

# Greg Ver Steeg

Phone: (626) 840-5901  
Email: gversteeg@gmail.com  
Homepage: <http://gregversteeg.com>

## Biographical sketch

Dr. Ver Steeg's research explores *practical methods for inferring meaningful structure in complex systems*. This work draws on a diverse set of connections between machine learning, information theory, causal inference, and physics. His recent work has applied these tools to solving outstanding problems in diverse domains including biology, natural language, and human behavior.

## Education

Ph.D., California Institute of Technology 2009  
Field: Physics, Area: Quantum Information and Computation  
Advisor: John Preskill, Feynman Professor of Theoretical Physics  
B.S., Drake University 2003  
Majors: Physics, Mathematics, Philosophy

## Experience

University of California, Riverside  
Computer Science & Engineering,  
Associate Professor (IV) 2023 - Present

University of Southern California  
Computer Science Department,  
Adjunct Research Associate Professor 2023 - Present  
Research Associate Professor 2019 - 2022  
Research Assistant Professor 2013 - 2019  
Information Sciences Institute,  
Senior Research Lead 2018 - 2022  
Research Lead 2015 - 2018  
Computer scientist 2011 - 2015  
Postdoctoral researcher 2009 - 2011  
USC Norris Comprehensive Cancer Center,  
Full member 2022 - Present

Amazon  
Visiting Academic 2021 - Present

ZEFR  
Scientific advisor 2019 - 2021

California Institute of Technology 2003-2009  
Graduate research

Summer internships	
Google Inc., Machine translation	2009
Brookhaven National Laboratory, Nuclear physics	2002
Los Alamos National Laboratory, Nuclear physics	2001
Drake University	2001-2003
Computational atomic physics research	

## Recent Sponsored Research

ISI Exploratory Research Award. "FairPRS: Fairly Predicting Genetic Risk Scores for Personalized Medicine"	2021-2022
Keck Foundation COVID Fund. "Model-based immuno-phenotype mapping to investigate signatures of immune response and disease severity in SARS-CoV-2 infection to inform precision medicine based therapeutic strategies"	2021
DARPA Reversible Quantum Machine Learning and Simulation Program	2021
DARPA NEED Program	2021
DARPA Cooperative Secure Learning	2020-2022
DARPA Learning with Less Labels	2019-2023
Amazon Research Award	2019
Keston Foundation Research Award	2018
James H. Zumberge Research and Innovation Fund Award	2017
DARPA Next Generation Social Science	2016-2020
DARPA Social Media in Strategic Communication grant "Situational Awareness for Social Media: Theories, Models and Algorithms"	2012-2016
AFOSR Young Investigator Award, "Bell Inequalities for Complex Networks"	2012-2015
Army Research Office, "Optimization via Open System Quantum Annealing"	2012-2015

## Awards

Best Paper Award, NeurIPS Deep Learning through Information Geometry Workshop	2020
Amazon Research Award	2019
ISI Institute Achievement Award	2018
IJCAI Early Career Spotlight	2017
AFOSR Young Investigator Award	2012-2015
Best paper runner-up at UAI	2011
National Science Foundation Graduate Research Fellowship	2005-2008
California Institute of Technology Richard Feynman Fellowship	2003
Barry Goldwater Scholarship	2002
Drake University Physics Scholarship	1999-2003

## Teaching and mentoring

### *Advising or co-advising*

#### PhD students

Hrayr Harutyunyan (Google Research)	2023
Robert Brekelmans (postdoc fellowship at Vector Institute)	2022
Sami Abu-al-haija (Google Research)	2022
Daniel Moyer (postdoc at MIT CSAIL, now Prof at Vanderbilt)	2019
Shuyang Gao, PhD (previously Amazon, now at an AI startup)	2018
David Kale, PhD (previously Netflix, now at an AI startup)	2018
Umang Gupta (Google Bard)	2023
Myrl Marmarelis	Expected 2024
Neal Lawton	Expected 2024
Xianghao Kong	
Han Li	

### *Teaching*

CS 224, Fundamentals of Machine Learning, UCR	Winter 2024
CS 229, Advanced Machine Learning, UCR	Spring 2023
CSCI-699 (Advanced Special Topics): Dynamics of Representation Learning, USC	Spring 2022

Designed a new course looking at different aspects of learning (optimization dynamics, models defined via dynamics, distribution shift) using tools from classical physics and non-equilibrium thermodynamics. Draft syllabus available at <https://sites.google.com/view/699dynamicsofrep/home>. 11/11 students evaluations rated my teaching as “excellent” (the highest score), with several students commenting that it was the best or most helpful graduate course they had taken.

CSCI-699 (Advanced Special Topics): Representation Learning: Theory and Practice, USC	Fall 2019
---	-----------

Designed and taught with Aram Galstyan and Sami Abu-El-Haija a course focusing on principles and practical aspects behind successful deep learning. Over 50 PhD students took the class. 44/45 student evaluations rated my teaching as “excellent” (the highest score). Syllabus available at <http://sami.haija.org/cs699/>.

CSCI-590: Directed research, USC	2014- Present
CSCI-599: Physics and Computation, USC	Spring 2012

Designed and taught with Aram Galstyan. Studied connections between theoretical computer science and machine learning, statistical physics of disordered systems, and information theory. Main text was Mezard and Montanari, “Information, Physics, and Computation”. First course to ever program and run a quantum annealing computer for a course project (USC was the first university to purchase a D-Wave computer). Two class projects led to publications.

