

DRAVYANSH SHARMA

dravy@ttic.edu ◊ <https://dravy.ttic.edu/>

Toyota Technological Institute at Chicago, 6045 S Kenwood Ave, Chicago, IL 60637, USA

EDUCATION

Carnegie Mellon University, Pittsburgh

2019 - 2024

PhD, Computer Science Department; Prof. Maria-Florina (Nina) Balcan

Thesis: [Data-driven algorithm design and principled hyperparameter tuning in machine learning](#)

Indian Institute of Technology Delhi, New Delhi

2011 - 2015

Bachelor of Technology in Computer Science and Engineering; GPA: 9.8/10 (first in class)

EMPLOYMENT

Postdoctoral Researcher, The Institute for Data, Econometrics, Algorithms, and Learning (IDEAL)

Toyota Technological Institute at Chicago, Northwestern University

2024-

- ◊ Research on machine learning theory and artificial intelligence.
- ◊ Taught a new course TTIC 31290, “Machine Learning for Algorithm Design”.

Student Researcher, Online learning with Bandit Feedback

Google Research, Bangalore

Summer 2022

- ◊ Algorithms for hyperparameter tuning in multi-armed bandits, with theoretical guarantees.

Senior Software Engineer, Speech Recognition & Speech Synthesis

Google (London, Mountain View)

2015-2019

- ◊ Improved latency and reduced embedded size of Google text-to-speech (TTS) component across ~40 languages, improving quality of over billion daily queries.
- ◊ Improvements to infrastructure and quality of pronunciation models used by Google’s automatic speech recognition (ASR) and TTS systems.
- ◊ Research on relevant problems in natural language processing; published at top conferences.
- ◊ Promoted with *Superb* (top 4-5% across Google) rating twice within first three years.
- ◊ Engineering Manager experience of over a year — managed three engineers and two interns.

Summer Intern, Information-theoretic Optimal Sensor Placement

Microsoft Research, Redmond

Summer 2014

- ◊ Designed, analyzed and tested new efficient and scalable active learning algorithms for optimal deployment of sensors. Obtained improved performance guarantees (published at ICML 2015).

Summer Intern, Exact and Approximate Learning of Finite Automata

Max Planck Institute for Software Systems, Germany

Summer 2013

- ◊ Worked with Dr. Rupak Majumdar (scientific director, MPI-SWS) on design of black box algorithms for characterizing the states of finite and linear automata.

HONORS AND ACHIEVEMENTS

Awards

- ◊ Invited talk at IJCAI 2025 in the Best Papers from Sister Conferences Track (Montreal).
- ◊ Top Reviewer Award at UAI 2025.
- ◊ [Outstanding Student Paper Award](#) at UAI 2024.

- ◇ Paper selected for Oral presentation (top 3.6% among 744 submissions) at UAI 2024.
- ◇ Paper selected for Oral presentation (top 4 among 28 accepted papers) at ICLR 2022 [SRML](#) workshop.
- ◇ Paper selected for Oral presentation (top 0.6% among 9122 submissions) at NeurIPS 2021.
- ◇ First place at [YinzOR 2019](#) (OR conference with participation from top US universities) poster competition for poster titled “Online optimization of piecewise Lipschitz functions in changing environments”.
- ◇ Invited (among 200 researchers worldwide) to attend the 7th [Heidelberg Laureate Forum](#) 2019.

Academic and Entrance Tests

- ◇ Ranked first (highest GPA) in Computer Science and Engineering Department class of 2015.
- ◇ Awarded second best “all-rounder” among all graduating students (across all departments) in 2015.
- ◇ Secured All India Rank 7 in IIT-JEE (Joint Entrance Examination) 2011 among 500,000 applicants.
- ◇ Twice invited and honored as the Prime Minister’s guest at Republic Day of India (2012, 2010) by the Indian Ministry of Education; All India first (second resp.) among junior college (high school) graduates.

National and International Olympiads

- ◇ Won Gold Medal at the 43rd International Chemistry Olympiad (IChO-2011) held at Ankara, Turkey.
- ◇ Among 35 students from all over India to be selected to attend the Orientation-cum-Selection Camp for the International Olympiad of Astronomy and Astrophysics (IOAA) 2011.
- ◇ Obtained State Ranks 1 and 3 in Regional Mathematics Olympiad 2010 and 2009.

