

EDWARD HU

edward@edwardjhu.com
Redmond, WA, USA

EDUCATION

Johns Hopkins University, Baltimore, MD Class of 2019

B.Sc. in Computer Science, Cognitive Science

Advised by Benjamin Van Durme

- Departmental Honors in Comp. Sci., Cog. Sci.
- Cumulative GPA: 3.96/4.00

RESEARCH EXPERIENCE

Microsoft Azure AI, Sept 2020 – Present

Microsoft Corporation, Redmond, WA

Researcher

- Work with OpenAI to develop and commercialize GPT-3, shipping to products including Microsoft Power Apps and Power BI
- Theory of infinite-width neural networks and empirical implications

Microsoft Research AI, Sept 2019 – Aug 2020

Microsoft Corporation, Redmond, WA

AI Resident

- Accelerate hyperparameter tuning for extremely large models by orders of magnitude, and identify scaling factors that improve training stability using infinite-width neural network theories
- Improve the state-of-the-art attacks under the Wasserstein threat model, and collaborate in a unified theory for randomized smoothing, a type of certified adversarial defenses

Center for Language and Speech Processing, Jan 2018 – Aug 2019

Johns Hopkins University, Baltimore, MD

Research Assistant

- Research in natural language paraphrasing and controlled generation
- Invented a novel, accelerated dynamic beam allocation algorithm which helped to create ParaBank, a large English paraphrase corpus

SELECTED PUBLICATIONS

- **LoRA: Low-Rank Adaptation of Large Language Models**
Edward Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, W. Chen Pre-print
- **Feature Learning in Infinite Width Neural Networks**
Greg Yang, *Edward Hu* ICML 2021
- **Improved Image Wasserstein Attacks and Defenses** (*Best Paper Award*)
Edward Hu, Adith Swaminathan, Hadi Salman, Greg Yang ICLR 2020 Workshop
- **Randomized Smoothing of All Shapes and Sizes**
G. Yang, T. Duan, *Edward Hu*, H. Salman, I. Razenshteyn, J. Li ICML 2020
- **Large-scale, Diverse, Paraphrastic Bitexts via Sampling and Clustering** (*Oral*)
Edward Hu, A. Singh, N. Holzenberger, M. Post, B. Van Durme CoNLL 2019
- **Improved Lexically-Constrained Decoding for Translation and Monolingual Rewriting**
Edward Hu, H. Khayrallah, R. Culkin, P. Xia, T. Chen, M. Post, B. Van Durme NAACL 2019
- **ParaBank: Monolingual Bitext Generation and Sentential Paraphrasing via Lexically-constrained Neural Machine Translation** (*Oral*)
Edward Hu, Rachel Rudinger, Matt Post, Benjamin Van Durme AAAI 2019

Updated on 9/4/2021