

CURRICULUM VITAE

ALEXANDER TERENIN

[HTTPS://AVT.IM/](https://avt.im/) · AVT28@CORNELL.EDU

EMPLOYMENT	<i>Assistant Research Professor</i>	2023 – 2026
	Center for Data Science for Enterprise and Society Cornell University	
	<i>Postdoctoral Research Associate</i>	2021 – 2023
	Computational and Biological Learning University of Cambridge	
EDUCATION	<i>PhD Imperial College London</i>	2018 – 2022
	Advised by Profs. Marc Deisenroth and Seth Flaxman	
	<i>MS University of California, Santa Cruz</i>	2014 – 2017
	Advised by Prof. David Draper	
	<i>BS, BA University of California, Santa Barbara</i>	2010 – 2014
	Ranked 1st in my primary department	
PREPRINTS	Q. Xie, L. Cai, A. Terenin , P. I. Frazier, and Z. Scully. Cost-aware Stopping for Bayesian Optimization . arXiv:2507.12453, 2025.	2025
	A. Terenin and J. Negrea. Bayesian Algorithms for Adversarial Online Learning: from Finite to Infinite Action Spaces . arXiv:2502.14790, 2025.	
PUBLICATIONS	P. Mostowsky, V. Dutordoir, I. Azangulov, N. Jaquier, M. J. Hutchinson, A. Ravuri, L. Rozo, A. Terenin , and V. Borovitskiy. The Ge-	2025

ometricKernels Package: Heat and Matérn Kernels for Geometric Learning on Manifolds, Meshes, and Graphs. JMLR 2025.

Z. Scully and [A. Terenin](#). The Gittins Index: A Design Principle for Decision-Making Under Uncertainty. INFORMS Tutorials 2025.

S. Holalkere, D. Bindel, S. Sellán, and [A. Terenin](#). Stochastic Poisson Surface Reconstruction with One Solve using Geometric Gaussian Processes. ICML 2025.

Q. Xie, R. Astudillo, P. I. Frazier, Z. Scully, and [A. Terenin](#). Cost-aware Bayesian Optimization via the Pandora’s Box Gittins Index. NeurIPS 2024. **INFORMS Data Mining Best Paper Competition Finalist.**

I. Azangulov, A. Smolensky, [A. Terenin](#), and V. Borovitskiy. Stationary Kernels and Gaussian Processes on Lie Groups and their Homogeneous Spaces II: non-compact symmetric spaces. JMLR 2024.

I. Azangulov, A. Smolensky, [A. Terenin](#), and V. Borovitskiy. Stationary Kernels and Gaussian Processes on Lie Groups and their Homogeneous Spaces I: the compact case. JMLR 2024.

J. A. Lin, S. Padhy, J. Antorán, [A. Terenin](#), J. M. Hernández-Lobato, and D. Janz. Stochastic Gradient Descent for Gaussian Processes Done Right. ICLR 2024.

L. Cosier, R. Iordan, S. N. T. Zwane, G. Franzese, J. T. Wilson, M. P. Deisenroth, [A. Terenin](#), and Y. Bekiroğlu. A Unifying Variational Framework for Gaussian Process Motion Planning. AISTATS 2024.

T. Brown, S. Marsden, V. Gopakumar, [A. Terenin](#), H. Ge, and F. Casson. Multi-objective Bayesian optimisation for design of Pareto-optimal current drive profiles in STEP. IEEE TPS 2024.

[A. Terenin](#),* D. R. Burt,* A. Artemev,* S. Flaxman, M. van der Wilk, C. E. Rasmussen, and H. Ge. Numerically Stable Sparse Gaussian Processes via Minimum Separation using Cover Trees. JMLR 2024.

*Equal contribution.

J. A. Lin, J. Antorán, S. Padhy, D. Janz, J. M. Hernández-Lobato, 2023
and [A. Terenin](#). [Sampling from Gaussian Process Posteriors using Stochastic Gradient Descent](#). NeurIPS 2023. **Selected for Oral Presentation.**

P. Rosa, V. Borovitskiy, [A. Terenin](#), and J. Rousseau. [Posterior Contraction Rates for Matérn Gaussian Processes on Riemannian Manifolds](#). NeurIPS 2023. **Selected for Spotlight Presentation.**

A. Östling, H. Sargeant, H. Xie, L. Bull, [A. Terenin](#), L. Jonsson, M. Magnusson, and F. Steffek. [The Cambridge Law Corpus: A Dataset for Legal AI Research](#). NeurIPS 2023.

