

Marius-Constantin Dinu

Dipl.-Ing. Dr. techn., BSc

Curriculum Vitae | Neuro-Symbolic AI, Reinforcement Learning & Research Automation

Wels / Linz, Austria | +43 676 924 43 82 | dinu.marius.constantin@gmail.com

linkedin.com/in/mariusconstantindinu | github.com/Xpitfire | [Google Scholar](https://scholar.google.com/citations?user=...) | dinu.at



RESEARCH PROFILE

My work sits on one thesis: a language model should be used as a *semantic parser* that turns unstructured intent into typed, checkable structure — not as an oracle whose output is trusted. Everything follows from that. In research it produced SymbolicAI, a framework composing generative models with symbolic solvers over an expression graph, and a line of work on parameter choice in unsupervised domain adaptation where no target labels exist to validate against. In practice it produced a company, a patent application, and production platforms on which autonomous agents run real work under enforced policy and bounded cost.

PROFESSIONAL EXPERIENCE

Founder & Chief Executive Officer | ExtensityAI FlexCo

Wels, Austria | Mar 2023 – Jun 2026

- Set the technical thesis and architected the core system: the language model emits typed expressions, which execute on a symbolic layer where every step is checkable. Implemented as contract-mediated calls — pre- and post-conditions bound to schemas, deterministic validators, bounded automatic repair before escalation.
- Delivered pipelines measured at up to **300×** more compute- and cost-efficient than comparable conventional LLM approaches on target workloads, by routing deterministic sub-tasks to solvers rather than re-prompting a large model.
- Raised a **EUR 767k** pre-seed, initiated a six-figure quantum-computing grant with PlanQC, and reached **EUR 900k** revenue within the first year of trading.
- Productised the stack as Extensity Research Services, an automated research platform, and shipped the Story Editor with story.one.
- Recruited and led an international team across AI/ML research, software engineering and business intelligence.
- Filed a US patent application on agent-based neuro-symbolic workflow automation and AI-for-science decision intelligence.

Founder & Principal Architect | Alpha Omega Labs — Cortex & omegaXiv

Austria | 2025 – Present

- Architect and principal engineer of two production platforms totalling roughly **800,000 lines** of first-party Python and TypeScript. Detailed under *Selected Systems*.
- Set the engineering doctrine both platforms enforce mechanically rather than by review: layered import direction checked in CI, consumer-driven contract tests that load the producer's real payload builder, and a harness that converts any skipped or empty test tier into a build failure.

AI Research Director & Advisor | Atlas

Vienna, Austria | Apr 2023 – Dec 2024

- Directed the applied research agenda across deep learning and reinforcement learning; advised leadership on AI strategy, model selection and scalable deployment.
- Built production-facing prototypes in Python using deep and recurrent architectures, from problem framing through evaluation to hand-off.

Research Scientist (Ph.D.) | Johannes Kepler University — LIT AI Lab

Linz, Austria | Sep 2019 – Sep 2024

- Solved hyper-parameter selection in *unsupervised* domain adaptation by extending weighted least squares to vector-valued deep networks, proving the aggregate's target error is asymptotically within a factor of two of the unknown optimal aggregation. Outperformed DEV and IWV on all seven benchmark datasets spanning text, images, EEG, body-sensor and mobile-phone signals. **ICLR 2023 Oral (notable, top 5%)**.
- Derived a justified regularisation-parameter rule for distance-regularised domain adaptation by recasting it as a regularised ill-posed inverse problem and applying the balancing principle — admitting, unlike prior art, source and target distributions with disjoint supports (**NeurIPS 2021**).
- Co-developed Align-RUDDER, replacing RUDDER's LSTM with a multiple-sequence-alignment profile model from bioinformatics, enabling reward redistribution from as few as two demonstrations; first in its class to mine a diamond in MineCraft ObtainDiamond (**ICML 2022**).

- Defined SACo and TQ, empirical measures of dataset exploration and exploitation that predict which offline-RL algorithm family will succeed on a given dataset (**CoLLAs 2021**).
- Investigated exploration strategies under continual domain shift, showing policy-gradient methods adapt faster than Q-learning in lifelong settings (**CoLLAs 2022**).

Senior Machine Learning Researcher | Dynatrace Research

Linz, Austria | Oct 2021 – Mar 2023

- Built deep learning and NLP prototypes at the LIT Open Innovation Center, collaborating across research and product to integrate ML capabilities into observability workflows.

Founder & Software Architect | Imagine Kara LLC

Delaware, USA | May 2018 – Dec 2018

- Partnered with Apollon to build a cloud masternode platform for automated coin hosting; owned architecture, testing and deployment for fintech applications in cryptocurrency and distributed computation.

Data Scientist | CELUM

Linz, Austria | Oct 2016 – Aug 2019

- Developed and deployed deep-learning image classification for automated tagging in a digital asset management product.

