

Umang Gupta

✉ umanggup@usc.edu — 🏠 umgupta.github.io — 📞 +1-2138222541 — [Google Scholar](#)

EDUCATION

- University of Southern California (USC), Los Angeles** **2017 - 2023**
Ph.D. in Computer Science (*Aug 2023*) GPA: 3.81/4
Thesis: Controlling Information in Neural Networks for Fairness and Privacy
Advisor: Prof. Greg Ver Steeg, Information Sciences Institute USC
[Jump to Publications](#)
- Indian Institute of Technology Delhi (IIT Delhi), India** **2010 - 2015**
B. Tech & M. Tech (Dual Degree), Electrical Engineering (**Silver Medalist**) GPA: 8.9/10
Thesis: Image Classification with Ontology and Deep Learning
Advisor: Prof. Santanu Chaudhury, Indian Institute of Technology, Delhi

RESEARCH INTERESTS & EXPERIENCE

- Fairness and Privacy in Machine Learning — Fairness in Language Models (LLMs), Fair Representation Learning, Differential Private Learning, Federated Learning
- Efficient & Robust Machine Learning — Efficient Tuning and Training of Transformer Models, Domain Adaptation & Generalization, Unsupervised or Self-Supervised Learning
- Applications of AI to Other Sciences & Healthcare (e.g., Neuroimaging)

WORK EXPERIENCE

Research Internships

- Machine Learning Center of Excellence, Morgan Stanley, New York** Jun'22 - Aug'22
Research Intern
- Contributed to several research projects focused on improving time-series forecasting.
 - Analyzed quantization error in transformers, leading to an efficient vector-quantized encoder-decoder transformer model whose complexity *scales linearly with sequence size*.
 - This collaboration resulted in two publications.
 - “Estimating transfer entropy under long ranged dependencies”, In Proc. of UAI 2022
 - “VQ-TR: Vector Quantized Attention for Time Series Forecasting”, In MILETS Workshop at KDD 2023

- Amazon Alexa-AI, Boston (Virtual)** May'21 - Aug'21
Applied Scientist Intern
- Studied biases in language model (e.g., GPT-2) outputs, especially the effect of model size on biases.
 - Devised a fair knowledge-distillation approach for reducing biases in large language models.
 - The results were published in the findings of ACL 2022 under the title “Mitigating Gender Bias in Distilled Language Models via Counterfactual Role Reversal.”

- Futurewei Technologies, Santa Clara** Jun'18 - Aug'18
Research Intern
- Investigated reinforcement learning algorithms for end-to-end training of seq2seq dialogue models (chatbots).

Software Development Experiences

- Visa Inc., Bangalore, India** Aug'15 - Jul'17
Senior Software Engineer
- Developed javascript web applications (ReactJS & BackboneJS) for card configuration management.
 - Created new dev tools and spearheaded the development of standard components to enhance productivity.
 - Received excellent rating at the year-end review for outstanding performance, participation, and developing deliverables on time.

Software Development Internships

- ↪ *MITACS Research Intern*, University of British Columbia, Vancouver, Canada May'14 - Jul'14
- ↪ *Software Development Intern*, Amagi Media Labs, Bangalore, India May'13 - Jul'13
- ↪ *Intern* Sohum Innovation Labs, Delhi, India May'12 - Dec'12

TECHNICAL SKILLS

Programming Languages	Python, Shell scripting, JavaScript, Java, C/C++
Libraries & Tools	PyTorch (among the top 10 answerers on StackOverflow), JAX, TensorFlow, scikit-learn, HuggingFace, deepspeed, L ^A T _E X, Git, React.js, Node.js,
Operating Systems	Linux, Unix, Windows

SELECTED PROJECTS

Deep Learning for Neuroimaging Sep'20 - Present

- Exploring applications of machine learning to neuroimaging in collaboration with Prof. Paul Thompson, Imaging and Genetics Center (USC).
- Introduced 2D-Slice-CNN models for learning from 3D MRIs that can exploit vision models pretrained on large-scale natural image datasets and outperform existing techniques.
- Resulted in two publications at ISBI 2021 and 2023.

Information-Theoretic Measures for Privacy & Fairness Jul'19 - Present

→ *Controlling Fairness via Mutual Information Minimization* [*Published at AAAI 2021*]

- Characterized an information-theoretic measure for statistical parity, a popular fairness measure.
- Demonstrated an effective method for controlling fairness through contrastive mutual information estimation.
- Showcased *better fairness-accuracy trade-offs* compared to recent competitive baselines.

→ *Minimizing Privacy Leakage during Training of the Neural Networks*

- Studied ways in which ML models may leak private training set information [*Published at MIDL 2021*].
- Investigating and improving techniques for differential private training of neural networks [*Ongoing*].