T.A. Physics 2, California Institute of Technology	2008
--	------

## Service

Open source code contributions:

**Github stats:** 17 public repositories with over 1500 “stars” and over 250 “forks”

(Not including advising on repositories developed by students)

Some code has been ported into autoML systems at <http://datadrivendiscovery.org>

Information-Theoretic Diffusion	2023
<a href="https://github.com/gregversteeg/InfoDiffusionSimple">https://github.com/gregversteeg/InfoDiffusionSimple</a>	
Hamiltonian Dynamics with Non-Newtonian Momentum for Rapid Sampling	2021
<a href="https://github.com/gregversteeg/esh_dynamics">https://github.com/gregversteeg/esh_dynamics</a>	
Linear factor models and covariance estimation with Linear CorEx	2017
<a href="https://github.com/gregversteeg/linearcorex">https://github.com/gregversteeg/linearcorex</a>	
Nonlinear CorEx models optimized for under-sampled, high-d biomedical data.	2017
<a href="https://github.com/gregversteeg/bio_corex">https://github.com/gregversteeg/bio_corex</a>	
Constructing topic models with CorEx	2016
<a href="https://github.com/gregversteeg/corex_topic">https://github.com/gregversteeg/corex_topic</a>	
The discrete information sieve	2016
<a href="https://github.com/gregversteeg/discrete_sieve">https://github.com/gregversteeg/discrete_sieve</a>	
Gaussianizing data	2015
<a href="https://github.com/gregversteeg/gaussianize">https://github.com/gregversteeg/gaussianize</a>	
Information-theoretic deep learning:	2014
<a href="https://github.com/gregversteeg/CorEx">https://github.com/gregversteeg/CorEx</a>	
Non-parametric information-theoretic estimation code:	2013
<a href="https://github.com/gregversteeg/NPEET">https://github.com/gregversteeg/NPEET</a>	
ISI AI Seminar Coordinator	2014
Tutorials: ICWSM (2013), IPAM (2016), IJCAI (2016)	
Co-organizer of NeurIPS Information Geometry Workshop	2020
Research faculty hiring committee	2021
Program Committees and reviewing: NeurIPS, ICML, ICLR, UAI, AISTATS, AAAI, Journal of Causal Inference, ICWSM, CSS, Physical Review E, <i>Entropy</i> , AFOSR proposals, Mitacs, <i>Neural Computation</i> , IEEE Transactions on Knowledge and Data Engineering, WebSci.	

## Recent pre-prints

arXiv:2312.14440, Asymmetric Bias in Text-to-Image Generation with Adversarial Attacks

Online Continual Learning for Progressive Distribution Shift (OCL-PDS): A Practitioner’s Perspective. <https://openreview.net/forum?id=gfzs0rkMT9>. Accepted to ICLR 2023 Domain Generalization Workshop (spotlight presentation)

Secure Federated Learning for Neuroimaging. <https://arxiv.org/pdf/2205.05249.pdf>

Federated Progressive Sparsification (Purge, Merge, Tune). <https://arxiv.org/abs/2204.12430>

Functional Connectivity in Visual Areas from Total Correlation.

Unsupervised domain harmonization of MRI Brain Scan using 3D Cycle GANs. <https://www.biorxiv.org/content/10.1101/2022.11.15.516349v1>.

## Peer-Reviewed Publications

- [1] Dheeraj Komandur, Umang Gupta, Tamoghna Chattopadhyay, Nikhil J Dhinagar, Sophia I Thomopoulos, Jiu-Chiuan Chen, Dan Beavers, Greg Ver Steeg, and Paul M Thompson. Unsupervised harmonization of brain mri using 3d cyclegans and its effect on brain age prediction. In *2023 19th International Symposium on Medical Information Processing and Analysis (SIPAIM)*, pages 1–5. IEEE, 2023.
- [2] Neal Lawton, Anoop Kumar, Govind Thattai, Aram Galstyan, and Greg Ver Steeg. Neural architecture search for parameter-efficient fine-tuning of large pre-trained language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, 2023.
- [3] Umang Gupta, Aram Galstyan, and Greg Ver Steeg. Jointly reparametrized multi-layer adaptation for efficient and private tuning. In *Findings of the Association for Computational Linguistics: ACL 2023*, 2023.
- [4] Tamoghna Chattopadhyay, Amit Singh, Neha Ann Joshy, Sophia I Thomopoulos, Talia M Nir, Hong Zheng, Elnaz Nourollahimoghadam, Umang Gupta, Greg Ver Steeg, Neda Jahanshad, et al. Predicting dementia severity by merging anatomical and diffusion mri with deep 3d convolutional neural networks. In *18th International Symposium on Medical Information Processing and Analysis*, volume 12567, pages 90–99. SPIE, 2023.
- [5] Nikhil J Dhinagar, Sophia I Thomopoulos, Priya Rajagopalan, Dimitris Stripelis, Jose Luis Ambite, Greg Ver Steeg, and Paul M Thompson. Evaluation of transfer learning methods for detecting alzheimer’s disease with brain mri. In *18th International Symposium on Medical Information Processing and Analysis*, volume 12567, pages 504–513. SPIE, 2023.
- [6] Umang Gupta, Tamoghna Chattopadhyay, Nikhil Dhinagar, Paul M Thompson, and Greg Ver Steeg. Transferring models trained on natural images to 3d mri via position encoded slice models. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2023.
- [7] Arghya Datta, Subhrangshu Nandi, Jingcheng Xu, Greg Ver Steeg, He Xie, Anoop Kumar, and Aram Galstyan. Measuring and mitigating local instability in deep neural networks. In *Findings of the Association for Computational Linguistics: ACL 2023*, 2023.
- [8] Myrl G Marmarelis, Elizabeth Haddad, Andrew Jesson, Neda Jahanshad, Aram Galstyan, and Greg Ver Steeg. Partial identification of dose responses with hidden confounders. In *Uncertainty in Artificial Intelligence*, pages 1368–1379. PMLR, 2023.
- [9] Elizabeth Haddad, Myrl G Marmarelis, Talia M Nir, Aram Galstyan, Greg Ver Steeg, and Neda Jahanshad. Causal sensitivity analysis for hidden confounding: Modeling the sex-specific role of diet on the aging brain. In *International Workshop on Machine Learning in Clinical Neuroimaging*, pages 91–101. Springer, 2023.
- [10] Xianghao Kong, Rob Brekelmans, and Greg Ver Steeg. Information-theoretic diffusion. In *International Conference on Learning Representations*, 2023.
- [11] Tamoghna Chattopadhyay, Amit Singh, Neha Ann Joshy, Sophia I. Thomopoulos, Talia M. Nir, Hong Zheng, Elnaz Nourollahimoghadam, Umang Gupta, Greg Ver Steeg, Neda Jahanshad, and Paul M. Thompson. Predicting dementia severity by merging anatomical and diffusion MRI with deep 3D convolutional neural networks. In Jorge Brieva, Pamela Guevara, Natasha Lepore, Marius G. Linguraru, Leticia Rittner, and Eduardo Romero Castro M.D., editors, *18th International Symposium on Medical Information Processing and Analysis*, volume 12567, page 125670C. International Society for Optics and Photonics, SPIE, 2023.

