

# MATT J. KUSNER

<http://mkusner.github.io/> ◊ [m.kusner@ucl.ac.uk](mailto:m.kusner@ucl.ac.uk) ◊ <https://github.com/mkusner>

## EMPLOYMENT

---

**University College London** *September 2019 - Present*

Associate Professor in Machine Learning, Department of Computer Science

**Google** *November 2021 - November 2022*

Visiting Faculty Researcher

**University of Oxford** *October 2018 - September 2019*

Associate Professor in Machine Learning, Department of Computer Science

Tutorial Fellow at Jesus College

**European Lab for Learning & Intelligent Systems (ELLIS)** *Jul 2021 - Present*

ELLIS Scholar

**The Alan Turing Institute**

Turing Fellow

Research Fellow

*October 2018 - Present*

*October 2016 - October 2018*

## EDUCATION

---

**Cornell University** *August 2015 - August 2016*

Visiting Ph.D. student in Dept. of Computer Science

Advisor: Kilian Q. Weinberger

**Washington University in St. Louis** *August 2011 - August 2016*

Ph.D. from Dept. of Computer Science & Engineering

Advisor: Kilian Q. Weinberger

## PROFESSIONAL ACTIVITIES

---

**Conference on Causal Learning and Reasoning (CLear)** *April, 2024*

*Communications Chair*

*Los Angeles, California*

**NeurIPS Workshop: Algorithmic Fairness through the Lens of Causality and Privacy** *2022*

*Advisory Committee*

**NeurIPS Workshop: Algorithmic Fairness through the Lens of Causality and Interpretability**

*2020 Co-organizer*

**NeurIPS Workshop: Machine Learning for Molecules** *December, 2020*

*Co-organizer*

**NIPS Workshop: Critiquing and Correcting Trends in Machine Learning** *December, 2018*

*Co-organizer*

*Montreal, Canada*

**NIPS Workshop: Machine Learning for Molecules and Materials** *December, 2017; 2018*

*Co-organizer*

*Long Beach, CA; Montreal, Canada*

**Uncertainty in Artificial Intelligence (UAI) Conference** *August, 2018*

*Publications Chair*

*Monterey, CA*

<b>NIPS Press Conference</b> <i>Invited Speaker</i>	December, 2017 Long Beach, CA
<b>NIPS Workshop on Machine Learning for Molecules and Materials</b> <i>Co-organizer</i>	December, 2017 Long Beach, CA
<b>ICML Conference</b> <i>Workflow Chair</i>	June, 2016 New York, NY
<b>ICML Workshop: Resource-Efficient Machine Learning</b> <i>Co-organizer</i>	July, 2015 Lille, France
<b>ICML Workshop: Learning with Test-Time Budgets</b> <i>Co-organizer</i>	June, 2013 Atlanta, GA
<b>AI2050 Early Career Fellowship Reviewer</b> <i>Schmidt Futures</i>	September, 2023
<b>ICLR 2023 Notable Area Chair</b>	May, 2023
<b>NWO Open Competition</b> <i>Dutch Research Council</i>	May, 2021
<b>Patterns, Predictions, and Actions</b> <i>MIT Press</i>	April, 2021
<b>ICML 2020 Top 33% Reviewer</b>	June, 2020
<b>ICML 2019 Top 5% Reviewer</b>	June, 2019
<b>NeurIPS 2018 Top 30% Reviewer</b>	December, 2018
<b>Area Chair</b> NeurIPS, ICLR, ICML	
<b>Program Committee</b> NeurIPS, ICML, ICLR, AISTATS, FAT*, JMLR, AAAI, KDD	

## PUBLICITY

---

<b>Fast Company</b> <a href="https://tinyurl.com/46y4ht8z">https://tinyurl.com/46y4ht8z</a>	August 2023
<b>Forbes</b> <a href="https://tinyurl.com/bdb84kd9">https://tinyurl.com/bdb84kd9</a>	November 2021
<b>UK Government Centre for Data Ethics and Innovation</b> <a href="https://tinyurl.com/76wfe7rb">https://tinyurl.com/76wfe7rb</a>	November 2020
<b>Harvard Business Review</b> <a href="https://tinyurl.com/3yatpnc4">https://tinyurl.com/3yatpnc4</a>	October 2020
<b>Wired</b> <a href="https://tinyurl.com/y29n58tl">https://tinyurl.com/y29n58tl</a>	February 2019
<b>Forbes</b> <a href="https://tinyurl.com/yxbrpwxxz">https://tinyurl.com/yxbrpwxxz</a>	March 2018
<b>The Guardian</b> <a href="https://tinyurl.com/y893qsto">https://tinyurl.com/y893qsto</a>	August 2017

