Georgios Kissas

Zürich, Switzerland · +41 78 969 95 77 · georgios.kissas@sdsc.ethz.ch
linkedin.com/in/georgioskissas · Google Scholar · georgioskissas.com

Summary

Machine learning researcher specializing in large-scale deep learning for scientific simulation. Hands-on experience pre-training billion-parameter transformers on multi-node GPU clusters, building generative models and tokenizers for large scale physical systems, and engineering end-to-end training pipelines over 14 TB of simulation data. 17 publications and preprints across NeurIPS, ICML, and JMLR; one US patent; six supervised MSc and PhD student projects, four published to date.

Professional Experience

Senior Data Scientist, Research – *Swiss Data Science Center, ETH Zürich* *Nov. 2025 – Present*

- Lead deep-learning research on TGSF-DA: large-scale AI for exascale cosmological astrophysics simulations with multi-million-particle systems.
- Develop generative modeling and discrete tokenization methods for large-scale particle simulations.
- Build multi-GPU, multi-node training pipelines on the Alps supercomputer over a 14 TB HDF5 training corpus, with high-throughput CUDA dataloaders and containerized Docker/enroot deployment.
- Deep learning surrogates for CFD applications in industrial design benchmarks (ShapeNet-Car, AirfRANS, Aircraft).
- Verification-first agentic research for math/code verification and math/code generation.

Postdoctoral Researcher & Lecturer – *ETH Zürich* *Jun. 2023 – Nov. 2025*

- ETH AI Center Postdoctoral Fellow, 1.5% acceptance rate; joint appointment across Applied Mathematics, Civil and Mechanical Engineering, and Computer Science.
- Pre-trained a 3B-parameter Llama-architecture and a 2.4B-parameter GPT-3-architecture language model with Megatron-LM on multi-node GPU (Blackwell) clusters on Euler, benchmarking attention mechanisms and training stabilizers for structurally stable transformers in Semigroup Attention, under review at NeurIPS 2026.
- Led a research line on neuro-symbolic scientific discovery: grammar-based discovery of differential equations and analytical solution synthesis, ICML 2026 and MSSP 2025, and data-driven constitutive model discovery, CMAME 2024.
- Developed unified simulation-estimation and discrete tokenization frameworks for PDEs: FUSE, NeurIPS 2024; Phaedra, arXiv 2026.
- Supervised six MSc and PhD student projects, four already resulting in publications, and lectured graduate-level machine learning courses.

Research Assistant – *University of Pennsylvania, Philadelphia* *Aug. 2018 – May 2023*

- Physics-informed machine learning for cardiovascular flow modeling; lead author of a CMAME paper listed among the journal's most-cited articles of 2019–2023.
- Co-developed attention-based neural operators, LOCA: JMLR 2022 and NeurIPS 2022 Spotlight, and nonlinear manifold decoders, NOMAD: NeurIPS 2022; co-inventor on the associated US patent.
- Bayesian inference, uncertainty quantification, and sensitivity analysis for time series and hemodynamics.

Education

PhD, Mechanical Engineering and Applied Mechanics, University of Pennsylvania *2018 – 2023*
Thesis: *Towards Digital Twins for Cardiovascular Flows: A Hybrid Machine Learning and CFD Approach.*
Advisor: Paris Perdikaris.

MSc, Computational Mechanics (Fluids), National Technical University of Athens (2nd-highest GPA) *2017*
Thesis: *A Computational Study of Self-Consistent Field Theory for Polymer Interfaces.*
Advisor: Doros Theodorou.

Diploma, Mechanical and Aeronautical Engineering, University of Patras (Honors) *2015*

Technical Skills

ML frameworks: PyTorch, JAX.
Large-scale computing: multi-GPU/multi-node distributed training, CUDA dataloaders, SLURM, Docker.
Methods: transformer pre-training and attention design, generative modeling, discrete tokenization, neural operators, neuro-symbolic methods, Bayesian inference and UQ, physics-based simulation.

