Kartik Ahuja

Contact Information

Email: ahujak@ucla.edu
Phone: +1 626-362-8188
Scholar, Website, Github

Research Interests

Machine Learning, Optimization and Causality

Work Experience

Research Scientist, FAIR, Meta **2022 - present**

Postdoctoral Fellow at Mila - Quebec AI Institute with Yoshua Bengio, Ioannis Mitliagkas, and Irina Rish **2020 -2022**

Research Scientist (AI Resident) at IBM Research, TJ Watson Research Center in Trustworthy AI team led by Kush Varshney **2019 -2020**

Education

University of California, Los Angeles **2013-2019**
PhD in Electrical and Computer Engineering

- GPA: **4.0/4.0**
- Thesis: "Optimization Methods for Resource Allocation and Machine Learning Applications."
- Advisor: Gregory J. Pottie

Indian Institute of Technology, Kanpur **2008-2013**
B. Tech - M. Tech Dual Degree in Electrical Engineering

- GPA: **8.8/10** (B. Tech) and **9.6/10** (M. Tech)
- Thesis: "Optimizing Signal Constellations."
- Advisor: Ajit K. Chaturvedi

Visiting Positions & Internship

University of California, Berkeley **Jan - Feb 2022**
Visiting Scholar

University of Oxford **Jan -Dec 2017**
Visiting PhD student

University of British Columbia **May -July 2012**
Internship

Publications

Machine Learning

K. Donhauser, C. Arnal, M. Pezeshki, V. Cabannes, D. Lopez-Paz, **K. Ahuja**[†]
Unveiling Simplicities of Attention: Adaptive Long-Context Head Attention
Submitted.
[†] Last author

D. Mahajan, M. Pezeshki, C. Arnal, **K. Ahuja**[†], P. Vincent[†]
Compositional Risk Minimization
Submitted.
[†] Last author

K. Ahuja, A. Mansouri
On Provable Length and Compositional Generalization
Submitted.

A Vysogorets, **K. Ahuja**, J. Kempe
DroP: Distributionally Robust Data Pruning
International Conference on Learning Representations (ICLR), 2025, **(Spotlight).**

K. Ahuja[*], A. Mansouri[*], Y. Wang
Multi-Domain Causal Representation Learning with Weak Distributional Invariances
International Conference on Artificial Intelligence and Statistics (AISTATS), 2024.
[*]Equal contribution

S. Gupta, S. Jegelka, D. Lopez-Paz, **K. Ahuja**[†]
Context is Environment
International Conference on Learning Representations (ICLR), 2024.
[†] Last author

V. Barin-Pacela, **K. Ahuja**, S. Lacoste-Julien, P. Vincent
On the Identifiability of Quantized Factors
Causal Learning and Reasoning Conference (CLeaR), 2024.

T. Nguyen, A. Mansouri, K. Madan, K.D. Nguyen, **K. Ahuja**, D. Liu, and Y. Bengio.
Reusable Slotwise Mechanisms
Neural Information Processing Systems (NeurIPS), 2023.

A. Dhurandhar, K. Ramamurthy, **K. Ahuja**, V. Arya.
Locally Invariant Explanations: Towards Stable and Unidirectional Explanations through Local Invariant Learning
Neural Information Processing Systems (NeurIPS), 2023.

K. Ahuja, D. Lopez-Paz
A Closer Look at In-Context Learning under Distribution Shifts
preprint, 2023, (arXiv). *ICML workshop on Foundation Models, 2023.*

K. Ahuja, D. Mahajan, Y. Wang, Y. Bengio
Interventional Causal Representation Learning
International Conference on Machine Learning (ICML), 2023 **(Oral presentation).** *Neuro-Causal and Symbolic AI workshop, NeurIPS 2022,* **(Oral presentation)**.

K. Chaudhuri[*], **K. Ahuja**[*], M. Arjovsky, D. Lopez-Paz
Why Does Throwing Away Data Improve Worst-Group Error?
International Conference on Machine Learning (ICML), 2023 **(Oral presentation).**
[*]Equal contribution.