TUTORIALS

- ◇ *Learning-augmented algorithms: Leveraging data for performance with theoretical guarantees.* COLT 2026 (June 2026, with Maryam Aliakbarpour and Sandeep Silwal).
- ◇ *New Frontiers of Hyperparameter Optimization: Recent advances and open challenges in theory and practice.* NeurIPS 2025 (December 2025, with Nina Balcan and Colin White).
- ◇ *Limitations of State-of-the-Art and a New Principled Framework for HPO and Algorithm Selection.* AutoML 2025 (September 2025).
- ◇ *Hyperparameter Optimization and Algorithm Selection: Practical Techniques, Theory, and New Frontiers.* UAI 2025 (July 2025).

WORKSHOPS

Led the organization of the following workshops.

- ◇ *Learning in an Agentic World: Foundations and Challenges* at COLT 2026 (June 2026, with Hedyeh Beyhaghi, Avrim Blum and Han Shao). <https://dravy.ttic.edu/law2026.html>
- ◇ *Learning-driven Algorithms and Machine-aided Proofs (LAMP)* at TTIC (August 2026, with Sandeep Silwal and Ellen Vitercik). <https://dravy.ttic.edu/lamp26.html>

SERVICE

Area Chair: ICLR 2027 (Main track) NeurIPS 2026 (Main track); ICML 2026 (Main track); NeurIPS 2025 (Main track); ICLR 2023 (Tiny Papers track).

Senior Program Committee: AAAI 2027

Recent Reviewer/Sub-reviewer: COLT 2026, FOCS 2026, JMLR 2026, TMLR 2026, UAI 2026, AIS-TATS 2026, SODA 2026, ITCS 2026.

Conference Reviewer: NeurIPS, COLT, SODA, FOCS, AISTATS, ICLR, UAI, AAAI, ALT, ITCS.

Journal Reviewer: JMLR, TMLR, TPAMI, SIMODS, Machine Learning.

Session Chair: AAAI 2025, IJCAI 2025, COLT 2026.

Workshop Proposal Reviewer: NeurIPS 2026.

Organized CMU Learning Theory reading group (Spring 2021, Fall 2021, Spring 2022, Fall 2022).

Co-organizer:

- ◊ IDEAL Annual meeting 2025, a one-day workshop at University of Illinois Chicago (UIC).
- ◊ 2025 IDEAL Get Ready for Research Workshop.

TEACHING

- ◊ Instructor for TTIC 31290 – Machine Learning for Algorithm Design (Fall 2025), co-taught with Avrim Blum.
- ◊ Head teaching assistant for 10-701 (Spring 2022) at CMU “Introduction to Machine Learning (PhD)”:
 - ▷ Designed and graded homeworks and exams, held office hours, ensured prompt responses on Piazza, co-ordinated tasks with other TAs.
 - ▷ Designed and delivered lecture to a class of graduate students (class size 138).
- ◊ Head teaching assistant for 10-422 (Spring 2023) at CMU “Foundations of Learning, Game Theory, and Their Connections”.
- ◊ Taught a two-day course on Machine Learning at Google (2016).
- ◊ Teaching assistant for COL106/CSL201 (Spring 2015) at IIT Delhi “Data Structures”.

PUBLICATIONS

(authors listed alphabetically by last name, except ◊)

Pre-prints

- ★ Balcan, M-F., Blum, A., Fragkia, K., Li, Z., & Sharma, D. *Online Learnability of Chain-of-Thought Verifiers: Soundness and Completeness Trade-offs*. Arxiv 2026.
- ★ Parameter Tuning with Generalization Guarantees for GPU-Accelerated Linear Programming. Prasad S. & Sharma, D. Arxiv 2026.
- ★ Balcan, M-F., & Sharma, D. *Learning accurate and interpretable tree-based models*. Arxiv 2025.
- ★ Balcan, M-F., Goyal S. & Sharma, D. *Distribution-dependent Generalization Bounds for Tuning Linear Regression Across Tasks*. Arxiv 2025.

Peer-reviewed Conferences

- ★ Balcan, M-F., Pegden W. & Sharma, D. *Online Optimization of Piecewise-Lipschitz Functions with Applications to Data-Driven Algorithm Design*. Invited Open Problem, COLT 2026.
- ★ Blum, A., Garicano, M., Ravichandran, K., & Sharma, D. *Algorithm Design and Stronger Guarantees for the Improving Multi-Armed Bandits Problem..* UAI 2026.
- ★ Sharma, D. *Gradient Descent with Provably Tuned Learning-rate Schedules*. AISTATS 2026.
- ★ Balcan, M-F., Blum, A., Li, Z., & Sharma, D. *On Learning Verifiers for Chain-of-Thought Reasoning*. NeurIPS 2025.
- ★ Balcan, M-F., Nguyen, A.T. & Sharma, D. *Sample complexity of data-driven tuning of model hyperparameters in neural networks with structured parameter-dependent dual function*. NeurIPS 2025.

