ROBERT GIAQUINTO

284 S Marengo Ave, Pasadena, CA 91101
robert-giaquinto.github.io
(505)·379·8990 ⋄ giaquinto.ra@gmail.com

RESEARCH INTERESTS

Machine Learning · Generative Models · Language Modeling · Representation Learning · Pretraining

EDUCATION

University of Minnesota - Twin Cities, Minneapolis, MN *2016 — 2021*
Ph.D. in Computer Science
Thesis: Advancing Probabilistic Models for Approximate and Exact Inference
Advisor: Arindam Banerjee

University of Minnesota - Twin Cities, Minneapolis, MN *2014 — 2016*
M.S. in Computer Science
Capstone: *Graphical Models for Data with Spatiotemporal Dependencies*

St. Olaf College, Northfield, MN *2006 — 2010*
B.A. in Mathematics, Statistics
Center for Interdisciplinary Research Fellowship with Julie Legler

RESEARCH EXPERIENCE

Autodesk AI Lab February 2024 — Present
Principal AI Research Scientist *Los Angeles, CA*

· Research and development of large foundation models for the 3D world.
· Focusing on the architecture, engineering, and construction (AEC) domain, where we work with large amounts of data detailing the design of buildings.
· AEC data have a hierarchical structure, are multimodal, and pose the challenge of extremely long context lengths.

AWS AI Labs June 2021 — January 2024
Applied Scientist *New York, NY*

· Applied research with Amazon's CodeWhisperer and Bedrock large language modeling (LLM) services.
  - Lead scientist in pretraining models for database interaction, including the design of a multi-task learning strategy combining structured and unstructured data.
  - Trained and evaluated tokenizers, lead data engineering efforts, and expanded evaluation of LLMs to numerous tasks and over 100 benchmark datasets.
· Created a toolkit to evaluate the brittleness and sensitivity of Amazon Kendra's search results.
· Additional proof-of-concept projects incorporating cutting-edge techniques—like contrastive learning, in-context learning, and continual learning.

Department of Computer Science, University of Minnesota Sep 2016 — July 2021
Research Assistant *Minneapolis, MN*

· Research focuses on deep generative models, approximate inference, and probabilistic models with applications to text and image data.
· Developed a gradient boosted approach for training normalizing flows, increasing the flexibility of a powerful class of deep generative models.

· Discovered a new probabilistic model of authors and the topics they write about over time.
  - Scaled model to billion word corpora trained on the Minnesota Supercomputing Institute's systems.

Adobe June 2019 — Aug 2019
Data Science Intern *San Jose, CA*

· Developed sequence-to-sequence models for Adobe's Sensei AI email marketing products.
· Presented findings to internal audience of researchers on time-series and rare-event prediction methods.

HRL Laboratories May 2018 — Aug 2018
Research Intern *Malibu, CA*

· Machine learning research on an Intelligence Advanced Research Projects Activity (IARPA) research program for integrating human and machine forecasts.
· Derived a novel graphical model to augment human forecasting of geopolitical, macroeconomic, and health events.

Thomson Reuters Labs May 2016 — Aug 2016
R&D Intern *Eagan, MN*

· Discovered compact representation of a large corpus of legal texts to facilitate fast search and information retrieval.
· Modeling of legal texts combined topic, language, and embedding models.

Institute for Health Informatics, University of Minnesota Feb 2015 — May 2016
Research Assistant *Minneapolis, MN*

· Built an automated system that extracts and shares key sections of doctor's notes with hospital patients.
· Transformed unstructured rich text files from doctor's notes using natural language processing into a structured dataset.
· Key sections of text were extracted using a semi-supervised classification algorithm, which incorporates hundreds of thousands of unannotated doctor's notes in the learning process.

Capella Education Company Aug 2010 — Feb 2015
Research Analyst *Minneapolis, MN*

· Developed an automated system to predict academic success of students applying to Capella University.
  - Predictions created focus for academic coaching, signal alerts for faculty, recommend students for targeted orientation courses, and shift marketing strategies.
· Built statistical models relating individual factors to a likelihood of defaulting on student loans.
  - Tailored results of model to prioritize financial aid counseling teams.

PUBLICATIONS

Conference Papers

1. **R. Giaquinto**, D. Zhang, B. Kleiner, Y. Li, M. Tan, P. Bhatia, R. Nallapati, and X. Ma. Multi-task pretraining with structured knowledge for text-to-sql generation. In *Association for Computational Linguistics (ACL)*, 2023.

2. B. Athiwaratkun, S. K. Gouda, Z. Wang, X. Li, Y. Tian, M. Tan, W. U. Ahmad, S. Wang, Q. Sun, M. Shang, S. K. Gonugondla, H. Ding, V. Kumar, N. Fulton, A. Farahani, S. Jain, **R. Giaquinto**, H. Qian, M. K. Ramanathan, R. Nallapati, B. Ray, P. Bhatia, S. Sengupta, D. Roth, and B. Xiang. Multi-lingual evaluation of code generation models. In *International Conference on Learning Representations (ICLR)*, 2023.