Selected Publications

*Selected publications and preprints; full list on Google Scholar; * denotes equal contribution.*

Kissas, G., Bonev, B., Oikonomou, O., Lingsch, L., Chamon, L. F. O. Attention, Generators, and Semigroups: Structurally Stable Deep Transformers. *Under review, NeurIPS 2026.*

Koren, N., Hofmann, T., **Kissas, G.** Advectra: Asymmetric Latent Transport for Non-Stationary Physics. *Under review, NeurIPS 2026.*

Lingsch, L., **Kissas, G.**, Jakubik, J., Mishra, S., 2026. Phaedra: Learning High-Fidelity Discrete Tokenization for the Physical Sciences. *arXiv:2602.03915.*

Oikonomou, O., Lingsch, L., Grund, D., Mishra, S., **Kissas, G.**, 2026. Neuro-Symbolic AI for Analytical Solutions of Differential Equations. *International Conference on Machine Learning (ICML).*

Lingsch, L. E., Grund, D., Mishra, S., **Kissas, G.**, 2024. FUSE: Fast Unified Simulation and Estimation for PDEs. *Neural Information Processing Systems (NeurIPS).*

Kissas, G., Mishra, S., Chatzi, E., De Lorenzis, L., 2024. The Language of Hyperelastic Materials. *Computer Methods in Applied Mechanics and Engineering*, 428, p. 117053.

Seidman, J. H., **Kissas, G.**, Pappas, G. J., Perdikaris, P., 2023. Variational Autoencoding Neural Operators. *International Conference on Machine Learning (ICML).*

Kissas, G.*, Seidman*, J. H., Guilhoto, L. F., Preciado, V. M., Pappas, G. J., Perdikaris, P., 2022. Learning Operators with Coupled Attention. *JMLR*; also presented at *NeurIPS 2022, Journal-to-Conference Track* (**Spotlight**).

Seidman*, J. H., **Kissas, G.***, Perdikaris, P., Pappas, G. J., 2022. NOMAD: Nonlinear Manifold Decoders for Operator Learning. *Neural Information Processing Systems (NeurIPS).*

Kissas, G., Yang, Y., Hwuang, E., Witschey, W. R., Detre, J. A., Perdikaris, P., 2020. Machine Learning in Cardiovascular Flows Modeling: Predicting Arterial Blood Pressure from Non-Invasive 4D Flow MRI Data Using Physics-Informed Neural Networks. *Computer Methods in Applied Mechanics and Engineering*, 358, p. 112623. **Among the journal's most-cited articles, 2019–2023.**

Patent

Computer Systems and Methods for Learning Operators. P. Perdikaris, G. J. Pappas, J. Seidman, **G. Kissas**. US Patent 2023/0214661 (link).

Awards & Fellowships

Gold Reviewer, ICML (2026) · **EUROMECH Colloquium 662** Early Career Award (2026) · **Top Reviewer**, NeurIPS (2024) · **Asuera Stiftung Scholarship**, ETH Zürich Foundation, CHF 120,000 (2023) · **ETH AI Center Postdoctoral Fellowship**, 1.5% acceptance rate (2023) · **Doctoral Fellowship Award**, University of Pennsylvania (2018) · **Limmat Stiftung Award**, National Technical University of Athens (2017).

Student Supervision

Six supervised student theses and projects; four have led to peer-reviewed publications or preprints, with manuscripts from the two 2026 projects in preparation.

- **Luca Mascherpa** (MSc, Politecnico di Milano, 2026) — addressing misspecification in PDE surrogate models; with S. Mishra, A. Manzoni → manuscript in preparation.
- **Simone Brivio** (PhD, Politecnico di Milano, 2026) — mathematical analysis of the Perceiver architecture; with S. Mishra, A. Manzoni → manuscript in preparation.
- **Karin Yu** (MSc, ETH Zürich, 2025) — formal language for ODE discovery; with E. Chatzi → MSSP 2025; ICML 2026 Workshop.
- **Levi Lingsch** (MSc, ETH Zürich, 2024) — flow matching neural operators for forward and inverse problems; with S. Mishra; grade 6.0/6.0 → FUSE, NeurIPS 2024.
- **Orestis Oikonomou** (MSc, UZH, 2024) — formal grammar foundation models for PDE discovery and solution; with S. Mishra, D. Scaramuzza; grade 6.0/6.0 → ICML 2026.
- **Diego Renner** (MSc, ETH Zürich, 2023) — differentiable simulations of haemodynamic systems; with S. Mishra, B. Moseley; grade 6.0/6.0 → arXiv:2412.14572 (2024).

Community Service

Reviewer for NeurIPS, ICML, ICLR, *Journal of Machine Learning Research*, *Computer Methods in Applied Mechanics and Engineering*, *Journal of Scientific Computing*, and *Nature Communications*.