<b>The New Scientist</b> <i><a href="https://tinyurl.com/14zfkv2">https://tinyurl.com/14zfkv2</a></i>	March 2017
<b>The Future of Life Institute</b> <i><a href="https://tinyurl.com/y3xgnmgy">https://tinyurl.com/y3xgnmgy</a></i>	December 2015
<b>OpenTable</b> <i><a href="https://tinyurl.com/y3ohyyw3">https://tinyurl.com/y3ohyyw3</a></i>	August 2015

## INVITED TALKS

---

<b>Babel</b> <i>The EU AI Act: Implications for the Technology Sector</i>	February 21, 2023
<b>Pint of Science</b>	May 11, 2022
<b>The Federal Reserve Banks of Cleveland and Philadelphia</b> <i>Causality for Fair Lending</i>	November 10, 2021
<b>RIKEN High-dimensional Statistical Modeling Team Seminar</b>	November 9, 2021
<b>ELLIS Workshop on Causethical ML</b>	July 26, 2021
<b>Cambridge Tech &amp; Society</b>	November 19, 2020
<b>Oxford CS Open Day</b> <i>Causality</i>	May 11, 2019
<b>Faculty of Law, Oxford University</b> <i>AI for English Law Conference</i>	March 18, 2019 Panel on AI and Ethics
<b>Birmingham UAS Student Conference</b>	March 19, 2019
<b>The Alan Turing Institute</b> <i>To support the Information Commissioners Office</i>	March 13, 2019 Explainability Expert Roundtable
<b>Royal Academy of Engineering</b> <i>To support the Department of Digital, Culture, Media and Sport</i>	November 5, 2018 Algorithmic Bias Roundtable
<b>Talking Machines Podcast</b>	November 1, 2018
<b>AI and ML in Cambridge (CamAIML)</b>	March 15, 2018
<b>Cambridge Centre for Mathematical Sciences</b>	February 20, 2018
<b>Oxford Computational Statistics and Machine Learning Seminar</b>	February 16, 2018
<b>The Royal Society</b>	February 12, 2018
<b>Cambridge University Engineering Department</b>	September 12, 2017

## TEACHING

---

<b>University College London</b>			
Applied Machine Learning	COMP0081	193 students (Masters and undergraduate)	Spring 2024
Applied Machine Learning	COMP0081	194 students (Masters and undergraduate)	Spring 2023
Applied Machine Learning	COMP0081	211 students (Masters and undergraduate)	Spring 2022
Applied Machine Learning	COMP0081	181 students (Masters and undergraduate)	Spring 2021
<b>University College London</b>			
Tutorial on Causal Inference			Fall 2020

## University of Oxford

Digital Systems	Tutorial	4 students (undergraduate)	Spring 2019
Discrete Mathematics	Tutorial	6 students (undergraduate)	Spring 2019
Imperative Programming Parts 1& 2	Tutorial	4 students (undergraduate)	Spring 2019
Linear Algebra	Tutorial	6 students (undergraduate)	Fall 2018
Discrete Mathematics	Tutorial	6 students (undergraduate)	Fall 2018

## Oxford Warwick Statistics Programme

Tutorial on Fairness and Causal Inference

Spring 2018

## GRADUATE SUPERVISION

---

### PhD alumni

Limor Gultchin	defended, co-advised with Varun Kanade	Spring 2024 (Oxford)
thesis: <i>Causal and Trustworthy Machine Learning: Methods and Applications</i>		
Qi Liu	defended, co-advised with Phil Blunsom	Spring 2023 (Oxford)
thesis: <i>Better conditioning on context for natural language processing</i>		
David Watson	defended, co-advised with Luciano Floridi	Spring 2021 (Oxford)
thesis: <i>Explaining black box algorithms: epistemological challenges and machine learning solutions</i>		

### PhD students under supervision

Yuchen (Caroline) Zhu	estimated completion: Spring 2025
Jean Kaddour	estimated completion: Spring 2025
Gbétondji Dovonon	estimated completion: Spring 2027
Leo Richter	estimated completion: Spring 2028