- [12] Marcin Abram, Keith Burghardt, Greg Ver Steeg, Aram Galstyan, and Remi Dingreville. Inferring topological transitions in pattern-forming processes with self-supervised learning. *npj Computational Materials*, 8(1):1–12, 2022.
- [13] Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Formal limitations of sample-wise information-theoretic generalization bounds. In *2022 IEEE Information Theory Workshop (ITW)*, pages 440–445. IEEE, 2022.
- [14] Rob Brekelmans, Sicong Huang, Marzyeh Ghassemi, Greg Ver Steeg, Roger Baker Grosse, and Alireza Makhzani. Improving mutual information estimation with annealed and energy-based bounds. In *International Conference on Learning Representations (ICLR)*, 2022.
- [15] Tigran Galstyan, Hrayr Harutyunyan, Hrant Khachatryan, Greg Ver Steeg, and Aram Galstyan. Failure modes of domain generalization algorithms. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [16] Elan Markowitz, Keshav Balasubramanian, Mehrnoosh Mirtaheri, Murali Annavaram, Aram Galstyan, and Greg Ver Steeg. Statik: Structure and text for inductive knowledge graph completion. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 604–615, 2022.
- [17] Myrl G Marmarelis, Greg Ver Steeg, and Aram Galstyan. A metric space for point process excitations. *Journal of Artificial Intelligence Research*, 2022.
- [18] Qiang Li, Greg Ver Steeg, Shujian Yu, and Jesus Malo. Functional connectome of the human brain with total correlation. *Entropy*, 24(12):1725, 2022.
- [19] Nikhil J Dhinagar, Sophia I Thomopoulos, Conor Owens-Walton, Dimitris Stripelis, Jose Luis Ambite, Greg Ver Steeg, and Paul M Thompson. Alzheimer’s disease detection with a 3d convolutional neural network using gray matter maps from t1-weighted brain mri. *Alzheimer’s & Dementia*, 18:e066446, 2022.
- [20] Umang Gupta, Jwala Dhamala, Varun Kumar, Apurv Verma, Yada Pruksachatkun, Satyapriya Krishna, Rahul Gupta, Kai-Wei Chang, Greg Ver Steeg, and Aram Galstyan. Mitigating gender bias in distilled language models via counterfactual role reversal. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 658–678, 2022.
- [21] Judith Gaspers, Anoop Kumar, Greg Ver Steeg, and Aram Galstyan. Temporal generalization for spoken language understanding. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 37–44, 2022.
- [22] Myrl G Marmarelis, Greg Ver Steeg, and Aram Galstyan. Bounding the effects of continuous treatments for hidden confounders. In *NeurIPS 2022 Workshop on Causal Machine Learning for Real-World Impact*, 2022.
- [23] Dimitris Stripelis, Umang Gupta, Nikhil Dhinagar, Greg Ver Steeg, Paul M Thompson, and José Luis Ambite. Towards sparsified federated neuroimaging models via weight pruning. In *International Workshop on Distributed, Collaborative, and Federated Learning, Workshop on Affordable Healthcare and AI for Resource Diverse Global Health*, pages 141–151. Springer, 2022.
- [24] Nikhil J. Dhinagar, Sophia I. Thomopoulos, Priya Rajagopalan, Dimitris Stripelis, Jose Luis Ambite, Greg Ver Steeg, and Paul M. Thompson. Evaluation of transfer learning methods for detecting Alzheimer’s disease with brain MRI. In Jorge Brieva, Pamela Guevara, Natasha Lepore, Marius G. Linguraru, Letícia Rittner, and Eduardo Romero Castro M.D., editors, *18th International Symposium on Medical Information Processing and Analysis*, volume 12567, page 125671L. International Society for Optics and Photonics, SPIE, 2023.