- ★ Sharma, D., & Sun, A. *Conservative classifiers do consistently well with improving agents: characterizing statistical and online learning*. NeurIPS 2025 (**Spotlight, top 3%**).
- ★ Attias, A., Blum, A., Naggita, K., Saleh, D., Sharma, D., Walter, M. *PAC Learning with Improvements*. ICML 2025.
- ★ Du, A., Huang E. & Sharma, D. *Tuning Algorithmic and Architectural Hyperparameters in Graph-Based Semi-Supervised Learning with Provable Guarantees*. UAI 2025.
- ★ Sharma, D. & Suggala A. *Offline-to-online hyperparameter transfer for stochastic bandits*. AAAI 2025.
- ★ Balcan, M-F., Nguyen, A.T. & Sharma, D. *Algorithm Configuration for Structured Pfaffian Settings*. AutoML 2025 (non-archival track, journal version accepted at TMLR 2025).
- ★ Balcan, M-F., Seiler C. & Sharma, D. *Accelerating ERM for data-driven algorithm design using output-sensitive techniques*. NeurIPS 2024.
- ★ Balcan, M-F., Blum, A., Sharma, D. & Zhang, H. *An analysis of the robustness of non-Lipschitz networks*. NeurIPS 2024 (journal to conference track).
- ★ Balcan, M-F. & Sharma, D. *Learning accurate and interpretable decision trees*. UAI 2024 (**Outstanding student paper award**).
- ★ Sharma, D. *No Internal Regret with Non-convex Loss Functions*. AAAI 2024.
- ★ Balcan, M-F., Pozzi, M. & Sharma, D. *Subsidy for repair in component maintenance games*. EMI/PMC 2024.
- ★ Balcan, M-F., Nguyen, A.T. & Sharma, D. *New Bounds for Hyperparameter Tuning of Regression Problems Across Instances*. NeurIPS 2023.
- ★ Balcan, M-F., Hanneke, S., Pukdee, R. & Sharma, D. *Reliable Learning for Test-time Attacks and Distribution Shift*. NeurIPS 2023.
- ◇ Sharma, D. & Jones, M. *Efficiently learning the graph for semi-supervised learning*. UAI 2023.
- ★ Balcan, M-F., Khodak, M., Sharma, D., & Talwalkar A. *Provably tuning the ElasticNet across instances*. NeurIPS 2022 [CMU MLD official [blog post](#)].
- ★ Balcan, M-F., Blum, A., Hanneke, S. & Sharma, D. *Robustly-reliable learners under poisoning attacks*. COLT 2022.
- ★ Balcan, M-F. & Sharma, D. *Data-driven semi-supervised learning*. NeurIPS 2021. (**Oral, top 1%**)
- ★ Balcan, M-F., Khodak, M., Sharma, D., & Talwalkar A. *Learning-to-learn non-convex piecewise-Lipschitz functions*. NeurIPS 2021.
- ★ Balcan, M-F., Dick, T., & Sharma, D. *Learning piecewise Lipschitz functions in changing environments*. AISTATS 2020.
- ◇ Sharma, D., Wilson, M., & Bruguier, A. *Better morphology prediction for better speech systems*. Interspeech 2019.
- ★ Sharma D. *On Training and Evaluation of Grapheme-to-Phoneme Mappings with Limited Data*. Proc. Interspeech 2018 (2018): 2858-2862.
- ◇ Bruguier, A., Bakhtin, A., & Sharma, D. *Dictionary Augmented Sequence-to-Sequence Neural Network for Grapheme to Phoneme Prediction*. Interspeech 2018.
- ◇ Sharma D., Kapoor A., & Deshpande A. *On greedy maximization of entropy*. ICML 2015.

Peer-reviewed Journals

- ★ Balcan, M-F., Nguyen, A.T. & Sharma, D. *Algorithm Configuration for Structured Pfaffian Settings*. TMLR 2025.
- ★ Balcan, M-F., Blum, A., Sharma, D. & Zhang, H. *An analysis of the robustness of non-Lipschitz networks*. JMLR 2023.
- ★ Maran, K., Reddy, S. P., Sharma, D., & Tripathi, A. *Some results on a class of mixed van der Waerden numbers*. Rocky Mountain Journal of Mathematics (2018), 48(3), 885-904.