## PHD THESIS COMMITTEES

---

**Gábor Melis**, *Towards Better Generative Models of Language*, DeepMind/UCL, Fall 2023.

**Márton Havasi**, *Advances in Compression using Probabilistic Models*, University of Cambridge, Spring 2021.

**James Townsend**, *Lossless Compression with Latent Variable Models*, UCL, Winter 2020.

**Matej Balog**, *Converting to Optimization in Machine Learning: Perturb-and-Map, Differential Privacy, and Program Synthesis*, University of Cambridge, Summer 2020.

## AWARDS

---

<b>Research Fellowship</b> <i>The Alan Turing Institute</i>	2016-2018 <i>London, UK</i>
<b>Turner Dissertation Award</b> <i>Washington University in St. Louis</i> (awarded yearly to the best Ph.D. dissertation)	December 2016 <i>St. Louis, MO</i>
<b>Konhauser Award for Mathematical Achievement</b> <i>Macalester College</i> (awarded yearly to the top senior student in computer science)	May 2011 <i>St. Paul, MN</i>

## PUBLICATIONS

---

Katherine Tsai, Stephen Pfohl, Olawale Salaudeen, Nicole Chiou, Matt Kusner, Alexander D'Amour, Sanmi Koyejo, Arthur Gretton

### Proxy Methods for Domain Adaptation

The International Conference on Artificial Intelligence and Statistics (AISTATS), 2024.

Jean Kaddour\*, Oscar Key\*, Piotr Nawrot, Pasquale Minervini, Matt J. Kusner

**No Train No Gain: Revisiting Efficient Training Algorithms For Transformer-based Language Models**

Neural Information Processing Systems (NeurIPS), 2023.

\* = authors contributing equally, listed in alphabetical order

Ibrahim Alabdulmohsin\*, Nicole Chiou\*, Alexander D'Amour\* Arthur Gretton\*, Sanmi Koyejo\*, Matt J. Kusner\*, Stephen R. Pfohl\*, Olawale Salaudeen\*, Jessica Schrouff\*, Katherine Tsai\*

**Adapting to Latent Subgroup Shifts via Concepts and Proxies**

The International Conference on Artificial Intelligence and Statistics (AISTATS), 2023.

Valentina Zantedeschi, Luca Franceschi, Jean Kaddour, Matt J. Kusner, Vlad Niculae

**DAG Learning on the Permutohedron**

The International Conference on Learning Representations (ICLR), 2023.

Kirtan Padh, Jakob Zeitler, David Watson, Matt J. Kusner, Ricardo Silva, Niki Kilbertus

**Stochastic Causal Programming for Bounding Treatment Effects**

Conference on Causal Learning and Reasoning (CLearR), 2023.

Jean Kaddour, Linqing Liu, Ricardo Silva, Matt J. Kusner

**When Do Flat Minima Optimizers Work?**

Neural Information Processing Systems (NeurIPS), 2022.

Natalie Maus, Haydn T. Jones, Juston S. Moore, Matt J. Kusner, John Bradshaw, Jacob R. Gardner

**Local Latent Space Bayesian Optimization over Structured Inputs**

Neural Information Processing Systems (NeurIPS), 2022.

Yuchen Zhu, Limor Gultchin, Arthur Gretton, Matt J. Kusner, Ricardo Silva

**Causal Inference with Treatment Measurement Error: A Nonparametric Instrumental Variable Approach**

The Conference on Uncertainty in Artificial Intelligence (UAI), 2022. *Oral Presentation*

Nitin Agrawal, James Bell, Adrià Gascón, Matt J. Kusner

**MPC-Friendly Commitments for Publicly Verifiable Covert Security**

*The Conference on Computer and Communications Security (CCS), 2021*

Jean Kaddour, Yuchen Zhu, Qi Liu, Matt J. Kusner, Ricardo Silva

**Causal Effect Inference for Structured Treatments**

*Neural Information Processing Systems (NeurIPS), 2021*

Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, Matt J. Kusner

**Unsupervised Point Cloud Pre-Training via View-Point Occlusion, Completion**

*The International Conference on Computer Vision (ICCV), 2021*

Valentina Zantedeschi, Matt J. Kusner, Vlad Niculae

**Learning Binary Decision Trees by Argmin Differentiation**

*The International Conference on Machine Learning (ICML), 2021*

Limor Gultchin, David Watson, Matt J. Kusner, Ricardo Silva

**Operationalizing Complex Causes: A Pragmatic View of Mediation**

*The International Conference on Machine Learning (ICML), 2021*

Afsaneh Mastouri\* Yuchen Zhu\*, Limor Gultchin, Anna Korba, Ricardo Silva, Matt J. Kusner, Arthur Gretton, Krikamol Muandet

