

# YO JOONG “YJ” CHOE

Postdoctoral Scholar, Data Science Institute, University of Chicago

✉ [yjchoe@uchicago.edu](mailto:yjchoe@uchicago.edu)    📄 [yjchoe.github.io](https://github.com/yjchoe)

## RESEARCH INTERESTS

---

1. **Game-Theoretic Statistics & Forecast Evaluation:** anytime-validity; sequential inference; e-values & e-processes; confidence sequences; testing-by-betting; applications to forecast/model comparison
2. **Transparency & Controllability of Large Language Models (LLMs):** (mechanistic) interpretability; model steering; linear representations; causal frameworks for understanding LLM representations

## EDUCATION

---

<b>Ph.D. in Statistics and Machine Learning</b> Carnegie Mellon University <i>Thesis:</i> “Comparing Forecasters and Abstaining Classifiers” <i>Advisor:</i> Aaditya Ramdas <i>Committee:</i> Aarti Singh, Edward Kennedy, Johanna Ziegel, Alexander D’Amour	2023
<b>M.S. in Machine Learning</b> Carnegie Mellon University	2017
<b>B.S. in Mathematics and Computer Science</b> University of Chicago	2015

## EMPLOYMENT

---

<b>Postdoctoral Scholar</b> Data Science Institute, University of Chicago <i>Mentor:</i> Victor Veitch	2023 – Present
<b>Research Scientist</b> Kakao Corporation (‘17–‘19) & Kakao Brain (‘19–‘20) <i>Research Areas:</i> Deep Learning & Natural Language Processing	2017 – 2020

## AWARDS

---

<b>Best Paper Award</b> ICML Workshop on Mechanistic Interpretability	2024
<b>Student Poster Award</b> (w/ \$10K Research Grant) Citadel Securities Inaugural Ph.D. Summit	2022
<b>Franklin V. Taylor Memorial Award</b> (Best Paper and Presentation) IEEE Systems, Man, and Cybernetics Society	2018
<b>Phi Beta Kappa</b> (Top 5% of Class) University of Chicago	2014

## PREPRINTS

---

- 2024 **Y. J. Choe** & A. Ramdas. "Combining Evidence Across Filtrations Using Adjusters." Under Review.  
[arXiv] [slides] [code]
- 2024 K. Park, **Y. J. Choe**, Y. Jiang, & V. Veitch. "The Geometry of Categorical and Hierarchical Concepts in Large Language Models." Under Review.  
**🏆 Best Paper Award, ICML 2024 Workshop on Mechanistic Interpretability** (link)  
[arXiv] [poster] [code]

## PUBLICATIONS

---

*Asterisks (\*) denote equal contribution as co-first author.*

- 2024 K. Park, **Y. J. Choe**, & V. Veitch. "The Linear Representation Hypothesis and the Geometry of Large Language Models." In *International Conference on Machine Learning (ICML)*.  
[proc] [arXiv] [slides] [poster] [code]
- 2023 **Y. J. Choe**, A. Gangrade, & A. Ramdas. "Counterfactually Comparing Abstaining Classifiers." In *Neural Information Processing Systems (NeurIPS)*.  
[proc] [arXiv] [slides] [poster] [code]
- 2023 **Y. J. Choe** & A. Ramdas. "Comparing Sequential Forecasters." *Operations Research (OR)*.  
**🏆 Student Poster Award w/ \$10K Research Grant, Citadel Securities Inaugural PhD Summit**  
[journal] [arXiv] [slides] [poster] [code]
- 2020 J. Ham\*, **Y. J. Choe\***, K. Park\*, I. Choi, & H. Soh. "KorNLI and KorSTS: New Benchmark Datasets for Korean Natural Language Understanding." In *Findings of the Association for Computational Linguistics: EMNLP (EMNLP-F)*.  
[proc] [arXiv] [code]
- 2020 **Y. J. Choe\***, K. Park\*, & D. Kim\*. "word2word: A Collection of Bilingual Lexicons for 3,564 Language Pairs." In *Language Resources and Evaluation Conference (LREC)*.  
[proc] [arXiv] [code]
- 2020 K. Park, **Y. J. Choe**, J. Ham. "Jejueo Datasets for Machine Translation and Speech Synthesis." In *Language Resources and Evaluation Conference (LREC)*.  
[proc] [arXiv] [code]
- 2019 J. Lim, S. Ryu, K. Park, **Y. J. Choe**, J. Ham & W. Y. Kim. "Predicting Drug-Target Interaction Using a Novel Graph Neural Network with 3D Structure-Embedded Graph Representation." *Journal of Chemical Information and Modeling (JCIM)*.  
[journal] [arXiv]
- 2019 **Y. J. Choe\***, J. Ham\*, K. Park\* & Y. Yoon\*. "A Neural Grammatical Error Correction System Built On Better Pre-training and Sequential Transfer Learning." In *Workshop on Innovative Use of NLP for Building Educational Applications*. (Peer-reviewed & published in proceedings. Publisher: ACL.)  
**🏆 Runner-up (Top-2 Submission), Workshop Shared Task** (link)  
[proc] [arXiv] [code]
- 2019 S. Na, **Y. J. Choe**, D.-H. Lee, & G. Kim. "Discovery of Natural Language Concepts in Individual Units of CNNs." In *International Conference on Learning Representations (ICLR)*.  
[proc] [arXiv] [poster] [code]

