

Charles C. Margossian

cmargossian@flatironinstitute.org • charlesm93.github.io • New York, NY

Education

- Ph.D. Statistics**, Columbia University, New York, NY. 2017– 2022
- Thesis: *Modernizing Markov chains Monte Carlo for scientific and Bayesian modeling*
 - Committee: Andrew Gelman (advisor), Aki Vehtari, Matt Hoffman, Sumit Mukherjee and Dave Blei
- B.Sci. Physics**, Yale University, New Haven, CT. 2011–2015
- Baccalauréat Scientifique**, Ecole Jeannine Manuel (High school), Paris, France. 2009–2011

Experience

- Research Fellow**, Flatiron Institute, New York, NY. 2022–
- Independent research in computational mathematics, 3 years funding
 - Computational Statistics and Machine Learning groups
- Core Developer**, Stan Development Team, <https://mc-stan.org/>. 2016–
- Open-source, Bayesian inference and automatic differentiation in C++
 - Main contributions: support for implicit functions (ODEs, algebraic equations...)
- Research Intern**, Google Research, New York, NY. Sum. 2021
- TensorFlow Probability team: inference on modern hardware (GPUs, TPUs)
 - Hosts: Matt Hoffman and Pavel Sountsov
- Visiting Doctoral Student**, Aalto University, Department of Computer Science, Espoo, Finland. Sum. 2019
- Probabilistic machine learning group
 - Advisor: Aki Vehtari
- Visiting Scientist**, Metrum Research Group, Tariffville, CT and Cambridge, MA. 2015–2017
- Co-creator of “Torsten: an extension of Stan for pharmacometrics modeling”
 - 2015–2016: Pharmacometrics bootcamp
 - Supervisor: Bill Gillespie
- Research Assist.**, Yale University, Department of Astronomy, New Haven, CT. 2013–2015
- Exoplanet group
 - Advisors: Debra Fischer and Ji Wang
- Patent Law Intern**, Leinweber & Zimmermann, Munich, Germany. Sum. 2014

Skills

Programming: R, Python, C++, Stan, \LaTeX , GitHub
Languages: English, French, German

Google Scholar: <https://scholar.google.com/citations?user=nPtLsvIAAAAJ&hl=en>

Preprints

- [1] [Charles C. Margossian](#), Loucas Pillaud-Vivien, and Lawrence K. Saul. Variational Inference for Uncertainty Quantification: an Analysis of Trade-Offs. *arXiv:2403.13748*, 2024.
- [2] Diana Cai, Chirag Modi, [Charles C. Margossian](#), Robert M. Gower, David M. Blei, and Lawrence K. Saul. EigenVI: score-based variational inference with orthogonal function expansions. *In preparation*, 2024.
- [3] [Charles C. Margossian](#). General adjoint-differentiated Laplace approximation. *arXiv:2306.14976*, 2023.
- [4] [Charles C. Margossian](#) and Michael Betancourt. Efficient Automatic Differentiation of Implicit Functions. *arXiv:2112.14217*, 2022.
- [5] [Charles C. Margossian](#) and Sumit Mukherjee. Simulating Ising and Potts models at critical and cold temperatures using auxiliary Gaussian variables. *arXiv:2110.10801*, 2021.
- [6] Andrew Gelman, Aki Vehtari, Daniel Simpson, [Charles C. Margossian](#), Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Bürkner, and Martin Modrák. Bayesian Workflow. *arXiv:2011.01808*, 2020.
- [7] Michael Betancourt, [Charles C. Margossian](#), and Vianey Leos-Barajas. The Discrete Adjoint Method: Efficient Derivatives for Functions of Discrete Sequences. *arXiv:2002.00326*, 2020.

Publications

*supervised student

- [1] [Charles C. Margossian](#), Matthew D. Hoffman, Pavel Sountsov, Lionel Riou-Durand, Aki Vehtari, and Andrew Gelman. Nested \hat{R} : Assessing the convergence of Markov chain Monte Carlo when running many short chains. *Bayesian Analysis (accepted)*, 2024.
- [2] [Charles C. Margossian](#) and Andrew Gelman. For how many iterations should we run Markov chain Monte Carlo? In *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC, (upcoming) 2nd edition, 2024.
- [3] [Charles C. Margossian](#) and David M. Blei. Amortized Variational Inference: When and Why? *Uncertainty in Artificial Intelligence*, PMLR (accepted), 2024.
- [4] David Heurtel-Depeiges, [Charles C. Margossian](#), Ruben Ohana, and Bruno Régaldo-Saint Blancard. Listening to the noise: Blind Denoising with