- [25] Fei Wang, Kuan-Hao Huang, Anoop Kumar, Aram Galstyan, Greg Ver Steeg, and Kai-Wei Chang. Zero-shot cross-lingual sequence tagging as seq2seq generation for joint intent classification and slot filling. In *Proceedings of the Massively Multilingual Natural Language Understanding Workshop (MMNLU-22)*, pages 53–61, 2022.
- [26] Greg Ver Steeg and Aram Galstyan. Hamiltonian dynamics with non-newtonian momentum for rapid sampling. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [27] Hrayr Harutyunyan, Maxim Raginsky, Greg Ver Steeg, and Aram Galstyan. Information-theoretic generalization bounds for black-box learning algorithms. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [28] Sami Abu-El-Haija, Hesham Mostafa, Marcel Nassar, Valentino Crespi, Greg Ver Steeg, and Aram Galstyan. Implicit svd for graph representation learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [29] Dimitris Stripelis, Hamza Saleem, Tanmay Ghai, Nikhil Dhinagar, Umang Gupta, Chrysovalantis Anastasiou, Greg Ver Steeg, Srivatsan Ravi, Muhammad Naveed, Paul M. Thompson, and Jose Luis Ambite. Secure neuroimaging analysis using federated learning with homomorphic encryption. In *International Symposium on Medical Information Processing and Analysis (SIPAIM)*, 2021.
- [30] Nikhil J. Dhinagar, Sophia I. Thomopoulos, Conor Owens-Walton, Dimitris Stripelis, Jose Luis Ambite, Greg Ver Steeg, Daniel Weintraub, Philip Cook, Corey McMillan, and Paul M. Thompson. 3d convolutional neural networks for classification of alzheimer’s and parkinson’s disease with t1-weighted brain mri. In *International Symposium on Medical Information Processing and Analysis (SIPAIM)*, 2021.
- [31] Vaden Masrani, Rob Brekelmans, Thang Bui, Frank Nielsen, Aram Galstyan, Greg Ver Steeg, and Frank Wood. q-paths: Generalizing the geometric annealing path using power means. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2021.
- [32] Umang Gupta, Dimitris Stripelis, Pradeep K. Lam, Paul Thompson, Jose Luis Ambite, and Greg Ver Steeg. Membership inference attacks on deep regression models for neuroimaging. In *Medical Imaging with Deep Learning*, 2021.
- [33] Mehrnoosh Mirtaheri, Sami Abu-El-Haija, Fred Morstatter, Greg Ver Steeg, and Aram Galstyan. Identifying and analyzing cryptocurrency manipulations in social media. *IEEE Transactions on Computational Social Systems*, 8(3):607–617, 2021.
- [34] Elan Sopher Markowitz, Keshav Balasubramanian, Mehrnoosh Mirtaheri, Sami Abu-El-Haija, Bryan Perozzi, Greg Ver Steeg, and Aram Galstyan. Graph traversal with tensor functionals: A meta-algorithm for scalable learning. In *International Conference on Learning Representations (ICLR)*, 2021.
- [35] Umang Gupta, Pradeep Lam, Greg Ver Steeg, and Paul Thompson. Improved brain age estimation with slice-based set networks. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2021.
- [36] Kyle Reing, Greg Ver Steeg, and Aram Galstyan. Influence decompositions for neural network attribution. In *The 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.
- [37] Kyle Reing, Greg Ver Steeg, and Aram Galstyan. Discovering higher-order interactions through neural information decomposition. *Entropy*, 23(1), 2021.
- [38] Umang Gupta, Aaron Ferber, Bistra Dilkina, and Greg Ver Steeg. Controllable guarantees for fair outcomes via contrastive information estimation. In *AAAI Conference on Artificial Intelligence (AAAI-21)*, 2021.

- [39] James O'Neill, Greg Ver Steeg, and Aram Galstyan. Layer-wise neural network compression via layer fusion. In *Asian Conference on Machine Learning (ACML)*, 2021.
- [40] Rob Brekelmans, Vaden Masrani, Frank Wood, Greg Ver Steeg, and Aram Galstyan. All in the exponential family: Bregman duality in thermodynamic variational inference. In *International Conference on Machine Learning (ICML)*, 2020.
- [41] Hrayr Harutyunyan, Kyle Reing, Greg Ver Steeg, and Aram Galstyan. Improving generalization by controlling label-noise information in neural network weights. In *International Conference on Machine Learning (ICML)*, 2020.
- [42] Daniel Moyer, Greg Ver Steeg, Chantal M. W. Tax, and Paul M. Thompson. Scanner invariant representations for diffusion mri harmonization. *Magnetic Resonance in Medicine*, 84(4):2174–2189, 2020.
- [43] Daniel Moyer, Greg Ver Steeg, and Paul M Thompson. Overview of scanner invariant representations. In *Medical Imaging with Deep Learning*, 2020.
- [44] Ayush Jaiswal, Daniel Moyer, Greg Ver Steeg, Wael AbdAlmageed, and Premkumar Natarajan. Invariant representations through adversarial forgetting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-20)*, 2020.
- [45] Sahil Garg, Guillermo Cecchi, Irina Rish, Palash Goyal, Sarik Ghazarian, Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. Modeling dialogues with hashcode representations: A nonparametric approach. In *AAAI Conference on Artificial Intelligence (AAAI-20)*, 2020.
- [46] Greg Ver Steeg, Hrayr Harutyunyan, Daniel Moyer, and Aram Galstyan. Fast structure learning with modular regularization. In *Advances in Neural Information Processing Systems*, pages 15567–15577, 2019.
- [47] Rob Brekelmans, Daniel Moyer, Aram Galstyan, and Greg Ver Steeg. Exact rate-distortion in autoencoders via echo noise. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [48] Sahil Garg, Aram Galstyan, Greg Ver Steeg, and Guillermo A Cecchi. Nearly-unsupervised hashcode representations for biomedical relation extraction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4017–4027, 2019.
- [49] Kyle Reing, Greg Ver Steeg, and Aram Galstyan. Maximizing multivariate information with error-correcting codes. *IEEE Transactions on Information Theory*, 2019.
- [50] Shirley Pepke, William M Nelson, and Greg Ver Steeg. Analyzing tumor gene expression factors with the corexplorer web portal. *JoVE (Journal of Visualized Experiments)*, (152):e60431, 2019.
- [51] Hrayr Harutyunyan, Hrant Khachatrian, David C. Kale, Greg Ver Steeg, and Aram Galstyan. Multitask learning and benchmarking with clinical time series data. *Nature Scientific Data*, 6(1):96, 2019.
- [52] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *International Conference on Machine Learning*, pages 21–29, 2019.
- [53] Sahil Garg, Aram Galstyan, Greg Ver Steeg, Irina Rish, Guillermo Cecchi, and Shuyang Gao. Kernelized hashcode representations for biomedical relation extraction. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.

