




Denamganai Kevin

Researcher: Compositional Generalization, Physical & Formal Reasoning in AI

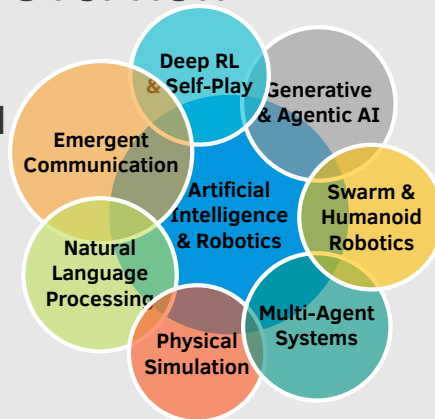
 kevidenamganai.netlify.app

 denamganai.kevin@gmail.com

 /kevin-denamganai-66131386/

 Near32

Overview



Programming

C • C++ • Python • \LaTeX

C# • Java • Mathematica

Language

French - Native

English - Bilingual (TOEFL 105)

Japanese - Professional Proficiency

Spanish - Conversant

German - Conversant

Research Summary

My research investigates the conditions under which AI systems acquire and deploy structured symbolic representations — beyond mere recombination of prelearned symbols — [towards in-context grounding of novel atomic symbols and their systematic recombination into unseen configurations](#). I formalised this capacity as **Compositional Learning Behaviours (CLBs)** — one of several **Symbolic Behaviours** that autonomous agents must exhibit to manipulate symbolic structures — operationalised them through the **Symbolic Behaviour Benchmark (S2B)** and its Meta-Referential Game framework, and demonstrated empirically that **CLB competency is a necessary but not sufficient condition for Olympiad-level formal mathematical verification**. This established a structural diagnostic gap in automated theorem proving literature. In parallel, I developed **EReLELA**, showing that emergent language abstractions can serve as structured exploration bonuses in sparse-reward RL, and the **Differentiable Language Model (DLM)** framework — a paradigm shift redefining Language Models as functions over sequences of probability distributions, enabling **end-to-end differentiability for any frozen LM** via Gumbel-Softmax gradient estimators. The unifying aim across these lines is to understand what must be true of AI systems for them to generalise reliably and co-operate efficiently with human beings — and to build the frameworks and tools that make that question answerable.

Projects

- **Symbolic Behaviour Benchmark (S2B)**: A benchmark to evaluate Compositional Learning Behaviours (CLBs) in AI agents via Meta-Referential Games — episodic tasks requiring in-context grounding and systematic recombination of novel symbolic structures [Denamganai et al., 2022].
- **S2B-LM**: An extension of the S2B **addressing the competence–performance distinction in Language Models (LMs)** by replacing numerical stimuli with categorical ones and introducing chain-of-thought scaffolding to elicit rather than merely probe latent CLB competency [Denamganai, 2026].
- **EReLELA**: A method to improve Exploration in Reinforcement Learning via Emergent Language Abstractions, leveraging referential game-trained representations as structured exploration bonuses in sparse-reward environments [Denamganai et al., 2024].
- **ReferentialGym**: Emergent Communication and Language Grounding via Referential Games, using PyTorch [Denamganai & Walker, 2020].
- **Archi**: Modular and reconfigurable building blocks for Deep Learning applications.
- **Regym**: A Single- and Multi-Agent / Self-Play Deep Reinforcement Learning framework.
- **DIPhyR-Gym**: An OpenAI Gym environment to investigate Deliberate & Intuitive Physics Reasoning in AI agents.
- **RelationalReasoning**: Deep Relational Reasoning algorithms, using PyTorch.
- **PyTorch_VAE**: Replication of many Disentangling β -VAE variants following [Higgins et al., 2018].
- **GazeboRL**: Deep Reinforcement Learning framework using ROS & Gazebo.
- **GazeboDomainRandom**: Domain Randomization tools following [Tobin et al. 2018], for object recognition.
- **Core**: Computer Algebra System in C/C++.
- **HaRo**: 3D printable MG995-based Raspberry Pi-powered humanoid robot.
- **EKF-DATMO**: Extended Kalman Filter-based solution to the DATMO problem.
- **SIMULATOR**: 3D Rigid-body Physics & Rendering Engine in C/C++ for video games.

