

# Hexiang (Frank) Hu

1002 Childs Way, Los Angeles, CA 90089  
hexiang.frank.hu@gmail.com  
<http://hexiang-hu.github.io>

---

<b>EDUCATIONS</b>	<b>University of Southern California</b> <i>Ph.D.</i> in Computer Science Advisor: Fei Sha	2017 - Present
	<b>University of California, Los Angeles (UCLA)</b> <i>Ph.D.</i> in Computer Science (Transfer) Advisor: Fei Sha	2016 - 2017
	<b>Simon Fraser University</b> <i>Bachelor of Science</i> in Computer Science First class with distinction	2013 - 2016
	<b>Zhejiang University</b> <i>Bachelor of Engineering</i> in Computer Science and Technology	2011 - 2013
<b>EMPLOYMENT</b>	<b>Research Assistant</b> <i>University of Southern California</i>	Jan 2017 - Present Los Angeles, CA
	<b>Research Intern</b> <i>Google Research</i>	Feb 2020 - Dec 2020 Mountain view, CA
	<b>Research Intern</b> <i>Intel Research Labs</i>	May 2019 - Dec 2019 Santa Clara, CA
	<b>Research Intern</b> <i>Facebook AI Research</i>	May 2018 - Dec 2018 New York, NY
	<b>Applied Scientist Intern</b> <i>Amazon Research</i>	May 2017 - Aug 2017 Palo Alto, CA
	<b>Research Assistant</b> <i>University of California, Los Angeles</i>	Sept 2016 - Dec 2016 Los Angeles, CA
	<b>Undergraduate Research Assistant</b> <i>Vision and Media Lab at SFU</i>	May 2015 - Dec 2015 Vancouver, BC
	<b>Software Engineer Intern</b> <i>Microsoft</i>	May 2014 - Aug 2014 Vancouver, BC
<b>PUBLICATIONS</b>	Bowen Zhang*, <b>Hexiang Hu*</b> , Vihan Jain, Eugene Ie, and Fei Sha. "Learning to Represent Image and Text with Denotation Graph". In: <i>EMNLP</i> . 2020	
	Wang Zhu*, <b>Hexiang Hu*</b> , Jiacheng Chen, Zhiwei Deng, Vihan Jain, Eugene Ie, and Fei Sha. "BabyWalk: Going Farther in Vision-and-Language Navigation by Taking Baby Steps". In: <i>ACL</i> . 2020	
	Han-Jia Ye, <b>Hexiang Hu</b> , De-Chuan Zhan, and Fei Sha. "Few-shot learning via embedding adaptation with set-to-set functions". In: <i>CVPR</i> . 2020	
	Risto Vuorio*, Shao-Hua Sun*, <b>Hexiang Hu</b> , and Joseph J Lim. "Multimodal Model-Agnostic Meta-Learning via Task-Aware Modulation". In: <i>NeurIPS</i> . 2019	

---

\*: Authors contributed equally

**Hexiang Hu**, Ishan Misra, and Laurens van der Maaten. “Evaluating text-to-image matching using binary image selection (BISON)”. in: *ICCV Workshop*. 2019

Bowen Zhang, Hexiang Hu, and Fei Sha. “Visual Storytelling via Predicting Anchor Word Embeddings in the Stories”. In: *ICCV Workshop*. 2019

Nelson Nauata, **Hexiang Hu**, Guang-Tong Zhou, Zhiwei Deng, Zicheng Liao, and Greg Mori. “Structured label inference for visual understanding”. In: *T-PAMI* (2019)

Kurt Shuster, Samuel Humeau, **Hexiang Hu**, Antoine Bordes, and Jason Weston. “Engaging image captioning via personality”. In: *CVPR*. 2019

**Hexiang Hu\***, Liyu Chen\*, Boqing Gong, and Fei Sha. “Synthesized policies for transfer and adaptation across tasks and environments”. In: *NeurIPS*. 2018

Soravit Changpinyo, **Hexiang Hu**, and Fei Sha. “Multi-task learning for sequence tagging: An empirical study”. In: *COLING*. 2018

Bowen Zhang\*, **Hexiang Hu\***, and Fei Sha. “Cross-modal and hierarchical modeling of video and text”. In: *ECCV*. 2018

**Hexiang Hu\***, Wei-Lun Chao, and Fei Sha. “Learning answer embeddings for visual question answering”. In: *CVPR*. 2018

Wei-Lun Chao\*, **Hexiang Hu\***, and Fei Sha. “Cross-dataset adaptation for visual question answering”. In: *CVPR*. 2018

Wei-Lun Chao\*, **Hexiang Hu\***, and Fei Sha. “Being negative but constructively: Lessons learnt from creating better visual question answering datasets”. In: *NAACL-HLT*. 2018

Chao-Yuan Wu, Manzil Zaheer, **Hexiang Hu**, R Manmatha, Alexander J Smola, and Philipp Krähenbühl. “Compressed video action recognition”. In: *CVPR*. 2018

**Hexiang Hu\***, **Shiyi Lan\***, Yuning Jiang, Zhimin Cao, and Fei Sha. “Fastmask: Segment multi-scale object candidates in one shot”. In: *CVPR*. 2017

