and Thomas Huang, 'Synchronized Submanifold Embedding for Person-Independent Pose Estimation and Beyond', IEEE Transaction on Image Processing (TIP), 2008.
- Shuicheng Yan, Huan Wang, Jilin Tu, Thomas S. Huang, and Xiaoou Tang, 'Mode-kn Factor Analysis for Image Ensembles', IEEE Transactions on Image Processing(TIP), 2008.

Thesis:

- Huan Wang, 'Dictionary Learning: Algorithms and Analysis', Computer Science Department, Yale University. (Ph.D Thesis)
- Huan Wang, 'Exploring Intrinsic Structures from Samples: Supervised, Unsupervised, and Semisupervised Frameworks', Department of Information Engineering, the Chinese University of Hong Kong. (M.Phil Thesis)
- Huan Wang, 'Design and Implementation of an Embedded Fiscal Register', Chu Kechen Honors College, Zhejiang University. (B.Eng Thesis)

U.S. Patents:

- Devansh Arpit, Huan Wang, Caiming Xiong, Learning Rich Nearest Neighbor Representations from Self-Supervised Ensembles, SFDC//102663/US, 2021
- Huan Wang, Caiming Xiong, Bram Wallace, Devansh Arpit, SYSTEMS AND METHODS FOR SELF-SUPERVISED TRAINING OF A CONVOLUTIONAL NEURAL NETWORK, 2021.
- Tong Niu, Huan Wang, Caiming Xiong. SYSTEMS AND METHODS FOR UNSUPERVISED PARAPHRASE GENERATION, U.S. patent application Serial Number 63/104,988, October 23, 2020.
- Devansh Arpit, Huan Wang, Caiming Xiong. NEURAL NETWORK BASED ANOMALY DETECTION FOR TIME-SERIES DATA, 31371-48298/US, Apr. 15, 2020.

- Wenhao Liu; Xinyi Yang; Tian Xie; Caiming Xiong; Huan Wang; Kazuma Hashimoto; Jin Qu; Feihong Wu; Yingbo Zhou. Customizing Chatbots using Policy Based Specification, 31371-46808/US, Sep 23, 2020.
- Wenhao Liu; Xinyi Yang; Tian Xie; Caiming Xiong; Huan Wang; Kazuma Hashimoto; Yingbo Zhou; Xugang Ye; Jin Qu; Feihong Wu. CONFIGURABLE CONVERSATION ENGINE FOR EXECUTING CUSTOMIZABLE CHATBOTS, 31371-46809/US, Sep 30, 2020.
- Tian Xie; Wenhao Liu; Xinyi Yang; Caiming Xiong; Huan Wang; Wenpeng Yin; Jin Qu. INTENT RESOLUTION FOR CHATBOT CONVERSATIONS WITH NEGATION AND COREFERENCES, 31371-46811/US, Aug 14, 2020.
- Wenhao Liu; Xinyi Yang; Tian Xie; Caiming Xiong; Huan Wang; Jin Qu; Soujanya Lanka; Chu Hong Hoi; Xugang Ye; Feihong Wu. ENTITY RESOLUTION FOR CHATBOT CONVERSATIONS, 31371-46810/US, Aug 14, 2020.

Academic Service

Reviewer:

Journal of Machine Learning Research (JMLR)
International Conference on Artificial Intelligence and Statistics (AISTATS)
International Conference on Machine Learning (ICML)
Neural Information Processing Systems Foundation (NIPS)
Association for the Advancement of Artificial Intelligence (AAAI)
IEEE Transactions On Image Processing
Neurocomputing (Elsevier)
Pattern Recognition Letters (Elsevier)
IEEE Trans. on Circuits and Systems for Video Technology
IEEE Trans. on Systems, Man and Cybernetics, Part B
International Journal of Computer Mathematics

Organizer:

"Big Data and Stream Analytics" Workshop @ SIAM Conference on Data Mining 2015

Honors & Awards

Chinese Government Award for Outstanding Self-Financed Students Abroad	2013
Best Paper Award at Conference Of Learning Theory (COLT)	2012
Award of Excellence (Stars of Tomorrow), Microsoft Research Asia	2010
Fellowship award, Yale University	2008
Studentship award, Chinese University of Hong Kong	2005-2007
Bachelor's degree with Honors , Zhejiang University	2004
Excellent B.Eng Thesis , Zhejiang University	2004
Yongqian Tang Outstanding Student Exchange Fellowship , Zhejiang Univ.	2003
Champion, University Chorus Competition , Zhejiang University	2002
Excellence award for moral, intellective and physical merits, Zhejiang University	2001
Outstanding Student Scholarship , Zhejiang University	2000-2003
Freshman Award for Academic Excellence , Zhejiang University	2000
Honorary Enrollment , Zhejiang University	2000

Extracurricular Activities & Global Experience

Invited presentation on IJCAI13 track on Best Papers in Sister Conferences	2013
Best paper presentation on COLT12 , Edinburgh, Britain	2012
Oral presentation on ICML07, Corvallis, OR, USA	2007
Oral presentation on IJCAI07, Hyderabad, India	2007
Cultural exchange, University of Leeds, Britain	2003
Piano accompanist for the college chorus	2001

Engineering Skills

Software Programming

C/C++/C#/Matlab/Java/Python, Pytorch/Tensorflow/Keras, Hadoop/Spark/Storm, MPI Programming, DirectX
Game Programming

Hardware Programming

Verilog HDL / VHDL, CPLD, FPGA programming, Single Chip Programming

Hardware Design

Printed Circuit Board (PCB) Design

Music Related

Piano Playing: Level 10 (Topmost Assessment by Chinese Ministry of Education)

Music Composition: my original song "[A Ripple of Love](#)" has **millions** of online playbacks