A. Ramé, **K. Ahuja**, J. Zhang, M. Cord, L. Bottou, and D. Lopez-Paz.
Model ratatouille: Recycling diverse models for out-of-distribution generalization.
International Conference on Machine Learning (ICML), 2023. ICLR workshop on domain generalization, 2023 **(Spotlight).**

J.C. Audet, **K. Ahuja**, M.J. Bayazi, P. Mousavi, G. Dumas, I. Rish.
WOODS: Benchmarks for Out-of-Distribution Generalization in Time Series
Transactions on Machine Learning Research, 2023. **(Featured Certification)**. *International Conference on Learning Representations (ICLR), 2024 as TMLR featured paper.*

H. Naganuma, **K. Ahuja**, S. Takagi, T. Motokawa, I. Mitliagkas, R. Yokota, K. Ishikawa, I. Sato
Empirical Study on Optimizer Selection for Out-of-Distribution Generalization
Transactions on Machine Learning Research, 2023.

K. Ahuja, J. Hartford, Y. Bengio.
Weakly Supervised Representation Learning with Sparse Perturbations
Neural Information Processing Systems (NeurIPS), 2022.

S. Gupta, **K. Ahuja**, M. Havaei, N. Chatterjee, Y. Bengio.
FL Games: A federated learning framework for distribution shifts
preprint, 2022. (arXiv) *Workshop on Federated Learning, NeurIPS 2022.* **(Oral presentation).**

K. Ahuja, J. Hartford, Y. Bengio.
Properties from mechanisms: An equivariance perspective on identifiable representation learning
International Conference on Learning Representations (ICLR), 2022. **(Spotlight)**

K. Ahuja [*], D. Mahajan[*], V. Syrgkanis, I. Mitliagkas.
Towards Efficient Representation Identification in Supervised Learning
Causal Learning and Reasoning Conference (CLeaR), 2022.
[*]Equal contribution.

A. Shah, K. Shanmugam, **K. Ahuja**.
Finding Valid Adjustments under Non-Ignorability with Minimal DAG Knowledge
International Conference on Artificial Intelligence and Statistics (AISTATS), 2022.

K. Ahuja, E. Caballero, D. Zhang, J.C. Audet, Y. Bengio, I. Mitliagkas, I. Rish.
Invariance Principle Meets Information Bottleneck for Out-of-Distribution Generalization
Neural Information Processing Systems (NeurIPS), 2021. **(Spotlight)**

P. Bashivan, R. Bayat, A. Ibrahim, **K. Ahuja**, M. Faramarzi, T. Laleh, B. Richards, I. Rish.
Adversarial Feature Desensitization
Neural Information Processing Systems (NeurIPS), 2021.

K. Ahuja, P. Sattigeri, K. Shanmugam, D. Wei, K.N. Ramamurthy, M. Kocaglu.
Conditional Independent Data Generation
Uncertainty in Artificial Intelligence (UAI), 2021.

D. Zhang, **K. Ahuja**, Y. Xu, Y. Wang, A. Courville.
Can Subnetwork Structure be the Key to Out-of-Distribution Generalization?
International Conference on Machine Learning (ICML), 2021. **(Oral presentation)**

Abhin Shah, **K. Ahuja**, K. Shanmugam, D. Wei, K. Varshney, A. Dhurandhar.
Treatment Effect Estimation Using Invariant Risk Minimization
IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2021.

K. Ahuja, K. Shanmugam, A. Dhurandhar.
Linear Regression Games: Convergence Guarantees to Approximate Out-of-Distribution Solutions
International Conference on Artificial Intelligence and Statistics Conference (AISTATS), 2021.

K. Ahuja, J. Wang, A. Dhurandhar, K. Shanmugam, K. Varshney.
Empirical or Invariant Risk Minimization? A Sample Complexity Perspective
International Conference on Learning Representations (ICLR), 2021.

K. Ahuja, K. Shanmugam, K. R. Varshney, A. Dhurandhar.
Invariant Risk Minimization Games
International Conference on Machine Learning (ICML), 2020.