Workshops

- ★ Balcan, M-F., Pozzi, M. & Sharma, D. *Subsidy design for better social outcomes*. Learning in an Agentic World: Foundations and Challenges (LAW), COLT 2026.
- ★ Attias, A., Blum, A., Naggita, K., Saless, D., Sharma, D., Walter, M. *PAC Learning with Improvements*. Learning in an Agentic World: Foundations and Challenges (LAW), COLT 2026.
- ★ Balcan, M-F., Blum, A., Fragkia, K., Li, Z., & Sharma, D. *Online Learnability of Chain-of-Thought Verifiers: Soundness and Completeness Trade-offs*. AI4Math Workshop at ICML 2026.
- ★ Balcan, M-F. & Sharma, D. *Learning reliably under adversarial attacks, distribution shifts and strategic agents*. Reliable ML from Unreliable Data (NeurIPS 2025).
- ★ Blum, A., Garicano, M., Ravichandran, K., & Sharma D. *Algorithm design and sharper bounds for improving-bandits*. OPT2025: 17th Annual Workshop on Optimization for Machine Learning (NeurIPS 2025).
- ★ Balcan, M-F., Goyal, S., & Sharma, D. *Distribution-dependent Generalization Bounds for Tuning Linear Regression Across Tasks*. DiffCoALG: Differentiable Learning of Combinatorial Algorithms (NeurIPS 2025).
- ★ Attias, A., Blum, A., Naggita, K., Saless, D., Sharma, D., Walter, M. *PAC Learning with Improvements*. 2nd Workshop on Social Choice and Learning Algorithms (IJCAI 2025).
- ★ Sharma, D. *Learning how to step in gradient-based optimization: beyond convexity and smoothness*. 3rd Workshop on High-dimensional Learning Dynamics (ICML 2025).
- ★ Sharma, D. *Gradient descent in presence of extreme flatness and steepness*. Methods and Opportunities at Small Scale (ICML 2025).
- ★ Balcan, M-F., Blum, A., Li, Z., & Sharma, D. *On Learning Verifiers for Chain-of-Thought Reasoning*. Workshop on Reliable and Responsible Foundation Models (ICML 2025).
- ★ Du, A., Huang E. & Sharma, D. *Theoretical Analyses of Hyperparameter Selection in Graph-Based Semi-Supervised Learning*. Workshop on Geometry-grounded Representation Learning and Generative Modeling (ICML 2024).
- ★ Balcan, M-F., Seiler, C. & Sharma, D. *Accelerating data-driven algorithm design using output-sensitive techniques*. Workshop on Learnable Optimization (AAAI 2024).
- ★ Balcan, M-F., Dick, T. & Sharma, D. *Shifting regret for tuning combinatorial algorithms with applications to clustering*. Workshop on Learnable Optimization (AAAI 2024), full version AISTATS 2020.
- ★ Balcan, M-F., Blum, A., Sharma, D. & Zhang, H. *Can non-Lipschitz networks be robust? The power of abstention and data-driven decision making for robust non-Lipschitz networks*. Socially Responsible Machine Learning Workshop (ICLR 2022) **Oral** (4/28 accepted).
- ◇ Sharma D., Sahai S.Y., Chaudhari N., Bruguier A. *Improved pronunciation prediction accuracy using morphology*. The ACL Special Interest Group on Computational Morphology and Phonology (ACL 2021).

- ◇ Sharma D., Sahai S.Y. *Predicting and Explaining French Grammatical Gender*. The ACL Special Interest Group on Typology (NAACL 2021).

Abstracts

- ★ Balcan, M-F., Blum, A., Li, Z. & Sharma, D. *On Learning Verifiers for Chain-of-Thought Reasoning (Extended Abstract)*. Workshop on Reliable and Responsible Foundation Models. (ICML 2025).
- ★ Balcan, M-F., Blum, A., Sharma, D. & Zhang, H. *An Analysis of Robustness of Non-Lipschitz Networks (Extended Abstract)*. Workshop on Safe AI (UAI 2025).
- ★ Balcan, M-F. & Sharma, D. *Learning accurate and interpretable decision trees (Extended Abstract)*. IJCAI 2025 (Best Papers from Sister Conferences).

