

IGOR BOGDANOV

ON • Canada • Open to Relocation • igor@isbogdanov.com

[GitHub](#) • [LinkedIn](#) • [Google Scholar](#) • [Personal Website](#)

AI Systems Researcher, Research Engineer & Builder: LLM Harnesses, Evaluation, Reliability, Interpretability

Python • PyTorch • JAX • Multi-Agent DRL • LLM Agents • Transformers • NLP • Adaptive Systems • Continual Learning

AI systems researcher, research engineer and builder with 15+ years of production software and systems experience, and recent MASc research in reliable LLM systems, compound agents, cost-aware evaluation, prompt-only self-improvement, and mechanistic interpretability of reasoning representations. I build full experimental systems: agent harnesses, evaluation pipelines, multi-provider LLM infrastructure, logging, and applied AI prototypes. My recent work includes two first-author ACM CAIS 2026 papers with all ACM reproducibility artifact badges, an ICML 2026 Mechanistic Interpretability Workshop paper, an ICLR 2026 workshop paper, and 2025 publications at NeurIPS MTI, EDM WLIEDM, and COMPSAC MediComp. Invited Presenter of my ACM CAIS work at AI Engineer World's Fair 2026 in San Francisco.

LIST OF SELECTED PUBLICATIONS

- **Bogdanov, I.**, & Huang, C. "Discovering Cross-Language Reasoning Invariance in LLMs with Geometry-Invariant Sparse Autoencoders." [ICML 2026 Workshop on Mechanistic Interpretability](#).
- **Bogdanov, I.**, Lung, C.-H., Kunz, T., Gao, J., Taylor, A., & Zaman, M. "Context, Reasoning, and Hierarchy: A Cost-Performance Study of Compound LLM Agent Design in an Adversarial POMDP." **ACM Conference on AI and Agentic Systems (CAIS)**, 2026. DOI: [10.1145/3786335.3813149](#).
- **Bogdanov, I.**, Lung, C.-H., Kunz, T., Gao, J., Taylor, A., & Zaman, M. "FORGE: Self-Evolving Agent Memory With No Weight Updates via Population Broadcast." **ACM Conference on AI and Agentic Systems (CAIS)**, 2026. DOI: [10.1145/3786335.3813155](#).
- **Bogdanov, I.**, Manakina, O., & Lung, C.-H. "Evaluation of Multi-Turn Consistency in LLM Agents: Survival Analysis and Failure-Rationale Taxonomy." [ICLR 2026 Workshop on LLM Logical Reasoning](#).
- Manakina, O., **Bogdanov, I.**, & Lung, C.-H. "Delay-of-Gratification as a Multi-Agent Survival Micro-Benchmark for Long-Horizon LLMs: Social Exposure, Personas, and Tool Use Budgets." [NeurIPS 2025 Workshop on Multi-Turn Interactions in Large Language Models](#).
- Manakina, O., & **Bogdanov, I.** "Learning to Grade Efficiently: A Bandit-Driven Prompt-Selection Framework for Low-Cost LLM Essay Scoring." [EDM WLIEDM](#), 2025.
- **Bogdanov, I.**, & Green, J. "Infant Care Video Dataset for Classification of Interventions Using Transformers." [COMPSAC MediComp](#), 2025. DOI: [10.1109/COMPSAC65507.2025.00299](#).

EDUCATION

CARLETON UNIVERSITY, Ottawa, ON, Canada (2024-2026)

Master of Applied Science, CGPA 11.6/12.0 (GPA 4.0/4.0) Major: Electrical and Computer Engineering.

Awards: Queen Elizabeth II Graduate Scholarship in Science and Technology (Awarded to only 19 students at Carleton)

Thesis: Design and Gradient-Free Adaptation of Compound LLM Agents in a Partially Observable Adversarial Networked Environment (Nominated for Senate Medal)

CARLETON UNIVERSITY, Ottawa, ON, Canada (2019-2024)

Bachelor of Engineering, Double Minor in Mathematics & Physics, CGPA 10.78/12.0

Major: Computer Systems Engineering. | **Minors:** Mathematics and Physics.

Awards: High Distinction, Dean's List & Scholarship: 2020, 2021, 2022

MOSCOW STATE REGIONAL UNIVERSITY, Moscow, Russia (2000-2005)

Specialist Diploma (5 years), Cumulative GPA 4.7/5 Major: Linguistics and Intercultural Communication

PROFESSIONAL & RESEARCH EXPERIENCE

GRADUATE AI RESEARCHER — Carleton University, Ottawa (May 2024-May 2026)

— *Realistic and Generalizable Training of Autonomous Cyber Agents:*

Designed and implemented compound LLM-agent systems for decision-making in complex, stochastic, partially observable environments. Built experiment infrastructure. Produced two first-author ACM CAIS 2026 papers and presented the work at AI Engineer World's Fair 2026.

FOUNDER & TECHNICAL DIRECTOR — Appalect, Ottawa (2018-Present, Part-Time)

— Led architecture and hands-on development across web, mobile, backend, automation, and integration systems.

— Built applied AI/LLM tooling and research infrastructure for multi-provider inference, experiment orchestration, configuration, retries, logging, and evaluation workflows.

— Worked across the full product lifecycle from concept and prototyping to implementation, deployment, and iteration.