K. Ahuja, K. Shanmugam, K. R. Varshney, A. Dhurandhar.
On the Equivalence of Bi-Level Optimization and Game-Theoretic Formulations of Invariant Risk Minimization
Inductive Biases, Invariances and Generalization in RL Workshop, International Conference on Machine Learning (ICML), 2020.

K. Ahuja, A. Dhurandhar, K. Shanmugam, K. R. Varshney.
Learning to Initialize Gradient Descent using Gradient Descent
preprint, 2020. (arXiv)

K. Ahuja
Estimating Kullback-Leibler Divergence Using Kernel Machines
53rd Annual Asilomar Conference on Signals, Systems, and Computers (ACSSC), 2019.

K. Ahuja, W. Zame, M. van der Schaar.
Optimal Piecewise Approximations for Model Interpretation
53rd Annual Asilomar Conference on Signals, Systems, and Computers (ACSSC), 2019.
(Second best student paper award)

K. Ahuja, M van der Schaar.
Joint Concordance Index
53rd Annual Asilomar Conference on Signals, Systems, and Computers (ACSSC), 2019, and *Neural Information Processing Systems workshop on Machine Learning for Health (NeurIPS, ML4H), 2017.*

K. Ahuja, W. Zame, M. van der Schaar.
Dpscreen: Dynamic Personalized Screening
Neural Information Processing Systems (NeurIPS), 2017.

Optimization & Networks

K. Ahuja, M. van der Schaar.
Dynamic Matching and Allocation of Tasks
ACM Transactions on Economics and Computation, vol. 7 no. 1, pp 1-27, 2019.

K. Ahuja, Y. Xiao, M. van der Schaar.
Efficient Interference Management Policies for Femtocell Networks
IEEE Transactions on Wireless Communications, vol. 14, no. 9, pp 4879-4893, 2015.
Featured in IEEE-spotlight , UCLA-news.

K. Ahuja, Y. Xiao, M. van der Schaar.
Distributed Interference Management Policies for Heterogeneous Networks
IEEE Journal on Selected Areas in Communications, vol. 33, no. 6, pp. 1112-1126, 2015.
Featured in IEEE MMTC letter, December, 2016.

Y. Xiao, **K. Ahuja**, M. van der Schaar.
Spectrum Sharing For Delay-Sensitive Applications with Continuing QoS Guarantees
IEEE Global Communications Conference (GLOBECOM), 2014.
Nominated for the best paper award (top 50 papers among 2100 submissions.)

K. Ahuja, M. Hasan, J. Hossain.
To Participate or Not in Spectrum Auctions with Entry fee: Bayesian Game Theoretic Approach
IEEE Wireless Communications and Networking Conference (WCNC), 2014.

K. Ahuja, M. van der Schaar, W. Zame.
Working alone and working with others: implications for the malthusian era
Economic Theory, pp.1-35, 2019.

A. Alaa, **K. Ahuja**, and Mihaela van der Schaar.
A Micro-Foundation of Social Capital in Evolving Social Networks
IEEE Transactions on Network Science and Engineering vol. 5, no. 1, pp. 14-31, 2017.

A. Alaa, **K. Ahuja**, M. van der Schaar.
Self-organizing Networks of Information Gathering Cognitive Agents
IEEE Transactions on Cognitive Communications and Networking, vol. 1. no. 1, pp 100-112, 2015.

K. Ahuja, S. Zhang, M. van der Schaar
Towards a Theory of Societal Co-Evolution: Individualism versus Collectivism
IEEE Global Conference on Signal and Information Processing (GlobalSIP), 2014.

K. Ahuja, S. Zhang, M. van der Schaar.
The Population Dynamics of Websites
Netecon workshop at ACM Conference on Economics and Computation (EC), 2015.

Patents

A. Dhurandhar, K.N. Ramamurthy, **K. Ahuja**, V. Arya.
Generating locally invariant explanations for machine learning
18/048,341. filled October 19, 2022.

K. Ahuja, A. Dhurandhar, K. Shanmugam, and K R. Varshney.
Learning Robust Predictors Using Game Theory
17/115,489, filed December 8, 2020.