- [54] Shuyang Gao, Robert Brekelmans, Greg Ver Steeg, and Aram Galstyan. Auto-encoding correlation explanation. In *Proceedings of the 22nd International Conference on AI and Statistics (AISTATS)*, 2019.
- [55] Daniel Moyer, Shuyang Gao, Rob Brekelmans, Greg Ver Steeg, and Aram Galstyan. Invariant representation without adversarial training. In *Advances In Neural Information Processing Systems*, 2018.
- [56] Brandalyn Cherish Riedel, Madelaine Daianu, Greg Ver Steeg, Adam Mezher, Lauren E Salminen, Aram Galstyan, and Paul Matthew Thompson. Uncovering biologically coherent peripheral signatures of health and risk for alzheimer’s disease in the aging brain. *Frontiers in Aging Neuroscience*, 10:390, 2018.
- [57] Daniel Moyer, Paul M Thompson, and Greg Ver Steeg. Measures of tractography convergence. In *International Conference On Medical Image Computing & Computer Assisted Intervention (MICCAI)*, 2018.
- [58] Yolanda Gil, Ke-Thia Yao, Varun Ratnakar, Daniel Garijo, Greg Ver Steeg, Pedro Szekely, Robert Brekelmans, Mayank Kejriwal, Fanghao Luo, and I-Hui Huang. P4ml: A phased performance-based pipeline planner for automated machine learning. In *ICML AutoML Workshop*, 2018.
- [59] Neal Lawton, Greg Ver Steeg, and Aram Galstyan. A forest mixture bound for block-free parallel inference. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018.
- [60] Greg Ver Steeg, Rob Brekelmans, Hrayr Harutyunyan, and Aram Galstyan. Disentangled representations via synergy minimization. In *55th Annual Allerton Conference on Communication, Control, and Computing*, 2017.
- [61] R.J. Gallagher, K. Reing, D. Kale, and G. Ver Steeg. Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics*, 2017.
- [62] Greg Ver Steeg. Unsupervised learning via total correlation explanation. In *IJCAI*, 2017.
- [63] David Stück, Haraldur Tómas Hallgrímsson, Greg Ver Steeg, Alessandro Epasto, and Luca Foschini. The spread of physical activity through social networks. In *Proceedings of the 26th International Conference on World Wide Web*, pages 519–528. International World Wide Web Conferences Steering Committee, 2017.
- [64] Greg Ver Steeg, Shuyang Gao, Kyle Reing, and Aram Galstyan. Sifting common information from many variables. *IJCAI*, 2017.
- [65] Shirley Pepke and Greg Ver Steeg. Comprehensive discovery of subsample gene expression components by information explanation: therapeutic implications in cancer. *BMC medical genomics*, 10(1):12, 2017.
- [66] Artemis Zavaliangos-Petropulu, Emily L Dennis, Greg Ver Steeg, Talin Babikian, Richard Mink, Christopher Babbitt, Jeffrey Johnson, Christopher C Giza, Robert F Asarnow, and Paul M Thompson. Variable clustering reveals associations between subcortical brain volume and cognitive changes in pediatric traumatic brain injury. In *12th International Symposium on Medical Information Processing and Analysis*. International Society for Optics and Photonics, 2017.
- [67] Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. Variational information maximization for feature selection. In *Advances In Neural Information Processing Systems*, pages 487–495, 2016.
- [68] Kyle Reing, David C. Kale, Greg Ver Steeg, and Aram Galstyan. Toward interpretable topic discovery via anchored correlation explanation. In *ICML Workshop on Human Interpretability in Machine Learning (WHI 2016)*, 2016.

- [69] Linhong Zhu, Dong Guo, Junming Yin, Greg Ver Steeg, and Aram Galstyan. Scalable temporal latent space inference for link prediction in dynamic social networks. *IEEE Transactions on Knowledge and Data Engineering*, 2016.
- [70] Greg Ver Steeg and Aram Galstyan. The information sieve. In *International Conference on Machine Learning (ICML)*, 2016.
- [71] Sarah Madsen, Greg Ver Steeg, Madelaine Daianu, Adam Mezher, Neda Jahanshad, Talia M. Nir, Xue Hua, Boris A. Gutman, Aram Galstyan, and Paul M. Thompson. Relative value of diverse brain mri and blood-based biomarkers for predicting cognitive decline in the elderly. In *SPIE Medical Imaging*, 2016.
- [72] Yoon-Sik Cho, Greg Ver Steeg, Emilio Ferrara, and Aram Galstyan. Latent space model for multi-modal social data. In *Proceedings of World Wide Web Conference (WWW)*, 2016.
- [73] Armen Allahverdyan, Greg Ver Steeg, and Aram Galstyan. Memory-induced mechanism for self-sustaining cascades in networks. *Physical Review E*, 2015.
- [74] Daniel Moyer, Boris Gutman, Gautam Prasad, , Greg Ver Steeg, and Paul Thompson. Mixed membership stochastic blockmodels for the human connectome. In *Proceedings of Bayesian and Graphical Imaging for Biomedical Imaging (BAMBI)*, 2015.
- [75] Madelaine Daianu, Greg Ver Steeg, Adam Mezher, Neda Jahanshad, Talia M. Nir, Xiaoran Yan, Gautam Prasad, Kristina Lerman, Aram Galstyan, and Paul M. Thompson. Information-theoretic clustering of neuroimaging metrics related to cognitive decline in the elderly. In *Proceedings of the MICCAI Workshop on Medical Computer Vision*, 2015.
- [76] Daniel Moyer, Boris Gutman, Gautam Prasad, Joshua Faskowitz, Greg Ver Steeg, and Paul Thompson. Blockmodels for connectome analysis. In *11th International Symposium on Medical Information Processing and Analysis (SIPAIM 2015)*, 2015.
- [77] Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. Estimating mutual information by local gaussian approximation. In *Uncertainty in Artificial Intelligence (UAI)*, 2015.
- [78] Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. Understanding confounding effects in linguistic coordination: an information-theoretic approach. *PLoS ONE*, 10(6): e0130167, 2015.
- [79] Sarah K. Madsen, Greg Ver Steeg, Adam Mezher, Neda Jahanshad, Talia M. Nir, Xue Hua, Boris A. Gutman, Aram Galstyan, and Paul M. Thompson. Information-theoretic characterization of blood panel predictors for brain atrophy and cognitive decline in the elderly. *IEEE International Symposium on Biomedical Imaging*, 2015.
- [80] Greg Ver Steeg and Aram Galstyan. Maximally informative hierarchical representations of high-dimensional data. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [81] Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. Efficient estimation of mutual information for strongly dependent variables. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [82] Nathan Hodas, Greg Ver Steeg, Joshua Harrison, Satish Chikkagoudar, Eric Bell, and Courtney Corley. Disentangling the lexicons of disaster response in twitter. In *The 3rd International Workshop on Social Web for Disaster Management (SWDM'15)*, 2015.
- [83] Christoph Adami and Greg L. Ver Steeg. Black holes are almost optimal quantum cloners. *Journal of Physics A*, 58, 2015.

