Hardik Tankaria, PhD

Kyoto, Japan
(+81) 80-2123-5511 • hardiktankaria1406@gmail.com • hardy-opt.github.io
hardy-opt

Summary

Optimization researcher and ML research engineer. PhD in Informatics (Mathematical Optimization) from Kyoto University, with published work on second-order and variance-reduced methods that make curvature information affordable at high dimension. I take algorithms from convergence proof to benchmarked implementation. Currently project owner for three concurrent R&D lines at MI-6: budget-aware Bayesian optimization for self-driving laboratories, ontology-driven descriptor selection shipped as a versioned software component, and cost-to-go search heuristics for retrosynthesis planning. **Target roles: AI/ML Research Engineer, ML/AI Research Scientist, Optimization Scientist, Algorithm Development, AI for Science Researcher.**

Experience

MI-6 Co., Ltd. **Tokyo, Japan**
Machine Learning Research Engineer *Nov 2025 - Present*

- **Project owner for three concurrent R&D lines** in optimization and AI for science, each taken from problem framing through to a measured, reproducible result.
- **Budget-aware Bayesian optimization for self-driving laboratories.** Designed a closed-loop experiment-selection system in which every proposal is priced by a cost model fitted to real instrument runs, charged against a ledger, and refused if the remaining budget cannot cover it. Because the cost function is a fitted closed form, the affordable set is known exactly before anything is spent, so the budget constraint is applied inside the acquisition optimization rather than checked afterwards. Also handles a per-experiment slot cap. Multiple budget-aware and budget-blind policies, interchangeable acquisition optimizers, and a one-command check that re-derives every committed figure from committed data.
- **Ontology-driven descriptor selection.** Shipped an installable, versioned software component that answers which molecular descriptors describe a chemical property, with confidence and provenance, and publishes it as a data contract consumed by a separate Gaussian-process optimization engine. Determinism enforced as a test, module boundary checked at process level, and continuous integration gating its own published result counts.
- **Retrosynthesis search and benchmark construction.** Built the evaluation infrastructure for double-ended synthesis planning (DESP), including large-scale route dataset generation with structural, template-replay and stock-closure validation, then used it to diagnose where the planner loses. Acted on the diagnosis by reusing the planner's learned cost-to-go estimate to reorder candidate reactions by remaining distance to available starting materials, improving solve rate at unchanged per-step compute. Validated on a held-out external benchmark verified disjoint from training. **First-author conference abstract.**
- **Agentic LLM framework.** Architecting an agentic framework over the descriptor ontology that makes cost-aware experimental decisions, selecting adaptively between a cheap observation and an expensive full characterisation inside a multi-fidelity loop.
- **Methodology across all lines:** pre-registered experiment designs, paired statistical tests with bootstrap confidence intervals, a stated detection floor beside every comparison, and negative results reported with their mechanism rather than dropped.

MI-6 Co., Ltd. **Tokyo, Japan**
Machine Learning Research Intern *Aug 2025 - Nov 2025*

- Built a neural-guided bidirectional search method for double-ended synthesis planning, extending a published planner to handle convergent routes via a multi-tree forest and a proof-tree-complete meeting criterion.
- Implemented Monte Carlo Tree Search with UCT/PUCT variants and graph-based heuristics for retrosynthetic pathway search; benchmarked neural-guided search against established baselines over public chemical databases.

Indian Institute of Technology Mandi **Himachal Pradesh, India**
Machine Learning Research Associate *Jan 2025 - Mar 2025*

- Developed a Nyström-based ADMM algorithm for federated learning, targeting the communication and memory cost of second-order information in distributed training.
- Built a Python framework for decentralised training across multiple nodes, with convergence monitoring and benchmarking. **Co-PI**; guided two M.Tech researchers.

Kyoto University, Dept. of Intelligence Science & Technology **Kyoto, Japan**
Research Assistant *Dec 2021 - Mar 2023*

- Developed the regularized Nyström method for large-scale optimization, reducing per-iteration cost from $O(nd^2 + d^3)$ to $O(ndm + dm^2)$ with sketch size m of roughly 5–20% of the feature dimension, and established global convergence under standard assumptions.
- Benchmarked against randomized subspace Newton and Newton-sketch on high-dimensional classification; applied curvature-aware optimization to brain-MRI tumour classification with transfer learning.

Hayakawa Co. Ltd. **Kyoto, Japan**
E-Commerce & Marketing Research Analyst (part-time) *May 2021 - Oct 2024*

- Held alongside doctoral research. Managed B2B client relationships and negotiations for international e-commerce; competitive analysis for listings and pricing; automated sales and inventory reporting.

Selected Publications

1. **A Regularized Limited Memory BFGS Method for Large-Scale Unconstrained Optimization and its Efficient Implementations.** H. Tankaria, S. Sugimoto, N. Yamashita. *Computational Optimization and Applications* **82**(1), 61–88, 2022. doi:10.1007/s10589-022-00351-5
2. **A Stochastic Variance Reduced Gradient using Barzilai–Borwein Techniques as Second Order Information.** H. Tankaria, N. Yamashita. *Journal of Industrial and Management Optimization* **20**(2), 525–547, 2024. doi:10.3934/jimo.2023089
3. **NysReg-Gradient: Regularized Nyström-Gradient for Large-Scale Unconstrained Optimization and its Application.** H. Tankaria, M. Yamada, D. Singh. *IEEE IJCNN 2026 (WCCI)*, Maastricht.
4. **Nys-Fed: Communication and Memory Efficient Nyström Approximated Curvature for Federated Learning.** S. Ghosh, H. Tankaria, A. Bhavsar, V. Dutt, D. Singh. *IEEE IJCNN 2026 (WCCI)*, Maastricht.
5. **Nys-Newton: Nyström-Approximated Curvature for Convex Optimization.** H. Tankaria, D. Singh, M. Yamada. *IEICE Technical Report (IBISML)* **122**(325), 134–141, 2022.