Education & Research

2024 - 2026

Postdoctoral Research Associate in Machine Learning and Physical Simulation

University of Edinburgh, Scotland, UK

- **Research Ownership (DLM/DLMI):** Conceived and developed a framework for **Differentiable Language Models (DLM)**, enabling gradient-based navigation of a frozen LM's output space. Redefined LMs as functions operating over sequences of probability distributions (SDoTs) and implemented end-to-end differentiability via Gumbel-Softmax estimators. Extended the **Language Model Inversion (LMI) problem to the distributional regime** and developed DLMI algorithm improving on the LMI problem at scales an order of magnitude beyond the state of the art. This work provides a foundation for gradient-based probing, which underpins gradient-based prompt optimisation, adversarial auditing, and controllable generation applications. (*To be submitted to TMLR, 2026*).
- **Research Infrastructure (DIPhyR-Gym):** Designed and developed an OpenAI Gym environment to investigate Deliberate & Intuitive Physics Reasoning in AI agents. Environments randomise physical boundary conditions — initial velocities, inertial properties — over a base MuJoCo XML configuration at each reset, enabling systematic study of agent generalisation across varied physical regimes. Phase space data and configurations are logged for offline analysis.
- **Research Contributions, Supervision & Technical Collaboration:** Co-contributed to and provided ongoing technical supervision for two projects led by PhD researcher Sean Memery, owning experimental and technical deliverables within each:
 - **CueTip (SIGGRAPH 2025):** Contributed to an interactive, physics-aware pool coaching assistant. Owned the design and execution of the LM-as-expert-rule-evaluator experiments (Sections 4.2.2 & 5.2.2, Figures 3, 6–11), and authored the LM prompts & implementation details appendices (Appendices B & C). Provided technical consultancy on the simulator–language model interface, advising on the design of natural language event traces to enable closed-loop coupling between simulator and Language Model.
 - **xInv:** Contributed to a methodology for producing human-interpretable explanations of iterative optimisation processes. Owned the experimental design and execution of the LM discrimination evaluation (Section 3.4, Figures 4 & 6 and related text), the LM prompts & implementation details appendix, and provided ongoing technical input throughout the project.
- **Ongoing Research Agenda:** Initiated and co-developed a research direction investigating whether language models' physical reasoning abilities may be enhanced by framing physics-reasoning tasks through domain-specific languages (DSLs) — drawing on principles of Emergent Communication and referential game frameworks to discover physics-specific languages serving as structured, human-interpretable intermediaries between natural language and physical simulation. This direction is the subject of an active grant proposal in preparation.

2018 – 2024 (thesis submitted March 2024; revisions due August 2026; award pending)

PhD Researcher (Project Lead: Emergent Communication & Human-AI Cooperation)

University of York, England, UK (IGGI CDT)

- **Principal Research Project:** *Emergent Languages as a Tool for Artificial Thoughts*. Directed the investigation into (artificial) language emergence and the systematicity of neural players in Referential Games. Developed the first formal nomenclature for the field's deep learning resurgence, along with PyTorch-based framework and toolset to lower barrier of entry into the subfield.
- **Status Update:** *Thesis submitted March 2024; formal extension for revisions granted through August 2026 due to mitigating health circumstances (autistic burnout); research output and postdoctoral continuity maintained throughout.*
- **Research Independence:** Acted as the sole specialist in Emergent Communication within the IGGI cohort. Autonomously defined the research agenda and technical roadmap, providing the primary impulse for all core hypotheses and technical execution while leveraging supervisors for high-level academic strategy.
- **Awarded 4-year EPSRC Fellowship:** Successfully secured and managed competitive funding via the IGGI Centre for Doctoral Training (>£100k equivalent in stipend and research support).

- **Research Infrastructure:**

- Designed and developed **ReferentialGym**, a PyTorch-based framework for language grounding and emergence [Denamganaï & Walker, 2020];
- Co-lead-developed **Regym**, a modular Multi-Agent / Self-Play Deep RL framework;
- Designed and developed **Archi**, a library of modular and reconfigurable building blocks for Deep Learning applications, integrated with Regym and ReferentialGym to enable rapid prototyping of complex neural architectures.
- Designed and implemented the **Symbolic Behaviour Benchmark (S2B)**, a benchmark to evaluate Compositional Learning Behaviours (CLBs) in AI agents via Meta-Referential Games — episodic tasks requiring in-context grounding and systematic recombination of novel symbolic structures, addressing a core bottleneck in **zero-shot human-AI cooperation** [Denamganaï et al., 2022]. Later extended to **S2B-LM** to better probe CLB competency in Language Models by addressing their competence–performance distinction [Denamganaï, 2026].

- **Algorithmic Innovation:**

- **Exploration via Abstraction (EReLELA)**: Engineered a method to improve Exploration in Reinforcement Learning via Emergent Language Abstractions; pioneered the comparative analysis of emergent vs. natural language in terms of hierarchical abstraction.
- **Alignment & Efficiency (ETHER)**: Developed Emergent Textual Hindsight Experience Replay to extend the HER paradigm, significantly enhancing sample-efficiency and instruction-following capabilities in communicative RL agents.

