

Zehua Jiang

Brooklyn, NY | +1 (212) 300-7821 | zehua.jiang@nyu.edu
zehua-jiang.github.io | github.com/JiangZehua | Google Scholar

Research Focus

NYU Ph.D. researcher studying LLM agents and reinforcement learning for controllable and open-ended content generation. Research focuses on how tool use, execution harnesses, and environmental feedback enable agents to combine functional constraint satisfaction with free-form language control. Uses game environments to investigate generation quality, execution reliability, and generalization.

Education

New York University | Ph.D. in Computer Science and Engineering Expected Jan 2027
Advisor: Julian Togelius. Thesis: *Learning and Adapting Agents for Procedural Content Generation*.

New York University | M.S. in Electrical Engineering Jan 2021 - Sep 2022

Sun Yat-sen University | B.S. in Optoelectronic Information Science and Engineering 2016 - 2020
Exchange: University of California, Berkeley (Spring 2019).

Research Experience

Game Innovation Lab, New York University Brooklyn, NY
Ph.D. Researcher, Sep 2022 - Present; Research Assistant, Jan - Apr 2022

Agentic PCG: Tool-Using LLMs for Game Level Design 2025 - 2026

- Developed Agentic PCG, a tool-using LLM framework with an execution harness for feedback-driven planning and iterative self-refinement in procedural content generation.
- Combined evaluator-gated editing with free-form language instructions to satisfy functional constraints without task-specific fine-tuning; evaluated target alignment and structural coherence.
- Demonstrated greater evaluator-call efficiency than evolutionary search baselines under limited evaluation budgets.
- In a human study (n=17), participants selected the prompt-matched level in 92.4% of paired comparisons (133/144).

Local LLM Harnesses for Procedural Content Generation 2026

- Investigated local open-weight LLM failure modes: invalid or truncated tool calls, missed design targets, and edits that undo prior progress.
- Designed and evaluated a harness with bounded response recovery, evaluator feedback, selective edit acceptance, early stopping, and optional language-to-target extraction, without model-weight updates.
- Increased Qwen3.8-27B's success in satisfying all Zelda design constraints from 24/64 to 64/64 paired targets and achieved 99.03% (714/721) on a larger target set; validated harness effectiveness with Gemma models.

Small LLM Post-Training for Agentic PCG Ongoing

- Investigating SFT and GRPO for small LLMs using environment experience and edit traces, including trajectories from trained RL agents.

LLM-Guided Quality-Diversity Search for Reusable PCG Tools Ongoing

- Exploring quality-diversity (QD) search with LLM feedback in the fitness function to evolve reusable patches as generation tools.

PCGRL: Generalization and Scaling in Level Generation via Reinforcement Learning 2023 - 2025

- Investigated RL-agent generalization to out-of-distribution map sizes and shapes, examining local observations, randomized training maps, and fixed-tile constraints.
- Studied cooperative level generation using shared policies and rewards.
- Contributed to JAX/GPU PCGRL implementation and experiments; the framework achieved approximately 15x training speedup in the reported setup.

Learning Controllable 3D Level Generators Jan - Apr 2022

- Extended RL-based level generation from 2D to 3D Minecraft environments; implemented PPO with alternative action representations and policy architectures.
- Implemented MAP-Elites and CMA-ME to evolve neural cellular automata generators for diverse, playable levels.

Industry Experience

Webull Financial

Research Intern

St. Petersburg, FL

Jun 2025 - Aug 2025

- Developed an internal LLM agent for brokerage back-office journal management, supporting retrieval and investigation of entries affected by upstream and downstream order discrepancies.
- Implemented batch review and correction through custom tool use; handed the system to engineering for deployment on AWS.

Tencent Morefun Studio

Research Intern

Shenzhen, China

Jun 2024 - Aug 2024

- Worked closely with the team developing F.A.C.U.L., an LLM-driven FPS companion for cooperative gameplay.

Geek+

Algorithm Intern

Beijing, China

Jul 2020 - Jan 2021

- Designed dynamic task assignment using predicted robot arrival times, reducing workstation idle time by 27.78% on average versus a greedy baseline; field-tested in a warehouse that shipped 13.19 million orders.

Selected Publications

- [1] **Zehua Jiang**, Sam Earle, Ahmed Khalifa, Julian Togelius. Agentic PCG: Procedural Content Generation via Tool-using LLMs. Preprint, 2026; under review.
- [2] **Zehua Jiang**, Sam Earle, Ahmed Khalifa, Julian Togelius. Enabling Agentic PCG with Harness-Controlled Local Models. Preprint, 2026; under review.
- [3] In-Chang Baek, Sung-Hyun Kim, Sam Earle, **Zehua Jiang**, Jin-Ha Noh, Julian Togelius, Kyung-Joong Kim. PCGRLLM: Large Language Model-Driven Reward Design for Procedural Content Generation Reinforcement Learning. *IEEE Transactions on Games*, 2026.
- [4] Sam Earle, **Zehua Jiang**, Eugene Vinitsky, Julian Togelius. Video Game Level Design as a Multi-Agent Reinforcement Learning Problem. *AIIDE*, 2025.
- [5] Sam Earle, **Zehua Jiang**, Julian Togelius. Scaling, Control and Generalization in Reinforcement Learning Level Generators. *IEEE CoG*, 2024.
- [6] **Zehua Jiang**, Sam Earle, Michael Cerny Green, Julian Togelius. Learning Controllable 3D Level Generators. *FDG*, 2022.

Technical Skills

LLMs and Agents: tool/function calling, execution harnesses, MCP, Hugging Face, AWS Bedrock; SFT and GRPO

Machine Learning: PyTorch, JAX, TensorFlow; PPO, multi-agent RL; MAP-Elites, CMA-ME

Engineering: Python, C#, JavaScript, SQL, C++; Git, Linux, Slurm, AWS; Unity, Unreal Engine

Teaching and Service

NYU, Game Design (CS-GY 6553 / CS-UY 4553): Student Adjunct Instructor, Spring 2025 and 2026; Head TA, Fall 2024 and 2025; TA, Spring 2023.

NYU, Software Engineering I (CS-GY 6063): TA, Fall 2025.

Invited tutorial speaker, “Controllable Content Generators via Reinforcement Learning”, IEEE CoG 2023; co-organizer and session chair, Special Session on Games, IEEE WCCI 2024. Reviewer (50+ papers): IEEE Transactions on Games, Knowledge-Based Systems, GECCO, CoG, WCCI, AIIDE.