- Gibbs Diffusion. *International Conference on Machine Learning*, PMLR (accepted), 2024.
- [5] Diana Cai, Chirag Modi, Loucas Pillaud-Vivien, Charles C. Margossian, Robert M. Gower, David M. Blei, and Lawrence K. Saul. Batch and match: black-box variational inference with a score-based divergence. *Selected for spotlight*, *International Conference on Machine Learning*, PMLR (accepted), 2024.
- [6] Charles C. Margossian and Lawrence K. Saul. The Shrinkage-Delinkage Trade-off: An Analysis of Factorized Gaussian Approximations for Variational Inference. *Selected for oral presentation*, *Uncertainty in Artificial Intelligence*, PMLR 216:1358–1367, 2023.
- [7] Lionel Riou-Durand, Pavel Sountsov, Jure Vogrinc, Charles C. Margossian, and Sam Power. Adaptive Tuning for Metropolis Adjusted Langevin Trajectories. *Artificial Intelligence and Statistics*, PMLR 206:8102–8116, 2023.
- [8] Chirag Modi, Charles C. Margossian, Yuling Yao, Robert M. Gower, David M. Blei, and Lawrence K. Saul. Variational Inference with Gaussian Score Matching. *Advances in Neural Information Processing Systems*, 37, 2023.
- [9] Charles C. Margossian, Yi Zhang, and William R. Gillespie. Flexible and efficient Bayesian pharmacometrics modeling using Stan and Torsten, Part I. *CPT: Pharmacometrics & Systems Pharmacology*, 11(9):1151–1169, 2022.
- [10] Philip Greengard, Jeremy Hoskins, Charles C. Margossian, Jonah Gabry, Andrew Gelman, and Aki Vehtari. Fast methods for posterior inference of two-group normal-normal models. *Bayesian Analysis*, 18(3):889–907, 2022.
- [11] Léo Grinsztajn*, Elizaveta Semenova, Charles C. Margossian, and Julien Riou. Bayesian workflow for disease transmission modeling in Stan. *Statistics in Medicine*, 40(27):6209–6234, 2021.
- [12] Charles C. Margossian, Aki Vehtari, Daniel Simpson, and Raj Agrawal. Hamiltonian Monte Carlo using an adjoint-differentiated Laplace approximation: Bayesian inference for latent Gaussian models and beyond. *Advances in Neural Information Processing Systems*, 34, 2020.
- [13] Anthony Hauser, Michel J. Counotte, Charles C. Margossian, Garyfallos Konstantinoudis, Nicola Low, Christian L. Althaus, and Julien Riou. Estimation of SARS-CoV-2 mortality during the early stages of an epidemic: a modeling study in Hubei, China and six regions in Europe. *PLOS Medicine*, 17(7), 2020.
- [14] Charles C. Margossian. A Review of automatic differentiation and its efficient implementation. *Awarded "Top WIREs articles in 2022"*, *WIREs: Data Mining and Knowledge Discovery*, 9(4), 2019.

- [15] Joseph R. Schmitt, Eric Agol, Katherine M. Deck, Leslie A. Rogers, Zachary J. Gazak, Debra A. Fischer, Ji Wang, Matthew J. Holman, Kian J. Jek, Charles Margossian, Mark R. Omohundor, Troy Winarski, John M. Brewer, Matthew J. Giguere, Chris Lintott, Stuart Lynn, Michael Parrish, Kevin Schawinski, Megan E. Schwamb, Robert Simpson, and Arfon M. Smith. Planet Hunters. VII. Discovery of a new low-mass, low-density planet (PH3 C) orbiting KEPLER-289 with mass measurements of two additional Planets (PH3 B and D). *Astrophysical Journal*, 795(2), 2014.

