Jeremy McMahan, PhD

jeremymmcmahan@gmail.com | www.jeremymmcmahan.com

Hybrid Frameworks for Optimization Under Uncertainty

Modern decision-making systems—from decentralized marketplaces to autonomous logistics—operate in a complex world of combinatorial structures and fluctuating information. While recent AI advances enable fast, online decision-making despite incomplete information, these methods often lack robust guarantees and struggle with combinatorial environments. Conversely, classical algorithmic approaches offer strong guarantees for combinatorial problems but cannot make real-time decisions or handle informational gaps imposed by real-world environments. My research, lying at the intersection of **Combinatorial Optimization, Algorithmic Game Theory, and Reinforcement Learning,** bridges these fields to develop efficient solutions for **Optimization Problems under Uncertainty.**

A data-driven world imposes new challenges and opportunities for algorithmic research:

- **Overcoming Worst-Case Barriers (Data as an Asset):** *When historical data or ML predictions are available, how can we leverage them to surpass worst-case theoretical limits?* My approach is to develop specialized learning-augmented algorithms that leverage structural hints—such as warm-starting via predictions or distributional knowledge—to produce solutions with near-optimal approximation ratios.
- **Overcoming Incomplete Information (Data as a Constraint):** *When input specifications are unavailable or unreliable, how can we derive meaningful and robust solutions?* My approach is to develop general-purpose frameworks that embed combinatorial constraints directly into agents. By applying Game Theory, I ensure these systems remain robust even when data is sparse, noisy, or manipulated by strategic and adversarial entities.

By unifying the structural insights of Combinatorial Optimization and the strategic analysis of Algorithmic Game Theory with the adaptive power of Learning, I design systems that are *theoretically sound, strategically robust, and practically deployable,* for the next generation of data-driven infrastructure.

Past Research Contributions

General Frameworks. In settings where domain knowledge is sparse, I view distinct challenges—from Stochastic Knapsack to Autonomous Vehicle Routing—through the unifying lens of the *Constrained Markov Game.*

- To solve this **general formulation**, my work developed hybrid algorithms that integrate *graph search*, *combinatorial constraint characterizations*, and *strategic rounding* directly into the reinforcement learning loop. In [4, 3, 2, 8][ICML, ICML, NeurIPS, AISTATS], we demonstrated that this approach yields algorithms that are provably optimal in terms of approximation guarantees, thereby **resolving several long-standing open problems**.

- Furthermore, addressing the impact of **adversarial attacks** and noisy data on these models in [5, 6, 7] [ICML, RLC, AAAI], we utilized *linear programming*, *dynamic programming*, *duality*, and *bilevel optimization* to devise the first provably efficient and robust algorithms against corrupted or noisy data.

These works demonstrate that by connecting Optimization, RL, and Game Theory, we can build solvers that function in the most general, unstructured environments.

Specialized Algorithms. Conversely, when specific structural or distributional knowledge is available, I exploit this information to move beyond general approximations and design specialized algorithms with stronger guarantees. Here, the goal is to leverage mathematical properties to break through worst-case barriers.

- For example, in the **Pandora's Box** problem, which abstracts sequential search and optimal stopping, we exploited the scenario model of correlated input distributions to establish a novel *recursive reduction* to a special case of the *Optimal Decision Tree problem*. This relationship yielded the first adaptive algorithms for correlated Pandora's box with near-optimal approximation ratios [1][APPROX].
- Similarly, when faced with a **single well-behaved constraint**, my work [2] [NeurIPS] devised a specialized Fully Polynomial-Time Approximation Scheme (FPTAS) using *packing-covering duality* and *approximate dynamic programming*, which circumvents the feasibility hardness of general constrained environments.

Together, these works demonstrate that with additional structure, we can achieve performance guarantees that are impossible in the general case.

Future Research Agenda

In the near term, I plan to push the limits of the Constrained Markov Game model by incorporating richer structural properties. For example, *submodularity* is fundamental to resource allocation and market design, yet remains under-explored in learning contexts. By leveraging known properties of submodularity in graph problems, I aim to derive approximation algorithms for **Submodular-Constrained Markov Games.** Simultaneously, I plan to bridge the gap between pessimistic worst-case bounds and empirical performance using **Smoothed Analysis.**

My long-term vision is to move beyond Markovian-based frameworks to creating a **Unified Theory of Learning-Augmented Optimization.** Since current combinatorial solvers start from scratch for every instance, while current learning agents struggle with combinatorial structures, I propose a hybrid paradigm:

- **Warm-Starting via Primal-Dual Methods:** I will investigate how learned predictions can initialize primal-dual schemas for online combinatorial search. The goal is to use past experience to jump-start the optimization process, achieving performance better than worst-case, while retaining robust fallbacks.
- **Automated Algorithm Design:** Instead of hand-crafting specialized algorithms, I aim to develop meta-algorithms that automatically detect problem structure (e.g., sparsity,

correlation) and adapt the solution strategy. This moves towards *Self-Optimizing Systems* that improve their efficiency over time as they ingest more domain data.

Ultimately, this agenda bridges theoretical rigor and practical automation. By explicitly incorporating details like submodularity or primal-dual structure into learning agents, we can move beyond black-box deployment to build systems that are *efficient, predictable, safe, and robust.*

References

[1] S. Chawla, E. Gergatsouli, J. McMahan, and C. Tzamos. Approximating Pandora's Box with Correlations. In N. Megow and A. Smith, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023)*, volume 275, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[2] J. McMahan. Deterministic policies for constrained reinforcement learning in polynomial time. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 94453–94489. Curran Associates, Inc., 2024.

[3] J. Mcmahan. Anytime-constrained equilibria in polynomial time. In A. Singh, M. Fazel, D. Hsu, S. Lacoste-Julien, F. Berkenkamp, T. Maharaj, K. Wagstaff, and J. Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 43399–43416. PMLR, 13–19 Jul 2025.

[4] J. Mcmahan. Polynomial-time approximability of constrained reinforcement learning. In A. Singh, M. Fazel, D. Hsu, S. Lacoste-Julien, F. Berkenkamp, T. Maharaj, K. Wagstaff, and J. Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 43417–43439. PMLR, 13–19 Jul 2025.

[5] J. Mcmahan, G. Artiglio, and Q. Xie. Roping in uncertainty: Robustness and regularization in Markov games. In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 35267–35295. PMLR, 21–27 Jul 2024.

[6] J. McMahan, Y. Wu, Y. Chen, J. Zhu, and Q. Xie. Inception: Efficiently computable misinformation attacks on Markov games. *Reinforcement Learning Journal*, 5:2345–2358, 2024.

[7] J. McMahan, Y. Wu, X. Zhu, and Q. Xie. Optimal attack and defense for reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(13):14332–14340, Mar. 2024.

[8] J. McMahan and X. Zhu. Anytime-constrained reinforcement learning. In S. Dasgupta, S. Mandt, and Y. Li, editors, *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 4321–4329. PMLR, 02–04 May 2024.