SELECTED TALKS

- ◇ *Principled hyperparameter tuning in core ML algorithms*. Tutorial on [Leveraging data for performance with theoretical guarantees](#) at COLT 2026.
- ◇ *On Learning Verifiers for Chain-of-Thought Reasoning*.
 - Chicago Junior Theorists Workshop.
- ◇ *New Frontiers of Hyperparameter Optimization: Recent advances and open challenges in theory and practice*.
 - [NeurIPS 2025 tutorial](#), joint with Maria-Florina Balcan and Colin White.
 - Lead organizer and panel moderator.
- ◇ *Distribution-dependent generalization for tuning high-dimensional linear regression*.
 - NITMB-IDEAL Kickoff Event, Fall 2025.
- ◇ *Configuring neural networks and tuning gradient descent hyperparameters*.
 - Algorithms Seminar at Northwestern University, Fall 2025.
- ◇ *Limitations of State-of-the-Art and a New Principled Framework for HPO and Algorithm Selection*. [AutoML 2025](#), [Tutorial 1](#).
- ◇ *PAC Learning with Improvements*.
 - Social choice and learning algorithms (SCaLA) workshop, IJCAI 2025.
 - ICML 2025 (recorded talk).
- ◇ *Hyperparameter Optimization and Algorithm Selection: Practical Techniques, Theory, and New Frontiers*. [UAI 2025](#), [Tutorial 1](#).
- ◇ *Principled Hyperparameter Tuning and Algorithm Selection*.
 - University of Illinois Chicago (UIC), IDEAL Annual Meeting, Summer 2025.
 - Capital Area Theory Seminar (CATS), University of Maryland, Summer 2025.
 - University of Montreal, Summer 2025.
 - Theory Seminar at Johns Hopkins University, Fall 2025.
 - Center for Mathematics and Artificial Intelligence Colloquium, George Mason University, Spring 2026.
- ◇ *Provable tuning of deep learning model hyperparameters*.
 - University of Chicago Theory Lunch, Spring 2025.

- IDEAL workshop on ‘Understanding the Mechanisms of Deep Learning and Generative Modeling’ at Northwestern University, Spring 2025.
- “Talks at TTIC” Seminar, Spring 2025.
- ◇ *Learning accurate and interpretable decision trees.*
 - Invited talk at IJCAI 2025 in the Best Paper from Sister Conferences track.
 - UAI 2024 Oral.
 - Workshop on Learning-Augmented Algorithms, TTIC (August 2024).
 - Invited talk at CMU AI seminar, Fall 2024.
- ◇ *Reliable learning under adversarial attacks.* CMU Theory Lunch, Spring 2024.
- ◇ *Provably tuning the ElasticNet across instances.*
 - CMU Theory Lunch, Fall 2022.
 - NeurIPS 2022 (recorded talk).
- ◇ *Near-optimal robustness for instance targeted poisoning and online meta-learning.*
 - DARPA Meeting, TTIC, Fall 2022.
- ◇ *Data-driven semi-supervised learning.*
 - Scalable Algorithms for Semi-supervised and Unsupervised Learning Workshop 2021 at Google.
 - CMU Theory Lunch.
 - NeurIPS 2021 Oral.
- ◇ *On the power of abstention and data-driven decision making for adversarial robustness.*
 - CMU Theory Lunch, Spring 2021.
 - ICLR 2022 Oral at SRML workshop.
- ◇ *Learning-to-learn non-convex piecewise-Lipschitz functions.*
 - NeurIPS 2021 (virtual).
- ◇ *Learning piecewise Lipschitz functions in changing environments.*
 - CMU AI Seminar.
 - AISTATS 2020 (virtual).
- ◇ *On greedy maximization of entropy.*
 - ICML 2015 Oral.

MENTORSHIP

- ◇ (2019-2026) Mentored several graduate and undergraduate students.
 - ▷ Yifan Wang (undergraduate at Tsinghua University, with Nina Balcan) → PhD at Georgia Tech
 - ▷ Maxwell Jones (undergraduate at CMU) → PhD at CMU
 - ▷ Christopher Seiler (undergraduate at CMU, with Nina Balcan) → Anthropic
 - ▷ Ally Yalei Du (undergraduate at CMU) → PhD at University of Washington
 - ▷ Marten Garicano (undergraduate at University of Chicago) → PhD at Columbia University

- ▷ Eric Huang (graduate at CMU)
- ▷ Anh Tuan Nguyen (graduate at CMU, with Nina Balcan)
- ◇ (2017-18) Manager to three engineers at Google Speech team, Mountain View.
- ◇ (2018) Hosted undergraduate interns at Google, NLP Research published at Interspeech and ACL.
 - ▷ Melissa Wilson → NSF AI2ES (Advances in AI across Earth Systems)
 - ▷ Neha Chaudhari → Google (full-time)

TECHNICAL SKILLS

Proficient: C++, Python (TensorFlow, PyTorch), Bash, AWS.

Experienced: Java, SQL, Javascript.