- **Ongoing (Independent Research)**: Leveraged S2B-LM to evaluate **Compositional Learning Behaviours (CLBs) in state-of-the-art formal theorem-proving language models, establishing CLB competency as a necessary but not sufficient condition for Olympiad-level formal mathematical verification** [Denamganaï, 2026] ; submitted to the *3rd AI for Math Workshop @ ICML 2026*).

2013 - 2017

Engineering Degree, Computer Science and Systems

Ecole Nationale Supérieure de l'Electronique et de ses Applications, France

- **Final-Year Project**: Development of a **3D rigid-body physics and rendering engine, entitled SIMULATOR**, for a video game project in C/C++.
- Designing and building of a **3D-printable humanoid robot, entitled HaRo**, using MG995 servomotors and Raspberry Pi.
- Development of a **Computer Algebra System, entitled Core** for automatic differentiation of Neural Networks architectures in C/C++.

2015 - 2017

Research MSc., Artificial Intelligence and Robotics

Université de Cergy-Pontoise, France

- **Thesis**: Visual Contexts for a Spatial Recognition System in Wide Environments
- Reviewed biologically-inspired robotic vision and focused on a neuronal architecture aiming at solving the online spatial recognition problem.
- Investigating a coarse-to-fine filtering scheme making use of hebbian-weighted adaptation and parallels visual information pathways integrated in a cognitive map.

2016 - 2017

MEng., Electrical Engineering and Information Science (GPA: 3.7/4)

Osaka Prefecture University, Japan

- **Thesis:** Adaptability Features in a Nonlinear System-based Swarm of Robots
- Proposed two obstacle avoidance behaviours designed as nonlinear system-based controllers in order to shore up the bridge between nonlinear systems and swarm robotics, following our laboratory's previous works.
- Investigating the synthesis potential of nonlinear system-based controllers with a deep learning-based controller in a deep reinforcement learning framework.

Certifications

2025 Reinforcement Fine-Tuning LLMs with GRPO

- DeepLearning.AI
- 2023 **Quality & Safety for LLM Applications**

- DeepLearning.AI

2023 LangChain for LLM Application Development

- DeepLearning.AI

2021 Associate Fellowship of the HEA (AFHEA)

⇒ **Portfolio**

- York Learning And Teaching Award Course

2014 Autonomous Navigation for Flying Robots (AUTONAVx)

- Edx

2015 Underactuated Robotics (6.832x)

- Edx

2024 LLMOps

- DeepLearning.AI
- 2023 **Building & Evaluating Advanced RAG**

- DeepLearning.AI

2023 Machine Learning in Weather and Climate

- European Centre for Medium-Range Weather Forecasts

2017 Deep Learning Foundation Nanodegree

- Udacity

2015 Autonomous Mobile Robots (AMRx)

- Edx

2014 Computational Neuroscience

- Coursera

Scientific Leadership & Community Service

✉ 2021–2026 **Program Committee & Reviewing**

NeurIPS, ICLR, ICML, AAAI

- **Global Peer Review Authority:** Reviewed **50+ manuscripts** across the world's premier AI venues, ensuring the technical rigor of the field's most influential research.
- **Recognition & Leadership:**
 - Recipient of a **Gold Reviewer Award** at ICML 2026;
 - Co-recipient of the **Best Reviewer Award** at ICLR 2022's Emergent Communication Workshop;
 - Served as **Co-Organizer** for the ICLR 2022 Workshop on Emergent Communication;
- **Strategic Influence:** Invited member of the **NeurIPS Workshop Proposal Review Committee** (2021), contributing to the selection and shaping of the conference's specialized research tracks.
- **Reviewing Track Record (Condensed):**
 - **2026:** ICML (6), AAAI AI Alignment (2), ICLR (4)
 - **2025:** NeurIPS (3), AAAI (6), ICLR (4)
 - **2024:** AAAI (2), ICLR (1), ICML (5), NeurIPS (4)
 - **2023:** AAAI (3), ICML (5), NeurIPS (2)
 - **2021–2022:** ICLR Workshop (6), ICML (5), NeurIPS (4), NeurIPS Workshop Proposals (5)

Experience

✕ 2023 6-months PhD Research Internship

Sony Interactive Entertainment Europe

Research and development into using **Emergent Language Abstractions** for better **Exploration in Reinforcement Learning**.

✕ 2021-2022 6-months PhD Research Internship

Digital Creativity Labs

Collaboration with **Revolution Software** to research and develop **Style Transfer** approaches for video game assets creation, using 3D aware representation learning approaches with Generative Adversarial Networks (GANs) and Diffusion Models.

✕ 2020-2023 Graduate Teaching Assistant

University of York, Computer Science Dept.

