Dominique C. Perrault-Joncas

Mercer Island, WA | (206) 661-1424 | dcpjoncas@gmail.com | Google Scholar

Summary

Principal Scientist with 13+ years applying reinforcement learning, causal inference, and LLM post-training (RLHF / preference optimization) to large-scale production decision systems. Led tech for cross-functional teams (10–15+), published at top venues (NeurIPS, TMLR), and delivered multiple nine-figure impact programs across experimentation, inventory placement, and forecasting at Amazon and Google.

Selected Impact

- Led tech for RL initiatives driving investment in engineering teams to scale a **nine-figure** opportunity in North America automated ordering systems.
- Demonstrated capacity-constrained RL ordering improving profit **+1.8%** under **10%** tighter constraints vs. production.
- Refined product launch criteria via meta-analysis of randomized experiments; backtests indicate **nine-figure** incremental impact vs. prior rules.
- Delivered an experimentation framework that reduced the risk of failures by 10×, mitigating potential **eight-figure** downside per failed test.

Technical Skills

Reinforcement Learning: Model-based RL, constrained RL, sim2real
LLM Post-Training: RLHF, preference optimization (DPO / C2-DPO), fine-tuning, evaluation
Game Theory & Market Design: marketplace equilibrium, competitive pricing mechanisms
Causal & Experimentation: interference-aware experimentation under constraints, optimal treatment effects estimation, Heavy-tailed Outcome Analysis
Engineering: Python, PyTorch, SQL/Spark, AWS

Experience

Amazon Jan 2020 – Present | Seattle, WA
Principal Research Scientist, Supply Chain Optimization Technologies

- **RL for Inventory Buying & Placement (Tech Lead, 2024–Present):** Co-led a team of 10+ scientists and engineers building reinforcement learning systems for inventory buying and placement decisions spanning multi-billion-dollar annual inventory spend; success drove follow-on investment in dedicated engineering teams to scale a nine-figure opportunity in North America.
- **Capacity-Constrained RL Buying:** Developed RL buying methods under capacity constraints; demonstrated +1.8% profit improvement even with 10% tighter buying constraints versus production baselines.
- **RL Architecture & Sim2Real:** Defined action space and value function design for RL-based placement, including a quadratic coordination interface across subsystems; launched first production labs for placement RL agents. Directed sim2real evaluation frameworks bridging simulation-to-production gaps.
- **Experimentation under Joint Constraints:** Developed and deployed an interference-aware experimentation methodology for resource-constrained environments; reduced risk of unsuccessful scaled experiments by 10×, mitigating potential eight-figure downside per failed large-scale test.
- **Meta-analysis & Launch Decisions:** Led meta-analysis of randomized experiments to improve treatment effect estimation for heavy-tailed outcomes; backtests across multiple orgs indicate updated launch decision criteria would have delivered nine-figure incremental impact versus prior rules (NeurIPS 2022 Workshop, with M.I. Jordan).

- **LLM Post-training / Preference Optimization:** Co-developed a unifying "implicit classification" framework for DPO-style algorithms; identified under-specification leading to winner–loser probability collapse and proposed C2-DPO, a constrained preference optimization method with an RLHF interpretation that improves alignment quality on standard preference datasets and multiple LLM checkpoints (arXiv 2025).
- **LLM-Augmented Forecasting:** Co-developed LLMForecaster, integrating unstructured text into seasonal demand forecasting via LLM embeddings (NeurIPS 2024 Workshop).
- **Leadership:** Tech lead for 15+ researcher RL group; built advisory committee governing experimental platform changes; Senior Area Chair for Amazon Consumer Science Summit (2021–26); mentored junior scientists.

Google Apr 2015 – Dec 2019 | Seattle, WA

Senior Quantitative Analyst / Statistician, Advanced Measurement Technologies

- **Cross-Media Measurement (Tech Lead, 2017–2019):** Led a team of 3 statisticians on a cross-media measurement panel initiative. Designed experiments for face recognition and gaze estimation systems for production deployment.
- **Indoor Positioning & Prediction:** Developed indoor positioning models using WiFi, Bluetooth, and sensor fusion. Built box office prediction models; designed methods for imputing missing user data at scale.

Amazon Nov 2012 – Apr 2015 | Seattle, WA

Research Scientist, Demand Forecasting

- Developed first demand forecasting models for Amazon Pantry at launch; scaled Bayesian and quantile regression to big data via MapReduce. Mentored scientists and led multiple research projects.

Submitted / Under Review

- H. Parikh, G. Levin-Konigsberg, **D. Perrault-Joncas**, A. Volfovsky. *Mind the Sim-to-Real Gap & Think Like a Scientist.* Submitted to **NeurIPS 2026**; arXiv.
- A. Wang, **D. Perrault-Joncas**, A. Maggiar, D. Foster, C. Eisenach. *Ready from Day 1: Population-Aware Coordination for Large-Scale Constrained Multi-Agent Systems.* Submitted to **NeurIPS 2026**; arXiv.
- H. Parikh, G. Levin-Konigsberg, **D. Perrault-Joncas**, A. Volfovsky. *TEA-Time: Transporting Effects Across Time.* Submitted to **NeurIPS 2026**; arXiv.
- R.J. George, C. Eisenach, U. Ghai, **D. Perrault-Joncas**, A. Anandkumar, D. Foster. *BRIDGE: Building Representations In Domain Guided Program Synthesis.* Submitted to **NeurIPS 2026**; arXiv.

Selected Publications

- K. Asadi, X. Xu, . . . , **D. Perrault-Joncas**, . . . , M. Ghavamzadeh. *C2-DPO: Constrained Controlled Classification DPO.* Accepted at **TMLR**, 2025. OpenReview.
- A. Maggiar, . . . , **D. Perrault-Joncas**. *Structure-Informed Deep RL for Inventory Management.* **NeurIPS 2025** Workshop.
- H. Zhang, . . . , **D. Perrault-Joncas**, . . . , A.G. Wilson. *LLMForecaster: Improving Seasonal Event Forecasts with Unstructured Textual Data.* **NeurIPS 2024** Workshop.
- N. Tripuraneni, **D. Perrault-Joncas**, D. Madeka, D. Foster, M.I. Jordan. *Meta-Analysis of Randomized Experiments with Applications to Heavy-Tailed Response Data.* **NeurIPS 2022** Workshop.
- T. Ding, **D. Perrault-Joncas**, O. Ronen, M.I. Jordan, D. Bergemann, D. Foster, O. Gottesman. *Marketplace Operators Can Induce Competitive Pricing.* arXiv, 2025.
- R. Pabari, U. Ghai, **D. Perrault-Joncas**, K. Torkkola, O. Ronen, D. Madeka, A. Rubinstein, D. Foster, O. Gottesman. *A Shared-Revenue Bertrand Game.* arXiv, 2025.

Also: 3 NeurIPS papers on manifold learning (2011, 2016, 2017); 25+ internal papers at Amazon.

Education: Ph.D. Statistics, UW (M. Meilă), 2012 | M.Sc. Applied Math, McGill, 2007 | B.Sc. Math & Physics, McGill, 2003