FOUNDER & CTO — IBCICO Development, Moscow (2009—2017)

- Founded and led a full-cycle software development agency, managing an engineering team through the entire software development life cycle building financial data systems, e-commerce and ticket booking platforms.

TEACHING ASSISTANT — Carleton University, Ottawa (May 2023-May 2026)

- *Intro to Machine Learning (2025-2026)*: Designed assignments/project, designed & conducted weekly hands-on AI/ML/DL tutorials for **100+** students;
- *Computer Systems Lab*: Supervised and mentored teams building a reactive autonomous robot vehicle;
- *Real-Time Concurrent Systems*: Supervised and mentored teams building a simulation of a concurrent system;
- *Intro to Python*: Grading and mentoring in a Python-based introductory data analysis course.

REVIEWER — NeurIPS 2026, GLOBECOM 2026, ICLR 2026 LLA and COLM 2026 LLA. (February 2026-Present)**AI/ML RESEARCH PROJECT HIGHLIGHTS****Reliable and Self-Improving LLM Agents for Adversarial Partially Observable Environments** — (2024–2026)

- Built a language-based LLM ReAct-style self-learning hierarchical reasoning harness and developed a population-based self-evolving protocol for CAGE Challenge 2, a stochastic adversarial POMDP benchmark, improving episode rewards by $7.7\times$ over baselines for the weakest model (Gemini-2.5-flash-lite). Produced two first-author ACM CAIS 2026 papers: one paper explains cost-performance trade-offs in compound LLM agent harnesses, and the other one introduces FORGE, a prompt-only memory-evolution protocol for self-improving agents without weight updates.

Text-Based Multi-Agent Gym Modeled as MDP/POMDP (2025)

- Built a benchmark gym to conduct experiments on multi-agent interactions with LLMs. Created an experimental setup, collected data for publication accepted to **NeurIPS 2025 MTI**.

Multi-armed Bandit Essay Grader Experiment (2025)

- Built a multi-armed bandit-based agent for efficient prompt selection (*78.4% reduction in LLM calls*). Created an experimental setup and collected data for publication accepted to **EDM 2025 WLIEDM**.

Federated-Learning Anomaly Detection (2024)

- Built a traffic generation containerized Mininet simulation and a federated autoencoder-based anomaly detection pipeline from network traffic logs.

Infant Care Video Dataset and Transformer Benchmark (2024)

- Created a **4,144-video, 12-class**, privacy-compliant NICU dataset; fine-tuned and benchmarked video transformers (TimeFormer/MotionFormer) reaching **94% top-1 accuracy** on neonatal interventions <https://www.infantcaredataset.org/>. Related publication was accepted to **COMPSAC 2025 MediComp**.

Physiotherapy Diagnostic Decision-Making Automation (2023)

- Researched and developed an expert system for self-diagnostics of patients, implemented a cloud-hosted app that converts admin-supplied decision trees into adaptive questionnaires, then outputs diagnosis & recommendations.

TECHNICAL SKILLS

Programming Languages, Storage & AI/ML Frameworks: Python, C/C++, Java, JavaScript, Verilog, SystemVerilog, Swift, Objective-C, Ruby, PyTorch, JAX, TensorFlow, Hugging Face, NumPy, Scikit-learn, Pandas, LangChain, OpenAI APIs, Anthropic APIs, Gemini APIs, CUDA/GPU, vector search (FAISS), RAG, SQL, MySQL, PostgreSQL, MongoDB, Firebase, SQLite, Redis, DuckDB;

ML, Post-Training & Adaptation Methods: Transformers, CNNs, GNNs, RNNs, RL (PPO, DPO, GRPO), LoRA, SFT, Prompt Engineering, In-context and Continual Learning;

AI Systems, CI/CD & Infrastructure: Multi-agent orchestration, GPU/Cloud, Bash, Linux/Unix, Make, Git, GitHub Actions, Docker, LXI, FPGA Prototyping, GPU and Network Optimization, Ethernet, AXI4, PCIe, TCP/IP.

ADDITIONAL HARDWARE PROJECTS**RISC-V DESIGN-VERIFICATION Testbench** (2025)

- Built UVM-inspired scripted and classic class-based UVM self-checking testbenches with *riscv-dv* constrained-random stimulus generation (8-seed regression log provided, 0 fails) • reached *88% line, 79% branch, 65% automatic Questa coverage* (UCDB merged → HTML Dashboard) • Formally verified two adder properties (correctness, commutativity) in SymbiYosys (depth 20);

RISC-V SINGLE-CYCLE CPU DESIGN & ISA Implementation (2022)

- Designed and built a synthesizable RV32I single-cycle Patterson-Hennessy-style CPU core in Verilog from the official RISC-V spec and implemented Xilinx simulation testbench comprising 20-instruction assembly regression covering arithmetic, branch, load/store, and jump paths, with GTKWave traces confirming the final state.

CMOS Pseudo-Random Sequence Generator — CUSoI SOI CMOS MPW (Fall 2022)

- Designed full-chip layout and 5-bit LFSR PRSG with on-chip two-phase clock and closed DRC/LVS in Cadence design suite; after fabrication, validated 31-state ≈ 2.7 kHz output on probe station and summarized in a project report.