**Proximal Causal Learning with Kernels: Two-Stage Estimation and Moment Restriction**

*The International Conference on Machine Learning (ICML), 2021*

\* equal contribution

Qi Liu, Matt J. Kusner, Phil Blunsom

**Counterfactual Data Augmentation for Neural Machine Translation**

*North American Chapter of the Association for Computational Linguistics (NAACL), 2021*

Niki Kilbertus, Matt J. Kusner, Ricardo Silva

**A Class of Algorithms for General Instrumental Variable Models**

*Neural Information Processing Systems (NeurIPS), 2020*

John Bradshaw, Brooks Paige, Matt J. Kusner, Marwin H. S. Segler, Jos Miguel Hernandez-Lobato

**Barking up the right tree: an approach to search over molecule synthesis DAGs**

*Neural Information Processing Systems (NeurIPS), 2020. Spotlight Presentation*

Matt J. Kusner, Joshua R. Loftus

**The long road to fairer algorithms**

*Nature (Comment), 2020*

Limor Gultchin, Matt J. Kusner, Varun Kanade, Ricardo Silva

**Differentiable Causal Backdoor Discovery**

*The International Conference on Artificial Intelligence and Statistics (AISTATS), 2020*

John Bradshaw, Brooks Paige, Matt J. Kusner, Marwin H. S. Segler, Jos Miguel Hernandez-Lobato

**A Model to Search for Synthesizable Molecules**

*Neural Information Processing Systems (NeurIPS), 2019*

Niki Kilbertus, Philip J. Ball, Matt J. Kusner, Adrian Weller, Ricardo Silva

**The Sensitivity of Counterfactual Fairness to Unmeasured Confounding**

*The Conference on Uncertainty in Artificial Intelligence (UAI), 2019*

Matt J. Kusner, Chris Russell, Joshua R. Loftus, Ricardo Silva

**Making Decisions that Reduce Discriminatory Impact**

*The International Conference on Machine Learning (ICML), 2019*

Nitin Agrawal\*, Ali Shahin Shamsabadi\*, Matt J. Kusner, Adrià Gascón

**QUOTIENT: Two-Party Secure Neural Network Training and Prediction**

*The Conference on Computer and Communications Security (CCS), 2019*

John Bradshaw, Matt J. Kusner, Brooks Paige, Marwin H. S. Segler, José Miguel Hernández-Lobato

**A Generative Model For Electron Paths**

*International Conference on Learning Representations (ICLR), 2019*

Amartya Sanyal, Matt J. Kusner, Adrià Gascón, Varun Kanade

**TAPAS: Tricks to Accelerate (encrypted) Prediction As a Service**

*International Conference on Machine Learning (ICML), 2018*

Niki Kilbertus, Adrià Gascón, Matt J. Kusner, Michael Veale, Krishna Gummadi, Adrian Weller

**Blind Justice: Fairness with Encrypted Sensitive Attributes**

*International Conference on Machine Learning (ICML), 2018*

David Janz, Jos van der Westhuizen, Brooks Paige, Matt J. Kusner, José Miguel Hernández-Lobato

**Learning a Generative Model for Validity in Complex Discrete Structures**

*International Conference on Learning Representations (ICLR), 2018*

Chirs Russell\*, Matt J. Kusner\*, Joshua R. Loftus, Ricardo Silva

**When Worlds Collide: Integrating Different Counterfactual Assumptions in Fairness**

*Neural Information Processing Systems (NIPS), 2017*

Matt J. Kusner\*, Joshua R. Loftus\*, Chirs Russell\*, Ricardo Silva  
**Counterfactual Fairness** (oral presentation)

*Neural Information Processing Systems (NIPS), 2017*

Matt J. Kusner, Brooks Paige, José Miguel Hernández-Lobato  
**Grammar Variational Autoencoder**

