

Shi Feng

shifeng@cs.umd.edu • <http://umiacs.umd.edu/~shifeng/>
4108 Iribe Center • 8125 Paint Branch Drive • College Park, MD 20742

Education

New York University Visiting Student, Center of Data Science Hosted by Prof. He He.	New York City Fall 2019
University of Maryland PhD candidate, Computer Science Advised by Prof. Jordan Boyd-Graber, member of CLIP lab.	College Park, Maryland 2016 –
Shanghai Jiao Tong University B.S. in Computer Science Member of the ACM Honor Class.	Shanghai, China 2012 – 2016

Publications

Universal Adversarial Triggers for Attacking and Analyzing NLP EMNLP Eric Wallace, Shi Feng , Nikhil Kandpal, Matt Gardner, Sameer Singh	2019
Misleading Failures of Partial-input Baselines ACL (short) Shi Feng , Eric Wallace, Jordan Boyd-Graber	2019
Understanding Impacts of High-Order Loss Approximations and Features in Deep Learning Interpretation ICML Sahil Singla, Eric Wallace, Shi Feng , Soheil Feizi	2019
Trick Me If You Can: Human-in-the-loop Generation of Adversarial Examples for Question Answering TACL Eric Wallace, Pedro Rodriguez, Shi Feng , Jordan Boyd-Graber	2019
How Pre-trained Word Representations Capture Commonsense Physical Comparisons Commonsense Inference in NLP workshop Pranav Goel, Shi Feng , Jordan Boyd-Graber	2019
Quizbowl: The Case for Incremental Question Answering In submission Pedro Rodriguez, Shi Feng , Mohit Iyyer, He He, Jordan Boyd-Graber	2019
What can AI do for me: Evaluating Machine Learning Interpretations in Cooperative Play IUI Shi Feng , Jordan Boyd-Graber	2019
Pathologies of Neural Models Make Interpretation Difficult EMNLP (oral) Shi Feng , Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, Jordan Boyd-Graber	2018
Interpreting Neural Networks with Nearest Neighbors EMNLP Workshop BlackboxNLP Eric Wallace*, Shi Feng* , Jordan Boyd-Graber	2018
The UMD Neural Machine Translation Systems at WMT17 Bandit Learning Task WMT Amr Sharaf, Shi Feng , Khanh Nguyen, Kianté Brantley, Hal Daumé III	2017
Improving Attention Modeling with Implicit Distortion and Fertility for Machine Translation COLING Shi Feng , Shujie Liu, Nan Yang, Mu Li, Ming Zhou, Kenny Q. Zhu code (Theano)	2016

Other Projects

Play With QANTA: Human-computer Cooperative QA play.qanta.org

An online interface where humans cooperate with a human-level QA system (QANTA) and compete with each other. The interface enhances this cooperation by explaining the model's predictions to the human using various interpretation methods. Our **IUI'19** paper is based on experiments conducted using this interface.

QANTA: Human-level Quizbowl System github.com/pinafore/qb

At HSNCT'17 we beat *top* human players for the first time ([video](#)). I'm mainly responsible for the *buzzer* of QANTA, which controls when to buzz and when to wait. The buzzer was trained with reinforcement learning using game history collected from Protobowl. This RL buzzer was first introduced to the system for HSNCT'17 and turned out to be crucial to the victory against human.

Talks

NLP Highlights Podcast	Apr 25 2019
Invited talk at UPenn	Mar 25 2019
Invited talk at UCSD	Mar 19 2019
Invited talk at UCI	Mar 18 2019

Awards and Service

Best reviewer award, EMNLP	2018
Reviewer, EMNLP	2018, 2019
Reviewer, AACL	2020

Work Experience

Microsoft Research, Research Intern <ul style="list-style-type: none"> Health AI team Domain adaptation for machine translation. A boosting approach to safely select in-domain data to adapt a general translation system to the medical domain. 	2018.6 – 2018.8
Microsoft Research Asia, Research Intern <ul style="list-style-type: none"> Natural Language Computing Group Built the first neural machine translation system with Theano for NLC group. Improved the attention mechanism, results published at COLING'16. Experimented sequence-to-sequence for many other tasks, including pos tagging, parsing, and chinese couplet completion (link). 	2015.8 – 2016.2