K. Ahuja, A. Dhurandhar, K. Shanmugam, and K R. Varshney.
Initializing Optimization Solvers
17/101,019, filed November 23, 2020.

K. Ahuja, P. Sattigeri, K. Shanmugam, D. Wei, M. Kocaglu, K.N. Ramamurthy
Conditional Independent Data Generation for Training Machine Learning Systems
filed June 26, 2021.

Academic Achievements

- IVADO postdoctoral fellowship (2021-2023).
- Co-authored a successful grant application to Microsoft research on Causal Machine Learning resulting in a funding of 54,000 CAD.
- Top 8 percent reviewer NeurIPS, 2021, top 10 percent reviewer NeurIPS, 2020 and ICML, 2021.
- Second best student paper award at the 53rd Annual Asilomar Conference on Signals, Systems, and Computers (ACSSC), 2019.
- UCLA Dissertation Year Fellowship (2018-19).
- Guru Krupa Foundation Fellowship by the ECE Department at UCLA (2013-14).
- Departmental Fellowship by the ECE Department at UCLA (Fall, 2013).
- All India Rank of 584 in Joint Entrance Examination 2008 (99.8 percentile) among more than 3,50,000 students.
- All India Rank of 1131 in All India Engineering Entrance Examination 2008 (99.8 percentile) among more than 7,50,000 students.

Invited Talks

- "Invariant Risk Minimization Games", Facebook Artificial Intelligence Research (FAIR), NY, 2020.
- "Invariant Risk Minimization Games", Computer Science Department, Yale University, 2020.
- "Out-of-Distribution Generalization: Invariance Principle and Beyond", Google Brain, 2022.
- "Out-of-Distribution Generalization: Invariance Principle and Beyond", University of Georgia, 2023.
- "Out-of-Distribution Generalization: Invariance Principle and Beyond", University of Texas, Austin, 2023.
- "Interventional Causal Representation Learning", Causal Representation Learning Workshop, Tübingen, 2023.
- "Interventional and Multi-Domain Causal Representation Learning", Bellairs Workshop on Causality, Barbados, 2024.
- "On Provable Length and Compositional Generalization", Seminar on Formal Language and Neural Networks.

Teaching Experience

- Teaching Assitant at UCLA: Digital Signal Processing, Network Economics and Game Theory, and Multimedia Communications.
- Teaching Assitant at IIT Kanpur: Representation and Analysis of Random Signals.

Students Mentored

- Divyat Mahajan (PhD student at Mila)
- Konstantin Donhauser (PhD student at ETH Zurich)
- Sharut Gupta (PhD student at MIT)
- Abhin Shah (PhD student at MIT)
- Amin Mansouri (PhD student at EPFL)
- Jun Wang (PhD student at Rensselaer Polytechnique)
- Dinghuai Zhang (PhD student at Mila)
- Jean-Christophe Gagnon-Audet (MSc. student at Mila)
- Hiroki Naganuma (PhD student at Mila)
- Marcel Nwaunka (BS student at University of Arkansas)

• Naveene Raya (BS student at California State University)

Professional Service

Reviewer

• Causal Learning and Reasoning (CLeaR) 2022, 2023

• Neural Information Processing Systems (NeurIPS) 2019, 2020, 2021, 2022, 2023

• International Conference on Machine Learning (ICML) 2019, ICML 2021 (**Expert Reviewer**) , ICML 2022, ICML 2023, ICML 2024

• Artificial Intelligence and Statistics Conference (AISTATS), 2021, 2022

• International Conference and Learning Representations (ICLR), 2021, 2022, 2023

• Journal of Machine Learning Research (JMLR)

• Transactions on Information Theory

• Association for Advancement of Artificial Intelligence (AAAI) conference 2020

• IEEE Global Communciations Conference (GLOBECOM)

• IEEE Journal of Selected Areas in Communications (JSAC)

• ACM/IEEE Transactions on Networking (TNET)

• National Conference on Communications (NCC), India

Area Chair

• Robustness of Few-shot and Zero-shot Learning in Foundation Models, NeurIPS 2023.

Programming Skills

Languages: Python, R, Java, Matlab
Frameworks for machine learning: Tensorflow, Keras.