Memo Ozdincer

memo@cs.toronto.edu • (647) 518-7370 • linkedin.com/in/memo-ozdincer • github.com/memo-ozdincer

Education

University of Toronto, BASc in Engineering Science (Eng. Physics + AI minor) – Dean's list, 3.7 GPA — Jun 2028

Experience

Jinesis AI Lab (Vector Institute) - AI Researcher (Safety) — Aug 2025 – Present
Advisors: Zhijing Jin, Bernhard Schölkopf

- Designed a weight-level LLM agent defense competitive with Meta SecAlign, OpenAI IH, Google CaMeL, and MELON on the ASR-utility Pareto across standard agentic benchmarks (<0.8% ASR on AgentDojo, InjecAgent). First-author NeurIPS submission.
- Built a 100K-trace dataset of entirely benign-looking prompt injections, plus a 566K JEPA-format training set, reverse-generating prompts from 15 jailbreak datasets × 20 augmentation methods (COLM, NeurIPS).
- Rolled out and evaluated >1M traces (1B+ tokens) across Llama, Qwen, Mistral, DeepSeek, and Gemma models (8B–685B) in Triton and TensorRT-LLM. Parallelized across 4xH100, 8xH200, and 2x/4xB200 nodes on bare metal/NVLink, never letting GPU utilization drop below 90%.

Matter Lab (Vector Institute + NVIDIA) - AI Researcher (Generative Chemistry) — Sep 2025 – Present
Advisor: Alán Aspuru-Guzik

- Beat Sella, the status-quo saddle-search algorithm, on highly noisy atomic potentials: 79% convergence at 1Å noise vs. Sella's 53%. Submitting to NeurIPS 2026.
- Redesigned search algorithms to be differentiable for generative model training, enabling adjoint backpropagation to train diffusion models on our custom vector field.
- Parallelized over 3.8M molecular dynamics simulations across 12 nodes of 4× NVIDIA A100 SXM4 (NVLink) using our HIP MLIP model, developed with NVIDIA.

National University of Singapore (SERIS) - ML Researcher (Comp. Physics) — May 2025 – Aug 2025
Advisors: Erik Birgersson, Armin Aberle

- Designed Physics-Informed ML model for solar cell characterization, replacing a slow coupled PDE solver (~4800 s/device) with millisecond inference, a 100,000x speedup. Deployed at SERIS for high-throughput screening.
- Trained on 500K simulations (31 quantum solar-cell properties spanning 21 OOM). Hand-derived physics priors from 71 quantum and optoelectronic relations.

aUToronto (UofT AutoDrive Team) - Software Engineer (Computer Vision) — May 2025 – Present

- Implemented BEVFusion-based sensor fusion pipeline (C++/Python) processing 1.1 GB/s from 6 cameras + 2 LiDARs at under 90ms latency for real-time tracking.
- Fine-tuned lightweight detection models (Twin+, YOLOPV2, YOLO11) with 35% perception speedup in adverse weather. Wrote LiDAR–camera calibration tooling in C++, reducing false-positive edge detections by 85%+.

Skills

AI/ML: LLM & Agent Safety, Fine-Tuning, Evals, Reinforcement Learning, Diffusion Models, CV (YOLO, BEVFusion)

Languages & Frameworks: Python, C++17, C, CUDA, PyTorch, vLLM, Triton, TensorRT-LLM, Hydra/OmegaConf, DVC, W&B, Tensorflow, Docker, Kubernetes, HuggingFace, AWS, SLURM/HPC