Conference posters and notebooks

*supervised student

- [1] Emmanuel Mokol* and Charles C. Margossian. Monitoring Nonstationary Variance to Assess Convergence of MCMC. *Best poster award*, *International Society of Bayesian Analysis (ISBA): World Meeting*, 2024.
- [2] Diana Cai, Chirag Modi, Charles C. Margossian, Robert M. Gower, David M. Blei, and Lawrence K. Saul. EigenVI: score-based variational inference with orthogonal function expansions. *International Conference on Machine Learning (ICML): Workshop on Structured Probabilistic Inference and Generative Modeling*, 2024.
- [3] Charles C. Margossian, Matthew D. Hoffman, Pavel Sountsov, Lionel Riou-Durand, Aki Vehtari, and Andrew Gelman. Assessing the convergence of Markov chains Monte Carlo when running many chains. *BayesComp*, 2023.
- [4] Stanislas du Ché* and Charles C. Margossian. Parallelization for Markov chains Monte Carlo with heterogeneous runtimes. *BayesComp*, 2023.
- [5] Charles C. Margossian, Lu Zhang, Sebastian Weber, and Andrew Gelman. Solving ODEs in a Bayesian context: challenges and opportunities. *Population Approach Group in Europe*, 2021.
- [6] Aurélien Marc, Marion Keriou, Charles Margossian, Julie Bertrand, Pauline Maisonnasse, Yoan Aldon, Rogier W. Sanders, Marit Van Gils, Roger Le Grand, and Jérémie Guedj. Developing a model of SARS-CoV-2 viral dynamics under monoclonal antibody treatment. *Population Approach Group in Europe*, 2021.
- [7] Johann D. Gaebler* and Charles C. Margossian. Propagating Derivatives through Implicit Functions in Reverse Mode Autodiff. *Stanford Institute for Computational & Mathematical Engineering*, 2021.
- [8] Charles C. Margossian, Aki Vehtari, Daniel Simpson, and Raj Agrawal. Approximate Bayesian inference for latent Gaussian models in Stan. *StanCon*, 2020.

- [9] Charles C. Margossian and Andrew Gelman. Bayesian model of planetary motion: exploring ideas for a modeling workflow when dealing with ordinary differential equations and multimodality. In *Stan Case Studies*, volume 7, 2020.
- [10] Charles C. Margossian. Computing Steady States with Stan's Nonlinear Algebraic Solver. *StanCon*, 2018.
- [11] Charles C. Margossian and William R. Gillespie. Gaining Efficiency by Combining Analytical and Numerical Methods to Solve ODEs: Implementation in Stan and Application to Bayesian PK/PD. *American Conference on Pharmacometrics*, 2017.
- [12] Charles C. Margossian and William R. Gillespie. Differential Equation Based Models in Stan. *StanCon*, 2017.
- [13] Charles C. Margossian and William R. Gillespie. Stan Functions for Pharmacometrics Modeling. *American Conference on Pharmacometrics*, 2016.

Software

Core developer, Stan: a probabilistic programming language, mc-stan.org.

Co-creator, Torsten: an extension of Stan for pharmacometrics modeling, [GitHub](#).

Contributor, mrgSolve: Simulation from ODE-Based Population PK/PD and System Pharmacology Models, [GitHub](#).

Contributor, bayesplot: Plotting for Bayesian Models in R, [Cran](#).

Awards and recognitions

ISBA best poster award, For *Monitoring Nonstationary Variance to Assess Convergence of MCMC*, presented at the International Society of Bayesian Analysis (ISBA): World Meeting. 2024

ICML spotlight, For *Batch and Match: black box variational inference with a score-based divergence*, top ~13% articles accepted at the International Conference on Machine Learning. 2024

UAI oral presentation, For *The Shrinkage-Delinkage Trade-off: An Analysis of Factorized Gaussian Approximations for Variational Inference*, top ~15% articles accepted at the conference on Uncertainty in Artificial Intelligence. 2023

WIRES top article, For *A Review of Automatic Differentiation and its Efficient Implementation*, which was amongst the top 10 most cited articles in the [2021 Journal Citation Report](#) for *WIRES: Data Mining and Knowledge Discovery*. 2022

Minghui Yu teaching assistant award, Department of Statistics, Columbia University, Awarded by the Director of Graduated Studies based on student feedback. 2022

AISTATS top reviewer, The top reviewers were selected based on the feedback received from the Area Chairs and comprise the top-10% of AISTATS reviewers. 2022

- Dean's fellowship**, Department of Statistics, Columbia University, 5 years funding for PhD degree. 2017
- Yale book award**, For "character and intellectual promise". 2010

Academic services

- Reviewer**, SIAM review (2024), Advances in Neural Information Processing Systems (2024, 2020), Statistics and Computing (2024), CPT: Pharmacometrics and Systems Pharmacology (2024), PeerJ (2023), Journal of Machine Learning Research (2023 (2)), PMLR: Artificial Intelligence and Statistics (2023, **Top Reviewer Award** 2021), Journal of Pharmacokinetics and Pharmacodynamics (2023, 2019), Computational Statistics (2022), Nature Geoscience (2021), Methods in ecology (2021), Journal of data science (2021). 2020–
- Elected Member**, Stan Governing Body, Two year term. 2022–
- Organizer**, StanCon 2024, Oxford, UK. 2024
- Organizer**, StanCon 2023, St Louis, MO. 2023
- Student representative**, PhD program in statistics at Columbia University. 2019 – 2020