- [84] Greg Ver Steeg and Aram Galstyan. Discovering structure in high-dimensional data through correlation explanation. *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [85] Greg Ver Steeg, Aram Galstyan, Fei Sha, and Simon DeDeo. Demystifying information-theoretic clustering. In *International Conference on Machine Learning (ICML)*, 2014.
- [86] Yoon Sik Cho, Greg Ver Steeg, and Aram Galstyan. Where and why users “check in”. In *Proc. of the Twenty-Eighth Conference on Artificial Intelligence (AAAI)*, 2014.
- [87] Siddhartha Santra, Greg Quiroz, Greg Ver Steeg, and Daniel Lidar. Max 2-sat with up to 108 qubits. *New J. Phys.*, 16:045006, 2014.
- [88] Yoon-Sik Cho, Greg Ver Steeg, and Aram Galstyan. Mixed membership blockmodels for dynamic networks with feedback. In *Handbook of Mixed Membership Models and Their Applications*. Chapman & Hall/CRC, 2014.
- [89] Greg Ver Steeg, Christopher Moore, Aram Galstyan, and Armen Allahverdyan. Phase transitions in community detection: A solvable toy model. *EPL (Europhysics Letters)*, 106(4):48004, 2014.
- [90] Christoph Adami and Greg Ver Steeg. Classical information transmission capacity of quantum black holes. *Classical and Quantum Gravity*, 31(7):075015, 2014.
- [91] Vasanthan Raghavan, Greg Ver Steeg, Aram Galstyan, and Alexander G Tartakovsky. Modeling temporal activity patterns in dynamic social networks. *IEEE TCSS*, 2014.
- [92] Greg Ver Steeg and Aram Galstyan. Statistical tests for contagion in observational social network studies. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2013.
- [93] Greg Ver Steeg and Aram Galstyan. Information-theoretic measures of influence based on content dynamics. In *Proceedings of the 6th International Conference on Web Search and Data Mining (WSDM)*. ACM, 2013.
- [94] Yoon Sik Cho, Greg Ver Steeg, and Aram Galstyan. Socially relevant venue clustering from check-in data. In *MLG’13: Mining and Learning from Graphs*, 2013.
- [95] Greg Ver Steeg and Aram Galstyan. Information transfer in social media. In *Proceedings of World Wide Web Conference (WWW)*, 2012.
- [96] Greg Ver Steeg and Aram Galstyan. A sequence of relaxations constraining hidden variable models. In *Proc. of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI)*, 2011. **Best paper runner-up award.**
- [97] Greg Ver Steeg, Rumi Ghosh, and Kristina Lerman. What stops social epidemics? In *Proc. 5th Int. AAAI Conf. on Weblogs and Social Media (ICWSM)*, 2011.
- [98] Yoon Sik Cho, Greg Ver Steeg, and Aram Galstyan. Co-evolution of selection and influence in social networks. In *Proc. of the Twenty-Fifth Conference on Artificial Intelligence (AAAI)*, 2011.
- [99] Aram Galstyan, Greg Ver Steeg, and Armen Allahverdyan. Statistical mechanics of semi-supervised clustering in sparse graphs. *J. Stat. Mech.*, (P08009), 2011.
- [100] A. E. Allahverdyan, Greg Ver Steeg, and A. Galstyan. Community detection with and without prior information. *EPL (Europhysics Letters)*, 90(1):18002, 2010.
- [101] Greg Ver Steeg. *Foundational aspects of nonlocality*. PhD thesis, California Institute of Technology, 2009.