- Mathematical Foundations of Computer Science (COM00013C)
- Formal Languages and Automata (COM00014C)
- Multi-Agent Interactions & Games (COM00009H)
- Intelligent Systems 1: Search & Representation (COM00020I)
- Computability and Complexity (COM00023I)
- Introduction to Software and Systems Engineering (COM00019I)

✕ 2016-Present Independent AI & Robotics Consultant & Engineer

- Development of a **Tool-Augmented Chatbot User Interface** to perform various task automation, powered by Open-Source Large Language Models (LLMs) and **Genetically-Evolved Prompts**, using **LangChain**
- Development of a **Tool-Augmented Chatbot User Interface**, using Retrieval-Augmented Generation (RAG) system on bespoke documents and responses to various API queries, powered by Open-Source Large Language Models (LLMs), **LangChain** and Google APIs
- Development of a **Q&A Chatbot User Interface on bespoke documents**, using Retrieval-Augmented Generation (RAG) system, powered by Open-Source Large Language Models (LLMs) and **LangChain**
- On-screen 2D gaze pose tracking system for hand-held devices with **PyTorch**.
- Domain randomization tools using **MakeHuman** and **Blender**.
- Semi-supervised GAN for car make and model classification using **TensorFlow**.
- 3D bot-human interface, using **Blender** and **Panda3D**'s Python API.
- Policy Neural Network for a Backgammon AI, using **TensorFlow**.
- Development of Navigation & Planning algorithms for a Roomba-like robot, using **ROS** & **Gazebo**.

Publications

- K. Denamganai, **On Compositional Learning Behaviours in Formal Mathematics**, preprint arXiv:2605.28512, 2026; submitted to the *3rd AI for Math Workshop @ ICML 2026*.
- K. Denamganai, and K. Subr, **Language Model Inversion through End-to-End Differentiation**, preprint arXiv:2602.11044; to be submitted to TMLR, 2026.
- S. Memery, K. Denamganai, A. Kapron-King, and K. Subr, **xInv: Explainable Optimization of Inverse Problems**, preprint arXiv:2506.11056, 2025.
- S. Memery, K. Denamganai, J. Zhang, Z. Tu, Y. Guo, and K. Subr, **CueTip: An Interactive and Explainable Physics-aware Pool Assistant**, SIGGRAPH2025, 2025.
- K. Denamganai, T. Bradley, P. Vito Amadori, S. Missaoui, G. Moss, and J. Walker, **EReLELA: Exploration in Reinforcement Learning via Emergent Language Abstractions**, rejected from NeurIPS2024 (ratings:3/4/4/7); OpenReview preprint, 2024; to be submitted to TMLR, 2026.

- K. Denamganaï, D. Hernandez, O. Vardal, S. Missaoui, and J. Walker, **ETHER: Aligning Emergent Communication for Hind-sight Experience Replay**, preprint arXiv:2307.15494 , 2023.
- K. Denamganaï, S. Missaoui, and J. Walker, **Visual Referential Games Further the Emergence of Disentangled Representations**, preprint arXiv:2304.14511, 2023.
- K. Denamganaï, S. Missaoui, and J. Walker. **Meta-Referential Games to Learn Compositional Learning Behaviours**, rejected from ICLR2025 (ratings:5/5/6/6), preprint arXiv:2207.08012, 2022; to be submitted to TMLR, 2026.
- K. Denamganaï and J. Walker, **On (Emergent) Systematic Generalisation and Compositionality in Visual Referential Games with Straight-Through Gumbel-Softmax Estimator**, in *4th NeurIPS Workshop on Emergent Communication*, preprint arXiv:2012.10776, 2020.
- K. Denamganaï and J. Walker, **ReferentialGym: A Nomenclature and Framework for Language Emergence & Grounding in (Visual) Referential Games**, in *4th NeurIPS Workshop on Emergent Communication*, preprint arXiv:2012.09486, 2020.
- D. Hernandez, K. Denamganaï, S. Devlin, S. Samothrakis and J. A. Walker, **A Comparison of Self-Play Algorithms Under a Generalized Framework**, in *IEEE Transactions on Games (ToG)*, doi: 10.1109/TG.2021. 3058898.
- D. Hernandez, K. Denamganaï, Y. Gao, P. York, S. Devlin, S. Samothrakis and J. Walker, **A Generalized Framework for Self-Play Training**, in *Proceedings of the 2019 IEEE Conference on Games (CoG)*, pp. 1-8, 2019.
- K. Denamganai, T. Nakamura, N. Hara and K. Konishi, **"Obstacle avoidance control law for two-wheeled mobile robots controlled by oscillators"**, in *Proceedings of the 61st Annual Conference of the Institute of Systems, Control and Information Engineers (ISCIE)*, 221-4, 2017.
- K. Denamganai, T. Nakamura, N. Hara and K. Konishi, **"Coupled Kuramoto oscillator-based control laws for both formation and obstacle avoidance control of two-wheeled mobile robots"**, *IEICE Technical Report*, NLP2017-44, pp. 87-91, 2017.