**Hexiang Hu**, Guang-Tong Zhou, Zhiwei Deng, Zicheng Liao, and Greg Mori. “Learning structured inference neural networks with label relations”. In: *CVPR*. 2016

Zhiwei Deng, Arash Vahdat, **Hexiang Hu**, and Greg Mori. “Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition”. In: *CVPR*. 2016

Chao-Yuan Wu, Manzil Zaheer, **Hexiang Hu**, R Manmatha, Alexander J Smola, and Philipp Krähenbühl. “Compressed video action recognition”. In: *CVPR*. 2018

**Hexiang Hu\***, **Shiyi Lan\***, Yuning Jiang, Zhimin Cao, and Fei Sha. “Fastmask: Segment multi-scale object candidates in one shot”. In: *CVPR*. 2017

**Hexiang Hu**, Guang-Tong Zhou, Zhiwei Deng, Zicheng Liao, and Greg Mori. “Learning structured inference neural networks with label relations”. In: *CVPR*. 2016

Zhiwei Deng, Arash Vahdat, **Hexiang Hu**, and Greg Mori. “Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition”. In: *CVPR*. 2016

## TECHNICAL REPORTS

**Hexiang Hu**, Ozan Sener, Fei Sha, and Vladlen Koltun. “Drinking from a Firehose: Continual Learning with Web-scale Natural Language”. In: *arXiv preprint arXiv:2007.09335* (2020)

Bowen Zhang\*, **Hexiang Hu**\*, Joonseok Lee\*, Ming Zhao, Sheide Chammas, Vishan Jain, Eugene Ie, and Fei Sha. “A Hierarchical Multi-Modal Encoder for Moment Localization in Video Corpus”. In: *arXiv preprint arXiv:2011.09046* (2020)

Jiacheng Chen\*, **Hexiang Hu**\*, Hao Wu, Yuning Jiang, and Changhu Wang. “Learning the Best Pooling Strategy for Visual Semantic Embedding”. In: *arXiv preprint arXiv:2011.04305* (2020)

**Hexiang Hu**\*, Zhiwei Deng\*, Guang-Tong Zhou, Fei Sha, and Greg Mori. “Labelbank: Revisiting global perspectives for semantic segmentation”. In: *arXiv preprint arXiv:1703.09891* (2017)

## PATENTS

R Manmatha, **Hexiang Hu**, and Deva Ramanan. *Compressed content object and action detection*. US Patent 10,528,819. Jan. 2020

## SERVICES

### Conference Reviewer

- Conference on Neural Information Processing Systems (NeurIPS)
- International Conference on Machine Learning (ICML)
- International Conference on Learning Representation (ICLR)
- Annual Meeting of the Association for Computational Linguistics (ACL)
- The North American Chapter of the ACL (NAACL)
- IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
- International Conference on Computer Vision (ICCV)
- European Conference on Computer Vision (ECCV)
- Asian Conference on Computer Vision (ACCV)
- Winter Conference on Computer Vision (WACV)
- AAAI Conference on Artificial Intelligence (AAAI)

### Journal Reviewer

- IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)

## TECHNICAL TALKS

“Understanding Language in Perception and Embodiment”

- Google, Mountain View, CA Feb 2021
- University of Washington, Seattle, WA Feb 2021

“Learning Visually Grounded Semantics with Denotation Graph”

- Google, Mountain View, CA Sept 2020

“Learning Generalizeable Neural Agent across Different Task Horizons”

- Google, Mountain View, CA Nov 2019

“Synthesized Policies for Transfer and Adaptation across Tasks and Environments”

- NeurIPS, Montreal, QC Dec 2018

“Being Negative but Constructively: Lessons from Creating Better VQA datasets”

- NAACL, New Orleans, LA Jun 2018

## TEACHING

*Course Assistant*, University of Southern California

CSCI-599: Deep Learning and its Applications  
CSCI-567: Machine Learning

Fall 2017 - 2018  
Spring 2020 - 2021

**AWARDS &  
HONORS**

**Awards and Prizes**

- NeurIPS Student Travel Award 2018, 2019
- Simon Fraser University President's Honour Roll 2014, 2015
- Zhejiang University First-class Academic Excellence Award 2011, 2012

**Scholarship & Fellowship**

- USC CS Department Fellowship 2017
- UCLA CS Department Fellowship 2016
- Simon Fraser University Entrance Scholarship 2013

**PRESS  
COVERAGE**

“BabyWalk AI breaks complex navigation into simple steps”  
by Kyle Wiggers, VentureBeat

May 12, 2020

**ADVISING &  
MENTORSHIP**

**Junior Ph.D Mentor**

- Bowen Zhang, Ph.D student at USC
- Liyu Chen, Ph.D student at USC

**Visiting Scholar Mentor**

- Wang Zhu, SFU, Next: Ph.D student at USC
- Jiacheng Chen, SFU, Next: Ph.D student at SFU
- Han-jia Ye, Nanjing University, Next: Assistant professor at Nanjing University
- Risto Vuorio, SK-Brain, Next: Ph.D student at Cambridge University
- Shiyi Lan, Fudan University, Next: Ph.D student at UMD

**Undergraduate Student Mentor**

- Yuna Lee, USC, Next: Software Engineer at Google