Applied Research Scientist | Siemens Corporate Research

Princeton, NJ, USA | Mar 2016 – Sep 2016

- Built a cross-platform handwritten-character recognition application combining SVMs and neural networks, with a C++ inference core bridged to Xamarin and the Android NDK through native interop.

Software Architect & Product Manager | Kontron

Linz, Austria | Dec 2010 – Mar 2016

- Led an international team delivering software for self-service coin-counting devices; designed the SOAP/web-service integration layer spanning Java and C#/.NET.

SELECTED SYSTEMS & OPEN SOURCE

Cortex — Governed Control Plane for AI Agents | Alpha Omega Labs

2025 – Present

- A self-hostable control plane that lets an organisation run AI coding agents from ordinary chat channels under enforced policy, with a full audit trail and a governed path to deployment. Roughly **182,600 lines** of first-party Python and TypeScript: 53 backend service modules, a 17-area React operations console, 643 Python test functions, 13 Compose stacks, 5 CI/CD pipelines.
- Two-plane authorisation: 15 control surfaces $\times$ 4 account groups $\times$ per-runtime capability policies, so a lower-privilege human can drive a higher-capability agent without escalation. Append-only audit log; TTL-scoped cross-organisation delegation.
- **Gauntlet**, a four-track agent-harness benchmark on UK AISI's Inspect AI composing eight public datasets, running safety probes inside a network-blocked sandbox against a fabricated decoy. Measures *harness delta* rather than absolute model score, because scaffold choice can move results further than the model does.
- Formal adversarial review of the agent wrapper with every finding closed and regression-tested: normalised command classifier resistant to several evasion classes, fail-closed policy bridge and approval gate, OS-level sandbox backstop.

omegaXiv — Autonomous Research Platform | Alpha Omega Labs

2025 – Present

- Nine-service platform on which AI agents run a scientific investigation end to end — plan, execute code in isolated sandboxes, self-check against quality contracts, publish a reviewable paper. Roughly **613,600 lines** of application code, 8,100 Python and 3,900 TypeScript tests behind an 80% coverage gate, 70 migrations, 11 CI/CD pipelines, 7 Terraform modules on Azure.
- Run engine: a deterministic *render* $\rightarrow$ *turn* $\rightarrow$ *parse* $\rightarrow$ *validate* $\rightarrow$ *repair* phase loop binding every model output to a per-phase schema, classifying failures into a six-value verdict and auto-repairing within bounded attempts before escalating.
- Spend governance in the run lifecycle: payment pre-authorisation before dispatch, settled by capture-or-cancel through a single idempotent finaliser; cross-replica model throttle; token-budgeted context packager.
- Autonomous supervisor detecting interrupted runs, deciding with an LLM, and applying a typed intervention or escalating — with row-level locking so a supervisor and a human cannot double-apply the same fix.

SymbolicAI — Open-Source Neuro-Symbolic Framework | ExtensityAI

2022 – Present

- Author and maintainer. Composes generative models with symbolic solvers over an expression graph, decomposing a complex task into small operations that are individually verifiable and recombining. **1,745 GitHub stars, 91 forks**, BSD-3-Clause; published at **CoLLAs 2024** (arXiv:2402.00854).
- Recognised publicly by Gary Marcus, Sepp Hochreiter, Jim Fan, an IBM Executive Architect and Microsoft's Data & AI Product Lead; covered by MarkTechPost, Science Times, Montreal.AI and mind-verse.

Story Editor — Human-Centric AI Publishing | story.one $\times$ ExtensityAI

2024 – 2025

- Led the AI architecture for an 18-month partnership productising the neuro-symbolic stack into a six-stage authoring workflow (search, read, summarise, write, improve, finalise) with structured outputs, cross-reference consistency checks and human-in-the-loop review. Launched ahead of the Frankfurt Book Fair 2025.

EDUCATION

Ph.D. (Dr. techn.), Artificial Intelligence — *passed with distinction* | Johannes Kepler University Linz, Austria | 2019 – 2024
 Thesis: *Parameter Choice and Neuro-Symbolic Approaches for Deep Domain-Invariant Learning*. Advisor: Sepp Hochreiter. Institute for Machine Learning, LIT AI Lab.

Executive MBA, Economics | WU Executive Academy & UMN Carlson School Vienna / Minneapolis | 2024 – 2026

M.Sc., Computer Science — Data Science | Johannes Kepler University Linz, Austria | 2017 – 2019

Exchange Semester, Machine Learning | National Taiwan University of Science and Technology Taipei, Taiwan | 2018 – 2019

B.Sc., Software Engineering | University of Applied Sciences Upper Austria Wels, Austria | 2013 – 2016