M. J. Hutchinson,* [A. Terenin](#),* V. Borovitskiy,* S. Takao,* Y. W. 2021
Teh, and M. P. Deisenroth. [Vector-valued Gaussian Processes on Riemannian Manifolds via Gauge Independent Projected Kernels](#). NeurIPS 2021.

N. Jaquier, V. Borovitskiy, A. Smolensky, [A. Terenin](#), T. Asfour, and L. Rozo. [Geometry-aware Bayesian Optimization in Robotics using Riemannian Matérn Kernels](#). CoRL 2021.

J. T. Wilson,* V. Borovitskiy,* [A. Terenin](#),* P. Mostowsky,* and M. P. Deisenroth. [Pathwise Conditioning of Gaussian Processes](#). JMLR 2021.

A. Hochlehnert, [A. Terenin](#), S. Sæmundsson, and M. P. Deisenroth. [Learning Contact Dynamics using Physically Structured Neural Networks](#). AISTATS 2021.

V. Borovitskiy,* I. Azangulov,* [A. Terenin](#),* P. Mostowsky, M. P. Deisenroth, and N. Durrande. [Matérn Gaussian Processes on Graphs](#). AISTATS 2021. **Best Student Paper Award.**

S. Cohen, G. Luise, [A. Terenin](#), B. Amos, and M. P. Deisenroth. [Aligning Time Series on Incomparable Spaces](#). AISTATS 2021.

*Equal contribution.

V. Borovitskiy,* [A. Terenin](#),* P. Mostowsky,* and M. P. Deisenroth. 2020
[Matérn Gaussian Processes on Riemannian Manifolds](#). NeurIPS
2020.

J. T. Wilson,* V. Borovitskiy,* [A. Terenin](#),* P. Mostowsky, and M. P. Deisenroth. [Efficiently Sampling Functions from Gaussian Process Posteriors](#). ICML 2020. **Honorable Mention for Outstanding Paper Award**.

S. Sæmundsson, [A. Terenin](#), K. Hofmann, and M. P. Deisenroth. [Variational Integrator Networks for Physically Structured Embeddings](#). AISTATS 2020.

[A. Terenin](#), D. Simpson, and D. Draper. [Asynchronous Gibbs Sampling](#). AISTATS 2020.

[A. Terenin](#), M. Magnusson, and L. Jonsson. [Sparse Parallel Training of Hierarchical Dirichlet Process Topic Models](#). EMNLP 2020.

[A. Terenin](#), S. Dong, and D. Draper. [GPU-accelerated Gibbs Sampling: A Case Study of the Horseshoe Probit Model](#). Stat. and Comp. 2019.

[A. Terenin](#), M. Magnusson, L. Jonsson, and D. Draper. [Pólya Urn Latent Dirichlet Allocation: A Doubly Sparse Massively Parallel Sampler](#). IEEE TPAMI 2019.

SOFTWARE

[GEOMETRICKERNELS](#) 2024
A Python framework for kernels and Gaussian process models on manifolds, graphs, and meshes

THESIS

[A. Terenin](#). [Gaussian Processes and Statistical Decision-making in Non-Euclidean Spaces](#). PhD Thesis, Imperial College London, 2022. Examined by Profs. [Michael Bronstein](#) and [Stefano Ermon](#), passed without corrections.

*Equal contribution.

HONORS AND AWARDS

<i>INFORMS Data Mining Best Paper Competition Finalist</i>	2024
Institute for Operations Research and the Management Sciences Four general-track finalists selected out of all total competition submissions	
<i>Oral Presentation</i>	2023
Conference on Neural Information Processing Systems	
<i>Spotlight Presentation</i>	
Conference on Neural Information Processing Systems	
<i>College Research Associate</i>	2021
Sidney Sussex College, University of Cambridge	
<i>Best Student Paper Award</i>	
Conference on Artificial Intelligence and Statistics Two awards (best paper, best student paper) given out of 1527 submissions and 455 accepted papers	
<i>Honorable Mention for Outstanding Paper Award</i>	2020
International Conference on Machine Learning Two awards and two honorable mentions given out of 4990 submissions and 1086 accepted papers	
<i>Best Presentation Award</i>	
Imperial College London Mathematics PhD Symposium	
<i>Best Poster Award</i>	2018
Imperial College London SIAM Student Conference	
<i>National Science Foundation Fellowship Honorable Mention</i>	2016
<i>Rama Thogarati Memorial Award</i>	2014
Department of Statistics and Applied Probability University of California, Santa Barbara Awarded to highest-ranking graduating student in department	
<i>Exceptional Academic Performance in the Psychology Major</i>	
University of California, Santa Barbara	