- [102] Greg Ver Steeg and N. C. Menicucci. Entangling power of an expanding universe. *Physical Review D*, 79(4), February 2009.
- [103] Greg Ver Steeg and S. Wehner. Relaxed uncertainty relations and information processing. *Quantum Information and Computation*, 9:0801–0832, 2009.
- [104] P. L. Bartlett, I. Bray, S. Jones, A. T. Stelbovics, A. S. Kadyrov, K. Bartschat, Greg Ver Steeg, M. P. Scott, and P. G. Burke. Unambiguous ionization amplitudes for electron-hydrogen scattering. *Physical Review A*, 68(2), August 2003.
- [105] K. Bartschat, M. P. Scott, P. G. Burke, T. Stitt, N. S. Scott, A. N. Grum-Grzhimailo, S. Riordan, Greg Ver Steeg, and S. I. Strakhova. Convergence of energy-differential ionization cross sections obtained from a T-matrix approach with R-matrix wave functions. *Physical Review A*, 65(6), June 2002.
- [106] K. Bartschat, S. Riordan, and Greg Ver Steeg. Extraction of energy-differential ionization cross sections in time-dependent calculations. *Physical Review A*, 65(6), June 2002.
- [107] Greg L. Ver Steeg, K. Bartschat, and I. Bray. Time-dependent model calculations for a molecular hydrogen ion in a strong ultra-short laser pulse. *Journal of Physics B Atomic Molecular Physics*, 36:3325–3336, August 2003.

Note: these papers span different disciplines including physics, computer science, and biology. Listed are all peer-reviewed publications including journal articles and conference proceedings. In physics, conferences typically do not have papers (only abstracts, not included here). Peer reviewed conference proceedings in CS and biology are included. In CS, conference proceedings are, depending on the conference, highly selective and considered top venues in the field (e.g., NeurIPS/NIPS, ICML, WWW, ICLR, CVPR). Journals in CS are typically used for longer version of work appearing originally in conferences. In neuroscience, conference papers are peer reviewed, but they are short and generally regarded as second tier publications mainly used for work in progress.

## Recent Invited Talks

DEMICS Workshop, Max Planck Institute, Dresden	2023
Disney Research	2023
ISI AI Seminar	2023
ITA Workshop	2023
UCR Data Science Seminar	2023
USC AI4Health Workshop	2022
Keynote invite for the AAAI-22 workshop on Information Theory for Deep Learning (declined)	2022
Amazon Machine Learning Conference	2021
Amazon Caltech Tech Talk	2019
Beyond Shannon Workshop at IAS Amsterdam	2019
UCSB Kavli Institute for Theoretical Physics, At the Crossroad of Physics and ML	2019
Santa Fe Institute	2018
Physics and Machine Learning Conference at Los Alamos	2018
Allerton Conference	2017
Berkeley DataEDGE conference	2017
Los Alamos National Lab.	2017
IJCAI Early Career Spotlight talk	2017
ITA Workshop	2017

Vicarious	2016
*W.V.T. Rusch Engineering Honors Colloquium	2016
Headspace	2016
Deep learning for health care workshop at MLHC	2016
StatPhys26 Workshop: Statistical physics methods in biology and computer science	2016
IPAM Cultural Analytics workshop	2016
ITA Workshop	2016
UCLA, Soatto group seminar	2016
UC Davis, Crutchfield group seminar	2016
Caltech, IST Seminar	2015
UC San Diego, AI seminar	2015
OpenX	2015
Brain Corporation	2015
*1st Collaborative Data Science Symposium hosted by Ameritrade	2015
Joint Symposium on Neural Computation at USC	2015
Machine Learning LA at eHarmony	2015
Information Theory and Applications Workshop	2015
*AFRL Rome, "Distinguished Speaker Seminar Series in Machine Intelligence and Autonomy"	2014
Santa Fe Institute, workshop on "Statistical Mechanics of Complexity"	2014
IPAM Mathematics of Social Learning Workshop	2014
ISI, Natural Language Seminar	2013
ID Analytics, "Information-Theoretic Tools for Social Media"	2013
Santa Fe Institute, workshop on "Structure, Statistical Inference and Dynamics in Networks"	2013
Sante Fe Institute, "Information-Theoretic Tools for Social Media"	2012
Keynote for "Making Sense of Microposts" workshop at WWW 2012.	2012
SAP Research, Singapore	2012
LARC, Singapore Management University seminar.	2012
UC Irvine, AI-ML seminar.	2012

\* *Talks with honorarium*

## Technical Reports

Ayush Jaiswal, Rob Brekelmans, Daniel Moyer, Greg Ver Steeg, Wael AbdAlmageed, Premkumar Natarajan. Discovery and Separation of Features for Invariant Representation Learning. <https://arxiv.org/abs/1912.00646>

Hrayr Harutyunyan, Daniel Moyer, Hrant Khachatrian, Greg Ver Steeg, and Aram Galstyan. Efficient Covariance Estimation from Temporal Data. <https://arxiv.org/abs/1905.13276>.

Wenzhe Li, Dong Guo, Greg Ver Steeg, Aram Galstyan. Unifying Local and Global Change Detection in Dynamic Networks. arXiv:1710.03035, 2017

## Workshop Presentations

NeurIPS 22, Federated Learning workshop, Federated Progressive Sparsification (Purge-Merge-Tune)+

NeurIPS 22, C4ML workshop, Bounding the Effects of Continuous Treatments for Hidden Confounders

N Mehrabi, U Gupta, F Morstatter, GV Steeg, A Galstyan. "Attributing Fair Decisions with Attention Interventions" arXiv:2109.03952, NAACL TrustNLP workshop, 2022.

TMS Annual Meeting, Inferring topological transitions in pattern-forming processes via self-supervised learning, 2022

Workshop on Privacy Enhancing Technologies for the Homeland Security Enterprise 2022

Attributing Fair Decisions with Attention Interventions. NAACL TrustNLP workshop, 2022.

Nikhil Dhinagar, . . . , Greg Ver Steeg, Paul Thompson "Inter and Intra-domain Pre-training for Alzheimer's Disease Classification using Brain MRI", OHBM, 2022.

S. Abu-El-Haija, V. Crespi, G. Ver Steeg, and A. Galstyan. Fast GRL with Unique Optimal Solutions. In Proceedings of the ICLR 2021 Workshop on Geometrical and Topological Representation Learning, 2021.