*International Conference on Machine Learning (ICML), 2017*

Gao Huang, Chuan Guo, Matt J. Kusner, Yu Sun, Kilian Q. Weinberger, Fei Sha  
**Supervised Word Mover's Distance** (oral presentation)

*Neural Information Processing Systems (NIPS), 2016*

Matt J. Kusner, Yu Sun, Karthik Sridharan, Kilian Q. Weinberger  
**Private Causal Inference** (oral presentation)

*Artificial Intelligence and Statistics (AISTATS), 2016*

Gustavo Malkomes, Matt J. Kusner, Wenlin Chen, Kilian Q. Weinberger, Benjamin Moseley  
**Fast Distributed k-Center Clustering with Outliers on Massive Data**

*Neural Information Processing Systems (NIPS), 2015*

Matt J. Kusner, Yu Sun, Nicholas I. Kolkin, Kilian Q. Weinberger  
**From Word Embeddings to Document Distances**

*International Conference on Machine Learning (ICML), 2015*

Matt J. Kusner, Jacob R. Gardner, Roman Garnett, Kilian Q. Weinberger  
**Differentially Private Bayesian Optimization**

*International Conference on Machine Learning (ICML), 2015*

Matt J. Kusner, Wenlin Chen, Quan Zhou, Zhixiang (Eddie) Xu, Kilian Q. Weinberger, Yixin Chen  
**Feature-Cost Sensitive Learning with Submodular Trees of Classifiers**

*AAAI Conference on Artificial Intelligence (AAAI), 2014*

Matt J. Kusner, Stephen Tyree, Kilian Q. Weinberger, Kunal Agrawal  
**Stochastic Neighbor Compression**

*International Conference on Machine Learning (ICML), 2014*

Jacob R. Gardner, Matt J. Kusner, Zhixiang (Eddie) Xu, Kilian Q. Weinberger, John P. Cunningham  
**Bayesian Optimization with Inequality Constraints**

*International Conference on Machine Learning (ICML), 2014*

Zhixiang (Eddie) Xu, Matt J. Kusner, Gao Huang, Kilian Q. Weinberger  
**Anytime Feature Learning**

*International Conference on Machine Learning (ICML), 2013*

Zhixiang (Eddie) Xu, Matt J. Kusner, Kilian Q. Weinberger, Minmin Chen  
**Cost-Sensitive Tree of Classifiers**

*International Conference on Machine Learning (ICML), 2013*

## **Journal Publications**

Mrinal Pahwa, Matt J. Kusner, Carl Hacker, David Bundy, Kilian Q. Weinberger, Eric Leuthardt  
**Optimizing the Detection of Wakeful and Sleep-Like States for Future Electrographic Brain Computer Interface Applications**

*PLOS ONE Journal, 2015*

Zhixiang (Eddie) Xu, Matt J. Kusner, Kilian Q. Weinberger, Minmin Chen, Olivier Chapelle  
**Classifier Cascades and Trees for Minimizing Feature Evaluation Cost**  
*Journal of Machine Learning Research (JMLR)*, 2014

## Preprints

Gbètondji Jean-Sebastien Dovonon, Michael M. Bronstein, Matt J. Kusner  
**Setting the Record Straight on Transformer Oversmoothing**, 2023.

Jean Kaddour, Aengus Lynch, Qi Liu, Matt J. Kusner, Ricardo Silva  
**Causal machine learning: A survey and open problems**, 2023.

Qi Liu, Matt J. Kusner, Phil Blunsom  
**A Survey on Contextual Embeddings**

Valentina Zantedeschi, Fabrizio Falasca, Alyson Douglas, Richard Strange, Matt J. Kusner, Duncan Watson-Parris  
**Cumulo: A Dataset for Learning Cloud Classes**  
NeurIPS Workshop Tackling Climate Change with Machine Learning, 2019. *Best Paper Award*.

Joshua R. Loftus, Chris Russell, Matt J. Kusner, Ricardo Silva  
**Causal Reasoning for Algorithmic Fairness**, May 2018

Jacob R. Gardner\*, Paul Upchurch\*, Matt J. Kusner, Yixuan Li, Kilian Q. Weinberger, Kavita Bala, John E. Hopcroft  
**Deep Manifold Traversal: Changing Labels with Convolutional Features**, March 2016