Engineering and Sciences Commencement Speaker
University of California, Santa Barbara

GRANTS

Low-Cost Active Climate Monitoring with Underactuated Agents 2025
2030 Project Fast Grant
Co-applicant, with Sarah Dean
Funder: Cornell Atkinson Center for Sustainability
Amount: \$26k

Student Travel Funding Sponsorship for NeurIPS Workshop on 2024
Bayesian Decision-making and Uncertainty
Co-applicant, with Natalie Maus, Renato Berlingheiri, and Zi Wang
Funder: Meta
Amount: \$2k

Research Conference and Workshop Sponsorship for NeurIPS Workshop on
Bayesian Decision-making and Uncertainty
Co-applicant, with Natalie Maus, Renato Berlingheiri, and Zi Wang
Funder: Google
Amount: \$2k

Research Conference and Workshop Sponsorship for NeurIPS Work- 2022
shop on Gaussian Processes, Spatiotemporal Modeling, and Decision-
making Systems
Co-applicant, with Geoff Pleiss, Elizaveta Semanova, and Zi Wang
Funder: Google
Amount: \$6k

MEDIA

[Understanding Uncertainty and the Value of Visualisation in AI](#) 2021
Science News at Imperial College London
Interview about the use of visuals in my research

[Aligning Time Series on Incomparable Spaces](#)
Data Skeptic Podcast
Postcast episode and interview covering one of my papers

STUDENTS

Sihwa Park. PhD, Cornell. Co-supervised with David Bindel	2024 – present
Jingsen Zhu. PhD, Cornell. Co-supervised with Ramin Zabih	2025 – present
Vinamr Jain. PhD, Cornell. Co-supervised with Sarah Dean	2025 – present
Tejal Nair. Undergraduate research, Cornell. Co-supervised with David Bindel Research currently in preparation for submission Continued to a PhD at William and Mary	2023 – 2026
Sidhanth Holalkare. Undergraduate research, Cornell. Co-supervised with David Bindel Research published at ICML Continued to a PhD at Stanford	2023 – 2025
Shashank Kalyanaraman. Undergraduate research, Cornell. Co-supervised with David Bindel	2024 – 2025
Zhiyao Xu. Undergraduate research, Cornell. Co-supervised with David Bindel Continued to a PhD at Purdue	2024 – 2025
Chad Yu. Undergraduate research, Cornell. Co-supervised with David Bindel and Ziv Scully Continued to a quantitative finance firm	2023 – 2025
Theo Brown. <i>Reinforcement Learning and Bayesian optimisation for Tokamak Plasma Control</i> , MEng, University of Cambridge. Co-supervised with Hong Ge Research published at plasma physics journal Continued to a PhD at University College London	2022 – 2023
Lucas Cosier. <i>Variational Gaussian Process Approximation for Motion Planning</i> , BSc, University College London.	2021 – 2022

Co-supervised with Yasemin Bekiroğlu and Marc Deisenroth
Research published at AISTATS
Continued to an MSc at ETH Zürich

Kai Biegun. *Robustness and Online Error Correction in Model-based Reinforcement Learning*, MSc, University College London. 2021
Co-supervised with Marc Deisenroth
Continued to robotics startup, then to a PhD at University College London

Andreas Hochlehnert. *A Contact-aware Symplectic Integrator Network*, MSc, University College London. 2021
Co-supervised with Steindór Sæmundsson and Marc Deisenroth
Research published at AISTATS
Continued to a PhD at University of Tübingen

TEACHING

Guest Lecturer 2025
Cornell University
Format: graduate-level course with approximately 30 students
Courses: Numerical Methods for Data Science

