

# Hao Li

4431 A.V.Williams Build  
College Park, MD 20742

+1-240-485-7596  
✉ [haoli@cs.umd.edu](mailto:haoli@cs.umd.edu)  
📄 <http://cs.umd.edu/~haoli>  
🌐 <http://github.com/ljk628>

## Education

<b>University of Maryland, College Park</b> <i>Ph.D. in Computer Science, Advisors: Prof. Hanan Samet and Tom Goldstein</i>	College Park, MD 2012 – 2017
<b>Institute of Computing Technology, Chinese Academy of Sciences</b> <i>M.S. in Computer Science, Advisor: Prof. Shiguang Shan</i>	Beijing, China 2009 – 2012
<b>Shandong University</b> <i>B.Eng. in Software Engineering</i>	Jinan, China 2005 – 2009

## Research Interests

Fast and Scalable Machine Learning, Deep Learning, HPC, Distributed System

## Experience

<b>University of Maryland Institute for Advanced Computer Studies</b> <i>Research Assistant with Prof. Hanan Samet and Prof. Tom Goldstein</i>	College Park, MD May 2013 - present
– Fast and scalable optimization methods for deep learning and reinforcement learning.	
– Near-duplicate image detection and streaming news image summarization.	
<b>Machine Learning Team, Nervana Systems</b> <i>Algorithms Engineer Intern</i>	San Diego, CA May - Aug 2016
– Developed a sub-module for Neon’s graph-backend, enabling executing the computation graph of TensorFlow with Neon.	
<b>Machine Learning Department, NEC Labs America</b> <i>Research Intern/Visiting Student</i>	Princeton, NJ Jan-Aug 2015/Jan-May 2016
– Research on fast and efficient convolutional neural networks.	
– Research on filter pruning for acceleration and compression of CNNs.	
<b>Storage System Department, NEC Labs America</b> <i>Research Intern</i>	Princeton, NJ May – Aug 2014
– Built MALT, a distributed data-parallel machine learning system with flexible peer-to-peer communication, utilizing one-side RDMA write over InfiniBand.	
<b>Institute of Computing Technology, Chinese Academy of Sciences</b> <i>Research Assistant with Prof. Shiguan Shan</i>	Beijing, China Jan 2011 - Jul 2012
– Multi-Person Tracking with Re-identification in Non-overlapping Camera Networks.	
<b>Internet Graphics Group, Microsoft Research Asia</b> <i>Research Intern with Ying-Qing Xu</i>	Beijing, China May – Nov 2010
– Bing Image Search Result Beautification.	
<b>Media Computing Group, Microsoft Research Asia</b> <i>Research Intern with Meng Wang</i>	Beijing, China Mar – Jul 2009
– Multimodal Graph-Based Reranking for Web Image Search.	

## Publications

- **Hao Li\***, Soham De\*, Zheng Xu, Christoph Studer, Hanan Samet, Tom Goldstein. *Towards a Deeper Understanding of Training Quantized Neural Networks*. ICML Workshop on Principled Approaches to Deep Learning (**PADL**), 2017 (**Google Best Student Paper**)
- Zheng Xu, Gavin Taylor, **Hao Li**, Mario Figueiredo, Xiaoming Yuan, and Tom Goldstein. *Adaptive Consensus ADMM for Distributed Optimization*. International Conference on Machine Learning, (**ICML**), 2017
- **Hao Li**, Asim Kadav, Igor Durdanovic, Hanan Samet, Hans Peter Graf. *Pruning Filters for Efficient ConvNets*. International Conference on Learning Representations (**ICLR**), 2017  
NIPS Workshop on Efficient Methods for Deep Neural Networks (**EMDNN**), 2016
- **Hao Li**, Shangfu Peng, Hanan Samet. *Streaming News Image Summarization*. International Conference on Pattern Recognition (**ICPR**), 2016 (Oral)
- **Hao Li**, Asim Kadav, Erik Kruus, Cristian Ungureanu. *MALT: Distributed Data-Parallelism for Existing ML Applications*. ACM European Conference on Computer Systems (**EuroSys**), 2015  
NIPS Workshop on Machine Learning Systems (**LearningSys**), 2015
- Shangfu Peng, Hong Wei, **Hao Li**, and Hanan Samet. *Simplification and Refinement for Speedy Spatio-temporal Hot Spot Detection Using Spark (GIS Cup)*. ACM Int. Conf. on Advances in Geographic Information Systems (**SIGSPATIAL**), 2016 (Short)
- Canxiang Yang, Shiguang Shan, Dan Wang, **Hao Li**, Xilin Chen. *View Adaptive Metric Learning for Multi-view Person Re-identification*. Asian Conference on Computer Vision (**ACCV**), 2014
- Meng Wang, **Hao Li**, Dacheng Tao, Ke Lu, Xindong Wu. *Multimodal Graph-Based Reranking for Web Image Search*. IEEE Transactions on Image Processing (**T-IP**), 2012
- **Hao Li**, Meng Wang, Zhisheng Li, Zheng-Jun Zha, Jialie Shen. *Optimizing Multimodal Reranking for Web Image Search*. ACM SIGIR Conference on Research and Development in Information (**SIGIR**), 2011 (Short)
- **Hao Li**, Meng Wang, Xian-Sheng Hua. *MSRA-MM 2.0: A Large-Scale Web Multimedia Dataset*. IEEE Int. Conference on Data Mining (**ICDM**) Workshop on Internet Multimedia Mining, 2009

## Awards

Google Best Student Paper Award, ICML PADL Workshop	2017
ICML Student Travel Award	2017
ICLR Student Travel Award	2017
Jacob K. Goldhaber Travel Grant, University of Maryland	2016
ACM EuroSys Student Travel Grant	2015
John D. Gannon Student Travel Grant, University of Maryland	2013
Dean's Fellowship, Department of Computer Science, University of Maryland	2012/13
ACM SIGIR Student Travel Grant	2011
Stars of Tomorrow Excellent Intern Award, Microsoft Research Asia	2010
Outstanding College Graduate of Shandong Province and Shandong University	2009
Meritorious Winner in Mathematical Contest in Modeling, USA	2008
National Second Prize in China Undergraduate Mathematical Contest in Modeling	2007

## Professional Services

Journal Reviewer for Neurocomputing'15, Multimedia Systems'14, IEEE T-Industrial Informatics'15  
Conference Reviewer for NIPS 2016, CIKM 2014, ICME 2013/2014

## Programming Skills

Language: Python, Lua, C/C++, Java, C#, Matlab, Go

Tools: Torch7, Pytorch, TensorFlow, Neon, MPI, GASPI, ZMQ, Hadoop

## Teaching Assistant Experience

CMSC 764 Advanced Numerical Optimization, with Prof. Tom Goldstein	Spring 2017
CMSC 828Z Reinforcement Learning, with Prof. Hal Daumé	Fall 2016
CMSC 216 Introduction to Computer Systems, with Dr. Larry Herman	Fall 2015
CMSC 426 Programming Handheld Systems, with Prof. Atif Memon	Spring 2013
CMSC 330 The Organization of Programming Languages, with Prof. Chau-Wen Tseng	Fall 2012/13

## Patents

Parallelized Machine Learning With Distributed Lockless Training. US Patent 20,160,103,901  
MALT: Distributed Data-Parallelism for Existing ML Applications. US Patent 20,160,125,316