- 2018 **Y. J. Choe**, S. Balakrishnan, A. Singh, J. M. Vettel, & T. Verstynen. "Local White Matter Architecture Defines Functional Brain Dynamics." In *IEEE International Conference on Systems, Man, and Cybernetics (IEEE SMC)*.  
**✂ Franklin V. Taylor Memorial Award (Best Paper and Presentation), IEEE SMC** (link)  
[proc] [arXiv] [slides]

## MISCELLANEOUS ARTICLES

---

- 2020 **Y. J. Choe**, J. Ham, & K. Park. "An Empirical Study of Invariant Risk Minimization." Presented at *ICML Workshop on Uncertainty and Robustness in Deep Learning*. [arXiv] [slides] [code]
- 2017 **Y. J. Choe\***, J. Shin\* & N. Spencer\*. "Probabilistic Interpretations of Recurrent Neural Networks." *Technical Report*. [report]
- 2016 **Y. J. Choe**. "A Statistical Analysis of Neural Networks." *Technical Report*. [report]
- 2016 M. Al-Shedivat, **Y. J. Choe**, N. Spencer, E. P. Xing. "Learning Diverse Overcomplete Dictionaries via Determinantal Priors." Presented at *ICML Workshop on Geometry in Machine Learning*. [abs]

## ACADEMIC SERVICES

---

### Area Chair

- NeurIPS Workshop, "Interpretable AI: Past, Present, and Future" 2024

### Session Chair

- Joint Statistical Meetings (Session: "New Advances in Nonparametric Hypothesis Testing - Part II") 2024

### Journal Referee

- JRSS-B (2024), ACM TALLIP (2023), JRSS-C (2022)

### Conference Reviewer / Program Committee

- ICLR (2025), AAAI (2025), NeurIPS (2024), ICML (2024)

### MLD Peer (Mentor for Junior PhD Students)

- Machine Learning Dept., Carnegie Mellon University (2 first-year students) 2022 – 2023

## TEACHING

---

### Postdoctoral Mentor

- Data Science Clinic, University of Chicago Fall 2023 – Spring 2024  
Our students won a **✂ 2023–2024 Award of Excellence in Data Science at UChicago!** (top-5 project; link)  
Student Project Title: "Operational Requirement Management Using Knowledge Graphs and LLMs"  
External Client: Argonne National Laboratory, AI for Operations Team

### Guest Lecturer

- "A Casual Introduction to AI, Machine Learning, and Large Language Models" [slides] Fall 2024  
*Ethical Challenges of AI, Ethics Lab, Georgetown University*

- “Model-Free Prediction for Stationary and Nonstationary Time Series in Python” Spring 2023  
*Financial Time Series Analysis, MS in Computational Finance Program, Carnegie Mellon University*
- “Basic Time Series Analysis and ARMA Modeling in Python” Spring 2022  
*Financial Time Series Analysis, MS in Computational Finance Program, Carnegie Mellon University*

### Graduate Teaching Assistant, Carnegie Mellon University

- UnitedHealth Group Bridges to Healthcare Technology Summer 2023  
*(Summer Research Program; link to students' final poster)*
- Financial Time Series Analysis Spring 2022; Spring 2023
- Simulation Methods for Option Pricing Spring 2023
- Data Science for Sports & Healthcare *(Summer Research Program)* Summer 2022
- Special Topics: Methods of Statistical Learning Spring 2021
- Statistical Machine Learning II Fall 2016
- Introduction to Probability Modeling Spring 2016
- Statistical Computing Fall 2015

### Undergraduate Teaching Assistant, University of Chicago

- Machine Learning and Large-Scale Data Analysis Spring 2015
- Foundations of Computational Data Analysis Winter 2015
- Calculus I-II-III Fall 2012 – Spring 2013

## SELECTED TALKS

---

### Invited

- “Mind the Filtration: E-processes vs. P-processes at Stopping Times” 2024  
*International Conference on Statistics and Data Science (ICSDDS)* Nice, France
- “Counterfactually Comparing Abstaining Classifiers” 2024  
*INFORMS Annual Meeting* Seattle, WA, USA
- “Comparing Sequential Forecasters” 2022  
*Citadel Securities Inaugural PhD Summit* Chicago, IL, USA
- “A Casual Introduction to Machine Learning & Deep Learning” 2019  
*Seoul Tech Society* Seoul, South Korea

### Contributed

- “Combining Evidence Across Filtrations Using Adjusters” 2024  
*Joint Statistical Meetings (JSM)* Portland, OR, USA
- “Counterfactually Comparing Abstaining Classifiers” [video] 2023  
*ICML Workshop on Counterfactuals in Minds and Machines* Honolulu, HI, USA
- “Comparing Sequential Forecasters” 2022  
*Workshop on Safe, Anytime-Valid Inference & Game-Theoretic Statistics* Eindhoven, The Netherlands
- “Local White Matter Architecture Defines Functional Brain Dynamics” 2018  
*IEEE International Conference on Systems, Man, and Cybernetics (SMC)* Miyazaki, Japan