*Teaching Assistant*¹ 2022
University of Cambridge
Format: small-group teaching with up to four students
Courses: Inference, Representation Learning on Graphs and Networks

Teaching Assistant 2016 – 2017
University of California, Santa Cruz
Format: discussion sections with 30-to-60 students
Courses: Classical and Bayesian Inference, Statistical Methods for the Biological, Environmental, and Health Sciences

EMPLOYMENT PRE-ACADEMIA

Research Scientist 2017 – 2018
Petuum (Carnegie Mellon University startup)
Worked on *DyNet*, a pre-TensorFlow deep learning framework, implemented counter-based RNGs using CUDA, published a NeurIPS workshop paper

¹This refers to supervision (and marking) in the University of Cambridge sense.

INTERNSHIPS

<i>Data Science Intern</i> Microsoft Interned with the <i>Bing Maps</i> team, built a user session replay tool	2019
<i>Data Science Intern</i> eBay Interned with the <i>Experimentation Analytics</i> and <i>Search Science</i> teams, built models for predicting A/B test results and search keyword intended category	2015 – 2016
<i>NSF Research Experience for Undergraduates</i> Lewis and Clark College Developed hashing-based pattern extraction techniques for <i>Orego</i> , an MCTS-based computer go algorithm, two years before go was solved by DeepMind	2013

SERVICE

Virtual Seminar Series on Bayesian Decision-making and Uncertainty Co-organized with Natalie Maus, Renato Berlingheiri, Zi Wang	2025
NeurIPS Workshop on Bayesian Decision-making and Uncertainty Co-organized with Natalie Maus, Renato Berlingheiri, Zi Wang	2024
Virtual Seminar Series on Gaussian Processes, Spatiotemporal Modeling, and Decision-making Systems Founding organizer, with Geoff Pleiss, Elizaveta Semanova, and Zi Wang	2022 – 2023
NeurIPS Workshop on Gaussian Processes, Spatiotemporal Modeling, and Decision-making Systems Co-organized with Geoff Pleiss, Elizaveta Semanova, and Zi Wang	2022
NeurIPS Workshop on Interpretable Inductive Biases and Physically Structured Learning Co-organized with Michael Lutter, Shirley Ho, and Lei Wang	2020
Area Chair ICML	2026
NeurIPS, ICML, AISTATS	2025
NeurIPS, AISTATS	2024
AISTATS	2023

Session Chair	
AISTATS	2024
Refereeing	
J. Gl. Opt., NeurIPS Symp. Prob. ML.	2026
ICLR, SIGGRAPH, JASA, CoRL, MCQMC, SIAM-UQ	2025
JMLR×3, SIAM-UQ, NeurIPS D&B, NeurIPS Workshop Prop.	2024
ICML, NeurIPS, CoRL, ICLR, JMLR	2023
JMLR×3, NeurIPS, NeurIPS Workshop Prop., AAAI	2022
AISTATS, ICML, JMLR×2	2021
ICLR, NeurIPS, Deep Learning Indaba	2020
NeurIPS	2019
EURASIP JASP	2018
PLOS One, Stat	2017
PhD External Examiner	
Eugene Golikov · EPFL	2025

TUTORIALS

The Gittins Index: A Design Principle for Decision-Making Under Uncertainty	2025
INFORMS Annual Meeting Co-presented with Ziv Scully	
Geometric Probabilistic Models	2024
Conference on Uncertainty in Artificial Intelligence Co-presented with Viacheslav Borovitskiy	

TALKS

Total of 20+ invited talks out of 70+ total talks given at 40+ institutions and conferences in 10+ countries. Full list available [online](#).

Selected Invited Talks

Dagstuhl Seminar on Bayesian Optimization · Invited Visit	11/05/2025
ISyE-MS&E-IOE Joint Rising Stars Workshop	04/25/2025
SIAM Conf. on Mathematics of Data Science · Invited Talk	10/24/2024
INFORMS · Invited Talk	10/20/2024
University of Waterloo · Invited Talk	10/08/2024
International Society of Bayesian Analysis · Invited Talk	07/06/2024

ETH Zürich · Invited Talk

06/21/2023

Google Brain · Invited Talk · BayesOpt Speaker Series

11/28/2022

PERSONAL

Languages

Python, JAX, CUDA, Rust, Scala, English (native), Russian (native)