PUBLICATIONS

380 citations | h-index 8 | i10-index 8 | Google Scholar

- **HyDRA: A Hybrid-Driven Reasoning Architecture for Verifiable Knowledge Graphs.** A. Kaiser, C. Leoveanu-Condrei, R. Gold, M.-C. Dinu, M. Hofmarcher. arXiv:2507.15917, 2025.
- **Primality Testing via Circulant Matrix Eigenvalue Structure: A Novel Approach Using Cyclotomic Field Theory.** M.-C. Dinu. arXiv:2505.00730, 2025.
- **Artificial Intelligence, Market Power and India in a Multipolar World.** S. Srivastava, Abhivardhan, M.-C. Dinu. 2025.
- **SymbolicAI: A Framework for Logic-Based Approaches Combining Generative Models and Solvers.** M.-C. Dinu, C. Leoveanu-Condrei, M. Holzleitner, W. Zellinger, S. Hochreiter. *CoLLAs*, 2024. arXiv:2402.00854.
- **Large Language Models Can Self-Improve at Web Agent Tasks.** A. Patel, M. Hofmarcher, C. Leoveanu-Condrei, M.-C. Dinu, C. Callison-Burch. arXiv:2405.20309, 2024.
- **Parameter Choice and Neuro-Symbolic Approaches for Deep Domain-Invariant Learning.** M.-C. Dinu. Ph.D. Thesis, Johannes Kepler University, 2024.
- **Addressing Parameter Choice Issues in Unsupervised Domain Adaptation by Aggregation.** M.-C. Dinu, M. Holzleitner, M. Beck, H. D. Nguyen, A. Huber, H. Eghbal-Zadeh et al. *ICLR*, 2023 — **Oral, notable top 5%**.
- **Align-RUDDER: Learning from Few Demonstrations by Reward Redistribution.** V. P. Patil, M. Hofmarcher, M.-C. Dinu, M. Dorfer, P. M. Blies, J. Brandstetter et al. *ICML*, 2022.
- **A Dataset Perspective on Offline Reinforcement Learning.** K. Schweighofer, M. Dinu, A. Radler, M. Hofmarcher, V. P. Patil et al. *CoLLAs*, 2022.
- **Reactive Exploration to Cope with Non-Stationarity in Lifelong Reinforcement Learning.** C. A. Steinparz, T. Schmied, F. Paischer, M.-C. Dinu, V. P. Patil, A. Bitto-Nemling et al. *CoLLAs*, 2022.
- **InfODist: Online Distillation with Informative Rewards Improves Generalization in Curriculum Learning.** R. Siripurapu, V. P. Patil, K. Schweighofer, M.-C. Dinu, T. Schmied et al. *Deep RL Workshop, NeurIPS*, 2022.
- **The Balancing Principle for Parameter Choice in Distance-Regularized Domain Adaptation.** W. Zellinger, N. Shepeleva, M.-C. Dinu, H. Eghbal-zadeh, H. D. Nguyen et al. *NeurIPS*, 2021.
- **XAI and Strategy Extraction via Reward Redistribution.** M.-C. Dinu, M. Hofmarcher, V. P. Patil, M. Dorfer, P. M. Blies, J. Brandstetter et al. *Intl. Workshop on Extending Explainable AI Beyond Deep Models*, 2020.

PATENTS

- **Agent-Based Neuro-Symbolic Methodologies for Scientific Discoveries and Workflow Automation.** M.-C. Dinu, C. Leoveanu-Condrei. US Patent Application Publication No. US 2026/0212154 A1; App. No. 19/032,587; filed 21 Jan 2025; published 23 Jul 2026; assignee ExtensityAI FlexCo (*patent pending*).

AWARDS & RECOGNITION

- **Sifted / Financial Times** — named among Europe’s top 11 AI startups to watch, 2024 (ExtensityAI).
- **ICLR 2023 Oral** — notable top 5% of accepted papers.
- **Ph.D. passed with distinction**, Johannes Kepler University Linz.
- **Author: *The Inflection Point*** (2025) — on the shift from “AI alchemy” to verifiable, architecture-driven AI systems.

SELECTED TALKS & MEDIA

- *Reinforcement Learning — Where We Are and What’s Next*. International RL Bootcamp, 2025.
- *The Science behind story.one’s Story Editor*. StoryOne — Human-Centric AI Publishing, 2025.
- *Extensity AI — The Tool That Automates Research Work*. startup.ro interview, 2025.
- Work referenced on Sabine Hossenfelder’s channel, *When Did LLMs Get This Good At Physics?*

COMMUNITY & SERVICE

- **AI Austria** — partner and organiser of the Reinforcement Learning community; co-organiser of the International RL Bootcamp (2nd edition, 2025).
- **Founder**, “AI Is All You Need” research community (450+ members); research outreach reaching 250,000+ views across X and LinkedIn.
- Open-source maintainer: SymbolicAI, Lighter (dependency injection for PyTorch).

CERTIFICATIONS & LANGUAGES

Udacity Nano Degree, Deep Learning Foundation (2017) | Oracle Certified Professional, Java SE 6 (2011) | Microsoft Certified IT Professional Administrator (2010)

Languages: German (native) | Romanian (native) | English (fluent, C1–C2)