Metrics: 42 citations, h-index 3 (Google Scholar, Oct 2026). Full list: hardy-opt.github.io/publications

Education

Kyoto University, Graduate School of Informatics — **Kyoto, Japan**
PhD in Informatics (Mathematical Optimization) — *Apr 2018 - Mar 2024*
Dept. of Applied Mathematics & Physics. Research areas: Mathematical Optimization; Optimization for Machine Learning; Second-Order & Stochastic Methods; Randomized Numerical Linear Algebra. Thesis: *The Utilization of Second-Order Information for Large-Scale Unconstrained Optimization Problems*. Funded by the Japan International Cooperation Agency (JICA). GPA 3.7/4.0.

Indian Institute of Technology Hyderabad — **Hyderabad, India**
MSc, Mathematics — *Aug 2014 - May 2016*
Thesis: *Hasse–Minkowski Principle for Quadratic Forms over Q*. Admitted via the Joint Admission Test (JAM), all-India rank 343.

D.K.V Arts & Science Collegs, Saurashtra University — **Jamnagar, India**
BSc, Mathematics — *Jul 2011 - May 2014*

Technical Skills

Programming: Python (expert), MATLAB, C++ (basic), SQL, LaTeX

Optimization: Convex & non-convex optimization; second-order and quasi-Newton methods (L-BFGS, Newton-sketch, Nyström); stochastic and variance-reduced methods (SGD, SVRG, Adam); randomized numerical linear algebra; convergence and complexity analysis; multi-objective and hyperparameter optimization

Supervised learning: Classification and regression: logistic regression, linear and L2-regularized models, support vector machines (hinge and squared hinge), random forests, gradient boosting (XGBoost, LightGBM), stacked ensembles. Feature engineering, class-imbalance handling (SMOTE, class weighting), cost-sensitive objectives, model selection, cross-validation, ROC-AUC, precision/recall, F1, confusion matrices, SHAP interpretability

Deep learning: Convolutional networks (ResNet), recurrent networks (LSTM), transfer learning and fine-tuning, CNN encoders for surrogate models, curvature-aware training of deep nets, medical image classification

Bayesian & probabilistic: Bayesian optimization, Gaussian processes, acquisition function design, cost-aware and multi-fidelity acquisition, sequential decision-making under budget

Federated & distributed: Federated learning across edge clients (benchmarked against FedAvg, FedProx, FedAdam, FedNew, FedNS, DONE); ADMM and one-pass ADMM; communication- and memory-efficient second-order updates; quantization of client updates; decentralised training

Search & planning: Shortest-path and best-first search with learned heuristics (A*, bidirectional A* with meeting-point selection); Monte Carlo Tree Search with UCT and PUCT; AND/OR graph search; cost-to-go value functions; branch-selection heuristics to reduce expansions; multi-objective route scoring

AI systems & agentic: LLM agents with knowledge grounding and tool use; sub-agent architectures communicating through versioned data contracts; cost-aware and multi-fidelity decision agents; evaluation of LLM components with pre-registered predictions, positive controls and instrument-validity gates

AI-assisted engineering: Claude Code and Codex in the development and review loop, including independent re-derivation of published results by a separate agent; AI-assisted refactoring, test generation and documentation

Libraries: PyTorch, NumPy, SciPy, scikit-learn, pandas, BoTorch, GPyTorch, Optuna, Gurobi, RDKit, MONAI, NiBabel

Research engineering: Reproducible research (determinism tests, CI-gated results), architecture decision records, experimental design and pre-registration, paired statistical testing (McNemar, bootstrap CI), detection floors, benchmark and dataset generation, model optimization

Platforms & tools: Linux, macOS; Git and GitHub Actions; pytest; Python packaging; conda; Google Colab; GCP; MS Office including Excel (VBA, pivot tables)

Spoken languages: Gujarati (native), Hindi (native), English (business), Japanese (elementary)

Funding, Awards & Service

2019–2020: **Co-investigator, KIOXIA Encouragement Research Grant** (JPY 1,000,000). *RL-BFGS for Big Data with Improved Global Convergence.* Named co-investigator as a second-year PhD student; responsible for numerical simulations and convergence analysis of the proposed method.

2017: **JICA Scholarship** for doctoral study at Kyoto University (full funding)

2017: **GSET**, qualified for Assistant Professor eligibility, Gujarat state universities, India

2015 / 2014: Merit-cum-Means Scholarship, IIT Hyderabad / JAM all-India rank 343

Service:

- **Vice President** of Japan Japan chapter of Indian Institute of Technology Hyderabad (IIT Hyderabad)
- Reviewer, ICML, IJCNN, ICASSP, ICDCIT.
- **President, Indian Association of Kyoto (2022–2024)**; Treasurer (2019–2022), led a volunteer team organising events for 200+ participants, hosting the Mayor of Kyoto and members of the Japanese House of Representatives.

Selected Talks

- **ICIAM 2023** (Waseda, Tokyo)
- **ICCOPT 2019** (Berlin)
- **OIST International Symposium on AI & Brain Science 2022** (Okinawa)
- **IBISML 2022**
- **JSIAM 2024** (Niigata)
- **Invited talk, IIT Mandi 2025**