Continual Learning with Neural Networks Aug'17 - Jan'18

Awarded Best Theory Project in the Fall 2017 Deep Learning course (CSCI 599)

- Inspired by biological learning mechanisms in humans, we proposed dual generative (memory) models to solve the problem of catastrophic forgetting in Neural Networks
- Demonstrated *better retention of previously seen concepts* when training samples arrive in a non-iid fashion

SELECTED PUBLICATIONS [\[GOOGLE SCHOLAR\]](#)

Fair Machine Learning

Umang Gupta, Jwala Dhamala, . . . , Rahul Gupta, Kai-Wei Chang, Greg Ver Steeg, Aram Galstyan
Equitable Text Generation with Distilled Language Models via Counterfactual Role Reversal
Findings of the Association for Computational Linguistics (ACL 2022)

Umang Gupta, Aaron Ferber, Bistra Dilkina, Greg Ver Steeg
Controllable Guarantees for Fair Outcomes via Contrastive Information Estimation
Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI 2021)

Ninareh Mehrabi, **Umang Gupta**, Fred Morstatter, Greg Ver Steeg, Aram Galstyan
Attributing Fair Decisions with Attention Interventions
TrustNLP Workshop @ NAACL 2022

Efficient, Private, and Distributed Learning

Umang Gupta, Aram Galstyan, Greg Ver Steeg
Jointly Reparameterized Multi-Layer Adaptation for Efficient and Private Tuning
Findings of the Association for Computational Linguistics (ACL 2023)

Kashif Rasul, **Umang Gupta**, Hena Ghonia, Anderson Schneider, Yuriy Nevmyvaka
VQ-TR: Vector Quantized Attention for Time Series Forecasting
Under review @ ICLR 2024 (preprint available upon request); MILETS KDD 2023 workshop

Umang Gupta, Dimitris Stripelis, Pradeep Lam, Paul Thompson, José Luis Ambite, Greg Ver Steeg
Membership Inference Attacks on Deep Regression Models for Neuroimaging
Proceedings of the 4th Conference on Medical Imaging with Deep Learning (MIDL 2021)

Dimitris Stripelis, **Umang Gupta**, Greg Ver Steeg, José Luis Ambite
Federated Progressive Sparsification (Purge, Merge, Tune)+
Federated Learning Workshop @ NeurIPS 2022

Others

Umang Gupta, Tamoghna Chattopadhyay, Nikhil Dhinagar, Paul Thompson, Greg Ver Steeg
Transferring Models Trained on Natural Images to 3D MRI via Position Encoded Slice Models
IEEE International Symposium on Biomedical Imaging (ISBI), 2023

Sahil Garg, **Umang Gupta**, Yu Chen, Syamantak Gupta, Yeshaya Adler, Anderson Schneider, Yuriy Nevmyvaka
Estimating Transfer Entropy under Long Ranged Dependencies
Uncertainty in Artificial Intelligence (UAI), 2022

Umang Gupta, Pradeep Lam, Greg Ver Steeg, Paul Thompson
Improved Brain Age Estimation with Slice-based Set Networks
IEEE International Symposium on Biomedical Imaging (ISBI), 2021

Nitin Kamra, **Umang Gupta**, Fei Fang, Yan Liu, Milind Tambe
Policy Learning for Continuous Space Security Games using Neural Networks
Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI 2018)

Nitin Kamra, **Umang Gupta**, Yan Liu
Deep Generative Dual Memory Network for Continual Learning
<https://arxiv.org/abs/1710.10368>

TEACHING EXPERIENCE

Teaching Assistant, USC

- Special Topics in Machine Learning, CSCI699 (Spring 2022)
- Machine Learning, CSCI567 (Summer 2020, Summer 2019, Spring 2018)
- Applied Natural Language Processing, CSCI544 (Spring 2019)
- Software Engineering, CSCI310 (Fall 2017)

Teaching Assistant, IIT Delhi

- Pattern Recognition, EEL709 (Spring 2015)
- Digital Signal Processing, EEL319 (Fall 2014)
- Circuit Theory, EEL202 (Fall 2013)

AWARDS AND HONORS

Program	Recipient of IITD Semester Merit Award for 4 of 8 semesters, 2010 - 2014.
Rank 1	Highest GPA in Dual Degree Program, Department of Electrical Engineering.
Research	Best Theory Project Award for arxiv:1710.10368 , CSCI-599: Deep Learning, USC (2017).
Best Essay	Among the top 20 National Winners, International Year of Forest, 2011 certification program.
Hackathons	First Prize, GS Quantify 2014, an annual computing competition organized by Goldman Sachs. First Prize, Bing Hackathon 2016, a machine learning contest organized by Microsoft Bing.
NIUS 2011	Among 30 students invited to research at HBCSE, Mumbai, under National Initiative for Undergraduate Science (NIUS) program. The research results were published in Monthly Notices of Royal Astronomical Society.
Competitive Exams	Recipient of KVPY scholarship 2010, Dept. of Science and Technology, Govt. of India. Top 1% in Physics (NSEP), Astronomy (NSEA), and Chemistry (NSEC) Olympiads in 2009. Secured All India Rank 410 in IIT-Joint Entrance Exam among 0.5 million candidates.

PROFESSIONAL SERVICES

Reviewing	Medical Imaging and Analysis (MedIA) Journal Neural Information Processing Systems (NeurIPS) – 2023 International Conference on Learning Representations (ICLR) – 2023, 2024 Association for Computational Linguistics (ACL) – 2023 Uncertainty in Artificial Intelligence (UAI) – 2023 European Conference on Artificial Intelligence (ECAI) – 2023 Asian Conference on Machine Learning (ACML) – 2022 Artificial Intelligence and Statistics (AISTATS) – 2021 Medical Imaging with Deep Learning (MIDL) – 2021, 2022, 2023 Workshops: Federated Learning Systems @ MLSys 2023; TrustNLP @ ACL 2023; {ENLSP, MedNeurIPS} @ NeurIPS 2022; TrustNLP @ NAACL 2022; FL4NLP @ ACL 2022;
-----------	---