Supervised students

- Manny Mokel (undergraduate student)**, Monitoring Nonstationary Variance to Assess Convergence of MCMC, Flatiron Institute, New York, NY. 2023
- Stanislas Du Ché (master student)**, Parallelization for Markov chain Monte Carlo with heterogeneous runtimes, Columbia University, New York, NY. 2022
- Johann Gaebker (PhD student)**, Propagating Derivatives through implicit functions in reverse mode automatic differentiation, Columbia University, New York, NY. 2020
- Léo Grinsztajn (master student)**, Bayesian Workflow for disease transmission models, Columbia University, New York, NY. 2020
- Hyunji (Angie) Moon (undergraduate student)**, Simulation-based Calibration for the embedded Laplace approximation, Columbia University, New York, NY. 2020

Teaching

- Instructor**, Monte Carlo Methods, Half-a-day course, Nordic Summer School on Probabilistic AI, Copenhagen, Denmark. 2024
- Instructor**, Bayesian Workflow for hierarchical and ODE-based models, Three day workshop, Summer School on Advanced Bayesian Methods, Leuven, Belgium. 2023
- Instructor**, Fundamentals of Stan, Half-day workshop, StanCon 2023, Washington University in St Louis, St Louis, MO. 2023
- Instructor**, Building, fitting, and criticizing Bayesian PK/PD models, One day workshop, University of Buffalo, Buffalo, NY. 2019–2023

Lecturer, Probability and Bayes, Lecture for PHC 506: Biometry in Pharmaceutics, University of Buffalo, Buffalo, NY. 2019–2023

Instructor, Stan for the people: an introductory workshop to Bayesian modeling, Two day workshop, McGill University, Montreal, Canada. 2019, 2020

Instructor, Population and ODE-based models using Stan and Torsten, Two day workshop, StanCon 2019, Cambridge University, Cambridge, UK. 2019

Guest Lecturer, Introduction to Bayesian Data Analysis with Stan, Lecture for STAT 2020: Bayesian Statistics, Harvard University, Cambridge, MA. 2017

Teaching Assistant, **Recipient of the Minghui Yu Teaching Assistant Award**, Courses at all levels (undergrad, masters and PhD), Columbia University, New York, NY. 2017–2022

Peer Tutor, Science, Technology, and Research Scholars (STARS) program, Yale University, New Haven, CT. 2013–2015

Presentations (selected)

Session chair, Monte Carlo methods using modern hardware, International Society of Bayesian Analysis (ISBA): World Meeting, Venice, Italy. 2024

Invited talk, Variational Inference for Uncertainty Quantification: An Analysis of Trade-offs, International Society of Bayesian Analysis (ISBA): World Meeting, Venice, Italy. 2024

Invited talk, The Wisdom of Automatic Differentiation, Applied and Computational Math Group Meeting, Courant Institute, New York University, New York, NY. 2023

Oral Presentation, The Shrinkage-Delinkage Trade-off: An Analysis of Factorized Gaussian Approximations for Variational Inference, Conference on Uncertainty in Artificial Intelligence, Pittsburgh, PA. 2023

Lecture, Solving ODEs in a Bayesian model, Flatiron-Wide Algorithms and Mathematics ($F_w(\alpha + m)!$), Flatiron Institute, New York, NY. 2022

Invited talk, Nested \hat{R} : Assessing convergence for Markov chains Monte Carlo when running many short chains, Center for Research in Economics and Statistics (CREST), École Polytechnique, Paris, France. 2022

Talk, Bayesian inference for latent Gaussian models: MCMC, approximate methods, and hybrids, Minghui Yu memorial conference, Columbia University, New York, NY. 2021

Invited talk, Developing a Bayesian modeling workflow for population PBPK, American Conference on Pharmacometrics. 2020

Invited talk, Developing a Bayesian workflow to model the Covid-19 outbreak, 12th Covid-19 symposium, Columbia University, New York, NY. 2020

Proposed talk, Computing steady states with Stan's nonlinear algebraic solver, StanCon, Pacific Grove, CA. 2018

Proposed talk, Differential equations based models in Stan, StanCon, Columbia University, New York, NY. 2017

Invited talk, Differential equations based models in Stan, Stan Meetup in Boston, Harvard University, Cambridge, MA. 2016

Lecture, Practice (and malpractices!) of Bayesian analysis, Metrum Journal Minute, Tariffville, CT. 2016

Talk, How stars and planets Interact: testing the effects of close-in giant planets on stellar magnetic activity, Davenport Mellon Forum, Yale University, New Haven, CT. 2015

Modified July 2024