

Peng Qi

121 Campus Dr., Apt. 3202B, Stanford, CA 94305
+1 (650) 862-9799 • pengqi@cs.stanford.edu • <http://qipeng.me/>

EDUCATION & PROFESSIONAL EXPERIENCES

Stanford University

PhD Student, Computer Science Sep. 2015 – present

- Research Assistant with Stanford Natural Language Processing Group (Advisor: Chris Manning)

Master of Science, Computer Science Sep. 2013 – Jun. 2015

- Research Assistant with Stanford Artificial Intelligence Laboratory (Advisor: Andrew Ng)

Master Student, Department of Statistics Apr. 2016 – present

Tsinghua University

Research Assistant, State Key Laboratory of Intelligent Technology & Systems (Advisor: Xiaolin Hu) Jul. 2011 – Jun. 2013

Bachelor of Engineering (magna cum laude), School of Software Sep. 2008 – Jul. 2012
Major GPA: 89.9 / 100 Major Rank: Top 5%

PUBLICATIONS

Andrew L. Maas, **Peng Qi**, Ziang Xie, Awni Y. Hannun, Daniel Jurafsky, and Andrew Y. Ng. Building DNN Acoustic Models for Large Vocabulary Speech Recognition. *Computer Speech & Language*, 2016.

Peng Qi and Xiaolin Hu. Learning nonlinear statistical regularities in natural images by modeling the outer product of image intensities. *Neural computation*, 26(4):693–711, 2014.

Xiaolin Hu, Jianwei Zhang, **Peng Qi**, and Bo Zhang. Modeling response properties of V2 neurons using a hierarchical k-means model. *Neurocomputing*, 134:198–205, 2014.

Peng Qi, Shuochen Su, and Xiaolin Hu. Modeling outer products of features for image classification. In *Advanced Computational Intelligence (ICACI)*, 2013.

Xiaolin Hu, **Peng Qi**, and Bo Zhang. Hierarchical k-means algorithm for modeling visual area V2 neurons. In *Neural Information Processing*, pages 373–381, 2012. *Best Paper Award*

EXPERIENCE

Stanford A.I. Lab Sep. 2013 – Jun. 2015

Large-Vocabulary Continuous Speech Recognition with Deep Learning

- Built an HMM-GMM speech recognition system on the largest speech corpus in academia (Fisher+Switchboard 2100-hour corpus) in Kaldi, and prepared the first evaluation pipeline for the RT-03 evaluation set in Kaldi
- Analyzed performances and various properties as DNN acoustic models scale up
- Investigated the effect of incorporating auxiliary information and varying network structure (CS229 Project [↗](#))

CS 229T Statistical Learning Theory Course Project Jan. 2013 – Feb. 2015

Empirically Efficient Methods for Training Deep Neural Networks

- Implemented and investigated the training efficiency of various adaptive gradient methods and accelerated gradient methods applied to deep neural networks

CS 224U Natural Language Understanding Course Project Mar. 2013 – present

Learning Unsupervised Semantic Word Vectors for Sentence Completion

- Implemented various unsupervised semantic word vector models and applied them to sentence completion
- Proposed a novel neural language model extending `word2vec` for sentence completion

Contributing Open Source Projects

Caffe [↗](#)

- Implemented a framework for generic solvers
- Implemented Nesterov's Accelerated Gradient solver and AdaGrad solver
- Contributed a number of neuron layers (Leaky ReLU, Mean-variance normalization)
- Compatibility issues / bug fixes for Mac OS

Kaldi [↗](#)

- Contributed the first training recipe for the Fisher/Switchboard mixed speech corpus (the largest speech corpus in use in academia)
- A bug fix for Kaldi's speaker identification for better speaker heldout training

ConvolutionalRBM.m [↗](#) (Owner)

- An implementation of Lee *et al.*'s convolutional restricted Boltzmann machine model in Matlab, MEX (C++/CUDA)

GraphCut [↗](#)

- A fast CPU/GPU hybrid implementation for min-cut problems in grid graphs

ACADEMIC SERVICE *Teaching:*

- TA, CS 224D Deep Learning for Natural Language Processing (Spring 2015)
- TA, CS 124 From Languages to Information (Winter 2015)
- TA, CS 224S Spoken Language Processing (Spring 2014)
- TA, CS 145 Introduction to Databases (Summer 2014)
- Tutor, CS 145 Introduction to Databases, CS 107 Computer Organization and Systems, CS 245 Database Systems Principles

Reviewer: IEEE TNNLS, ICACI 2013

MISCELLANEOUS

Honors: China's National Scholarship (top 3% university-wide at Tsinghua), Freshman Scholarship (provincial top 10 in college entrance exam), and other merit-based awards

Programming: Experienced with C/C++, CUDA, Matlab/MEX, Python, Java; Working knowledge of *nix Bash, Javascript (node.js/jQuery), Haskell, PHP, C#, HTML5/CSS3, Perl, etc.

Leadership: Vice president of Students' Union, School of Software, Tsinghua University (2010 – 2011)

Part-time activities: Chinese Translation [↗](#) of Philip J. Guo's *The Ph.D. Grind* [↗](#)