# Publication List

In theoretical work, authors are typically listed alphabetically. Papers using alphabetical authorship are indicated with "(alphabetical order)."

## Journal Papers

[J2] **Chao-Kai Chiang** and Masashi Sugiyama, "Unified Risk Analysis for Weakly Supervised Learning," *Transactions on Machine Learning Research*, 2025. ISSN 2835–8856.
(Regular journal paper)

[J1] Chia-Jung Lee, **Chao-Kai Chiang**, Mu-En Wu, "Online Learning Problems against Dynamic Strategies in Gradually Evolving Worlds," *Journal of Information Hiding and Multimedia Signal Processing* 8(4): 869–879 (2017).
(Regular journal paper)

## Refereed Conference Papers

[C13] Okan Koc, Alexander Soen, **Chao-Kai Chiang**, and Masashi Sugiyama, "Domain Adaptation and Entanglement: an Optimal Transport Perspective," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 3034–3042, 2025.
(Regular conference paper)

[C12] Jongyeong Lee, **Chao-Kai Chiang**, and Masashi Sugiyama, "The Choice of Noninformative Priors for Thompson Sampling in Multiparameter Bandit Models," in *AAAI Conference on Artificial Intelligence (AAAI)*, pp. 13383–13390, 2024.
(Regular conference paper)

[C11] Xin-Qiang Cai, Yu-Jie Zhang, **Chao-Kai Chiang**, and Masashi Sugiyama, "Imitation Learning from Vague Feedback," in *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 48275–48292, 2023.
(Regular conference paper)

[C10] Jongyeong Lee, Junya Honda, **Chao-Kai Chiang**, and Masashi Sugiyama, "Optimality of Thompson Sampling with Noninformative Priors for Pareto Bandits," in *International Conference on Machine Learning (ICML)*, pp. 18810–18851, 2023.
(Regular conference paper)

[C9] Zhiyun Lu, Liyu Chen, **Chao-Kai Chiang**, and Fei Sha, "Hyper-parameter Tuning under a Budget Constraint," in *International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 5744–5750, 2019.
(Regular conference paper)

[C8] Virginia Smith, **Chao-Kai Chiang**, Maziar Sanjabi, and Ameet Talwalkar, "Federated Multi-Task Learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 4424–4434, 2017.
(Regular conference paper)

[C7] Peter Auer and **Chao-Kai Chiang** (alphabetical order), "An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits," in *Annual Conference on Learning Theory (COLT)*, pp. 116–120, 2016.
(Regular conference paper; the full version is [P1])

[C6] Peter Auer, **Chao-Kai Chiang**, Ronald Ortner, and Madalina M. Drugan (alphabetical order for the first three authors), "Pareto Front Identification from Stochastic Bandit Feedback," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 939–947, 2016.
(Regular conference paper)

[C5] Chia-Jung Lee, **Chao-Kai Chiang**, and Mu-En Wu, "Resisting Dynamic Strategies in Gradually Evolving Worlds," in *International Conference on Robot, Vision and Signal Processing (RVSP)*, pp. 191–194, 2015.
(Short conference paper)

[C4] Ku-Chun Chou, Hsuan-Tien Lin, **Chao-Kai Chiang**, and Chi-Jen Lu, "Pseudo-reward Algorithms for Contextual Bandits with Linear Payoff Functions," in *Asian Conference on Machine Learning (ACML)*, pp. 344–359, 2014.
(Regular conference paper)

[C3] **Chao-Kai Chiang**, Chia-Jung Lee, and Chi-Jen Lu (alphabetical order), "Beating Bandits in Gradually Evolving Worlds" in *Annual Conference on Learning Theory (COLT)*, pp. 210–227, 2013.
(Regular conference paper)

[C2] **Chao-Kai Chiang**, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu, "Online Optimization with Gradual Variations," in *Annual Conference on Learning Theory (COLT)*, pp. 6.16.20, 2012.
(Regular conference paper, **Mark Fulk Best Student Paper Award**)

[C1] **Chao-Kai Chiang** and Chi-Jen Lu, "Online Learning with Queries," in *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 616–629. SIAM, 2010.
(Regular conference paper)

## Preprints

[P4] **Chao-Kai Chiang**, Takashi Ishida, and Masashi Sugiyama, "LLM Routing with Dueling Feedback," arXiv:2510.00841.
(Preprint; under double-blind review)

[P3] Jongyeong Lee, **Chao-Kai Chiang**, and Masashi Sugiyama, "Asymptotically Optimal Thompson Sampling Based Policy for the Uniform Bandits and the Gaussian Bandits," arXiv:2302.14407v1.
(Preprint)

[P2] Zhiyun Lu, **Chao-Kai Chiang**, and Fei Sha, "Hyper-parameter Tuning under a Budget Constraint," arXiv:1902.00532.
(Preprint)

[P1] Peter Auer and **Chao-Kai Chiang** (alphabetical order), "An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits," arXiv:1605.08722.
(Preprint; the full version of [C7])

## Supplementary Information on Non-Regular Publications

[C5] Chia-Jung Lee, Chao-Kai Chiang, and Mu-En Wu, "Resisting Dynamic Strategies in Gradually Evolving Worlds," in *International Conference on Robot, Vision and Signal Processing (RVSP)*, pp. 191194, 2015. (Short conference paper)

The short paper [C5] presents a minor extension to the Hedge algorithm.

[P4] Chao-Kai Chiang, Takashi Ishida, and Masashi Sugiyama, "LLM Routing with Dueling Feedback," arXiv:2510.00841. (Preprint; under double-blind review)

This preprint is currently under a double-blind conference review.

[P3] Jongyeong Lee, Chao-Kai Chiang, and Masashi Sugiyama, "Asymptotically Optimal Thompson Sampling Based Policy for the Uniform Bandits and the Gaussian Bandits," arXiv:2302.14407v1. (Preprint)

This preprint is a preliminary version of [C12].

[P2] Zhiyun Lu, Chao-Kai Chiang, and Fei Sha, "Hyper-parameter Tuning under a Budget Constraint," arXiv:1902.00532. (Preprint)

This preprint is a preliminary version of [C9].

[P1] Peter Auer and Chao-Kai Chiang (alphabetical order), "An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits," arXiv:1605.08722. (Preprint; the full version of [C7])

This preprint is the full version of [C7].