Rob Brekelmans, Vaden Masrani, Thang D Bui, Frank Wood, Aram Galstyan, Greg Ver Steeg, Frank Nielsen. "Annealed Importance Sampling with q-Paths" <https://openreview.net/forum?id=ZBJ20FRVDP> (Best paper award at NeurIPS 2020 Workshop on Deep Learning through Information Geometry)

Rob Brekelmans, Frank Nielsen, Aram Galstyan, Greg Ver Steeg. "Likelihood Ratio Exponential Families " <https://openreview.net/forum?id=JceRB2cd7G> NeurIPS Deep Learning with Information Geometry 2020

Tigran Galstyan, Hrant Khachatrian, Greg Ver Steeg, Aram Galstyan. "Robust Classification under Class-Dependent Domain Shift". arXiv:2007.05335 (ICML 2020 workshop on Uncertainty and Robustness in Deep Learning)

Tigran Galstyan, Hrant Khachatrian, Greg Ver Steeg, Aram Galstyan "Robust Classification under Class-Dependent Domain Shift". ICML 2020 Workshop on Uncertainty & Robustness in Deep Learning.

NeurIPS 2019 Workshop on Information Theory and Machine Learning

Hrayr Harutyunyan, Kyle Reing, Greg Ver Steeg, Aram Galstyan. Reducing overfitting by minimizing label information in weights. NeurIPS Workshop on Information Theory in Machine Learning, 2019.

Rob Brekelmans, Aram Galstyan, Greg Ver Steeg. Understanding Thermodynamic Variational Inference. NeurIPS Workshop on Information Theory in Machine Learning, 2019.

Amol Kapoor, Sami Abu-El-Haija, Bryan Perozzi, Aram Galstyan, Greg Ver Steeg, Hrayr Harutyunyan, Kristina Lerman and Nazanin Alipourfard. MixHop: Higher-Order Graph Convolutional Architectures via Sparsified Neighborhood Mixing. Workshop on Machine Learning with Graphs, 2019.

Mehrnoosh Mirtaheri, Sami Abu-El-Haija, Fred Morstatter, Greg Ver Steeg, Aram Galstyan. Identifying and Analyzing Cryptocurrency Manipulations in Social Media. Cybersafety workshop, 2019.

Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Hrayr Harutyunyan. A Higher-Order Graph Convolutional Layer. NIPS Relational Representation Learning Workshop, 2018.

Sahil Garg, Guillermo Ceechi, Irina Rish, Shuyang Gao, Bhavana Bhaskar, Greg Ver Steeg, Palash Goyal and Aram Galstyan. "Therapeutic Dialogue Modeling via Locality Sensitive Hashing". CompPsy-18 workshop at IJCAI.

- Greg Ver Steeg, Stefan Fuertinger, Kristina Simonyan, Aram Galstyan, “Increasing Statistical Power in fMRI Studies with Maximally Informative Factors”, Organization for Human Brain Mapping (OHBM) 2018.
- Hrayr Harutyunyan, Hrant Khachatryan, David Kale, Greg Ver Steeg, and Aram Galstyan. A Public Benchmark for Clinical Prediction and Multitask Learning. NIPS Machine Learning in Healthcare Workshop, 2017.
- Jahanshad, Ver Steeg, et al. “Enhancing genetic correlations between blood and brain using latent factors of correlated blood measures”, SPIE 2016.
- Kyle Reing, David Kale, Greg Ver Steeg, Aram Galstyan. ICML Workshop on Interpretability in Machine Learning, 2016.
- Madsen, Ver Steeg, et al. “Predicting Cognitive Decline with Information-Theoretic Clustering of Brain MRI and Blood Tests.” Society of Biological Psychiatry, 2015.
- NIPS workshop on Machine Learning in Computational Biology (MLCB), 2014.
- NIPS workshop on Machine Learning and Interpretation in Neuroimaging (MLINI), 2014.
- Greg Ver Steeg and Aram Galstyan. “Minimal Assumptions Tests for Observational Social Network Studies”. In *Workshop on Information in Networks*, 2013.
- Shuyang Gao, Greg Ver Steeg and Aram Galstyan. “Explaining Away Stylistic Coordination”. In *Workshop on Information in Networks*, 2013.
- Greg Ver Steeg “Optimization via Open System Quantum Annealing.” *ARO Quantum Computing Review*, 2013.
- Vasanthan Raghavan, Alexander G. Tartakovsky, Aram Galstyan, and Greg Ver Steeg. Coupled Hidden Markov Models for User Activity in Social Networks. In *IEEE 2nd International Workshop on Social Multimedia Research*, 2013.
- Marco Huesch, Greg Ver Steeg and Aram Galstyan. Vaccination (Anti-) Campaigns in Social Media in *AAAI-13 Workshop on Expanding the Boundaries of Health Informatics Using AI (HIAI)*, 2013.
- Greg Ver Steeg and Aram Galstyan. Information-Theoretic Measures of Influence Based on Content Dynamics. In *Workshop on Information in Networks(WIN)*, 2012.
- Greg Ver Steeg and Aram Galstyan. Information transfer in social media. In *Workshop on Information in Networks(WIN)*, 2011.
- Greg Ver Steeg, Armen Allahverdyan, and Aram Galstyan. Semi-supervised clustering in sparse networks. In *International Conference on Complex Systems(ICCS)*, 2011.
- Yoon Sik Cho, Greg Ver Steeg, and Aram Galstyan. Co-evolving mixed membership blockmodels. In *NIPS workshop on Networks Across Disciplines*, 2010.
- Greg Ver Steeg and Aram Galstyan. Ruling out latent homophily in social networks. In *NIPS workshop on Social Computing*, 2010.
- Armen Allahverdyan, Aram Galstyan, and Greg Ver Steeg. Clustering with prior information. In *NIPS workshop: Clustering: Science or Art? Towards Principled Approaches*, 2009.