3. **R. Giaquinto** and A. Banerjee. Gradient boosted normalizing flows. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

4. **R. Giaquinto** and A. Banerjee. DAPPER: Scaling the DAP topic model to billion word corpora. In *IEEE International Conference on Data Mining (ICDM)*, 2018.

5. **R. Giaquinto** and A. Banerjee. Topic modeling on health journals with regularized variational inference. In *AAAI Conference on Artificial Intelligence*, 2018.

6. R. Bjarnadottir, S. Maganti, M. J. Kreitzer, M. Mathiason, **R. Giaquinto**, and K. Monsen. Discovering the value of the omaha system for knowledge representation and data extraction in health intelligence. In *AAAI Joint Workshop on Health Intelligence (W3PHIAI)*, 2018.

Journal Articles

7. R. Cazzolla Gatti, P. B. Reich, J. G. Gamarra, T. Crowther, C. Hui, A. Morera, J.-F. Bastin, S. De-Miguel, G.-J. Nabuurs, J.-C. Svenning, and others. The number of tree species on earth. *Proceedings of the National Academy of Sciences*, 119(6):e2115329119, 2022.

8. C. E. Smith, Z. Levonian, **R. Giaquinto**, H. Ma, G. Lein-McDonough, Z. Li, S. O'Conner-Von, and S. Yarosh. "I Cannot Do All of this Alone": Exploring instrumental and prayer support in online health communities. *Transactions on Computer-Human Interaction (ToCHI)*, 2020.

9. H. Ma, C. E. Smith, L. He, S. Narayanan, **R. Giaquinto**, R. Evans, L. Hanson, and S. Yarosh. Write for life: Persisting in online health communities through expressive writing and social support. *Proceedings of the ACM on Human-Computer Interaction (CSCW)*, 1:73:1–73:24, 2017.

Preprints

10. R. Kwiatkowski, Z. Wang, **R. Giaquinto**, V. Kumar, X. Ma, A. Deoras, B. Xiang, and B. Athiwaratkun. FusionToken: Enhancing compression and efficiency in language model tokenization.

11. **R. Giaquinto** and A. Banerjee. Probabilistic super-resolution with normalizing flows.

12. **R. Giaquinto** and A. Banerjee. Compressive normalizing flows.

13. **R. Giaquinto** and T.-C. Lu. Promoting discussions among users in human-machine collaborative forecasting.

Thesis

14. **R. Giaquinto**. *Advancing Probabilistic Models for Approximate and Exact Inference*. PhD thesis, Jul 2021.

TEACHING AND INVITED TALKS

2020 NeurIPS - Gradient Boosted Normalizing Flows.

2019 & 2020 Teaching Assistant for Introduction to Artificial Intelligence.

2019 Teaching Assistant for Advanced Algorithms and Data Structures.

2019 Adobe, San Jose, CA - Intern research presentation.

2018 Teaching Assistant for Algorithms and Data Structures.

2018 HRL Laboratories, Malibu, CA - Promoting Discussions Among Users in Human-Machine Collaborative Forecasting.

2018 ICDM - DAPPER: Scaling the DAP Topic Model to Billion Word Corpora.

2018 Minnesota Supercomputing Institute Research Exhibition - Scaling Inference on Massive Corpora to Supercomputing Scales.

2018 AAAI - Topic Modeling on Health Journals with Regularized Variational Inference.

2018 CaringBridge Research Collaborative Ideation Workshop - Discovering Topics on CaringBridge Journals.

SOFTWARE

Gradient Boosted Normalizing Flows (Python package). 2019 — Present

Dynamic Author Persona topic models (Python package). 2017 — Present

See `http://github.com/robert-giaquinto/` for addition projects.

TECHNICAL STRENGTHS

Machine Learning	PyTorch, Tensorflow, Transformers, Deepspeed, AWS
Programing Languages, Proficient	Python, C, C++, CUDA, R, Regex, MATLAB, LaTeX, Bash
Programing Languages, Basic	Julia, Java, HTML, CSS, AWK
Databases	MySQL, PostgreSQL, Oracle, SQLite, MongoDB
Tools	Git, Docker, Terminal, Microsoft Suite
Operating Systems	Mac OSX, Windows, Linux

COMMUNITY SERVICE

Publicity Chair: International Conference on Artificial Intelligence and Statistics (AISTATS), 2021.

Reviewer: ICML (2017), KDD (2018), NeurIPS (2018), ICLR (2020), AISTATS (2021).

REFERENCES

Available on request.