

Peng Qi

353 Serra Mall, Gates Computer Science, 2A Wing, Stanford, CA 94305
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 of Science, Department of Statistics Apr. 2016 – Mar. 2017

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

PUBLICATIONS (* = Equal Contribution)

Zhilin Yang*, **Peng Qi***, Saizheng Zhang*, Yoshua Bengio, William W. Cohen, Ruslan Salakutdinov, and Christopher D. Manning. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.

Yuhao Zhang*, **Peng Qi***, and Christopher D. Manning. Graph Convolution over Pruned Dependency Trees Improves Relation Extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.

Peng Qi*, Timothy Dozat*, Yuhao Zhang*, and Christopher D. Manning. Universal Dependency Parsing from Scratch. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, 2018.

Urvashi Khandelwal, He He, **Peng Qi**, and Dan Jurafsky. Sharp Nearby, Fuzzy Far Away: How Neural Language Models Use Context. In *Proceedings of the 56th Annual Conference of Association of Computational Linguistics (ACL)*, 2018.

Peng Qi and Christopher D. Manning. Arc-swift: A Novel Transition System for Dependency Parsing. In *Proceedings of the 55th Annual Conference of Association of Computational Linguistics (ACL)*, 2017.

Timothy Dozat, **Peng Qi**, and Christopher D. Manning. Stanford's Graph-based Neural Dependency Parser at the CoNLL 2017 Shared Task. *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. **First place**

Arun Chaganty*, Ashwin Paranjape*, Jason Bolton*, Matthew Lamm*, Jinhao Lei*, Abigail See*, Kevin Clark, Yuhao Zhang, **Peng Qi**, and Christopher D. Manning. Stanford at TAC KBP 2017: Building a Trilingual Relational Knowledge Graph. In *Proceedings of the Text Analysis Conference (TAC)*, 2017.

Yuhao Zhang*, Arun Chaganty*, Ashwin Paranjape*, Danqi Chen*, Jason Bolton*, **Peng Qi**, and Christopher D. Manning. Stanford at TAC KBP 2016: Sealing Pipeline Leaks and Understanding Chinese. In *Proceedings of the Text Analysis Conference*

(TAC) - Knowledge Base Population, 2016.

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 NLP Group Sep. 2015 – present
Universal Dependency Parsing from Scratch

- One unified neural model trained end-to-end for word tokenization, sentence segmentation, part-of-speech tagging, and dependency parsing

Explainable Multi-hop Question Answering in the Wild (work-in-progress, Facebook’s ParlAI Research Awardee)

- Collecting a question answering dataset focusing on multi-hop factoid questions based on Wikipedia
- Proposed a cost-effective method for efficient crowdsourcing of high-quality data

Facebook AI Research (New York) Jun. 2017 – Sep. 2017
Learning to Teach through Communication (Mentor: Jason Weston)

- Studied emergent teaching behavior of machine learning agents in a constrained communication setting
- Proposed a reinforcement learning-based method for agents to learn to teach
- Implemented the proposed method in ParlAI with positive results on image classification tasks

Contributing Open Source Projects

StanfordNLP [↗](#) (documentation [↗](#))

- PyTorch implementation of Stanford’s full system in the 2018 CoNLL Shared Task on Universal Dependency Parsing from Raw Text
- Includes neural network models for tokenization, part-of-speech tagging, lemmatization, and dependency parsing in 50+ languages with pretrained models available and intuitive Python interface
- Main contributor and maintainer

ParlAI [↗](#)

- Implemented various features to enable teacher training in ParlAI
- Bugfixes and stability improvements

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)

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: EMNLP 2018 (Best Reviewer Award), ACL 2018, ACL 2017, IEEE TNNLS, ICACI 2013

MISCELLANEOUS

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

Programming: Experienced with Python, Java, C/C++, CUDA, Matlab/MEX, Lua, Bash; Working knowledge of Javascript, Haskell, PHP, C#, HTML5/CSS3, Perl, etc.