JASON ZHI LIANG, Ph.D.

San Francisco, CA

Contact: jasonzliang@utexas.edu | (510) 378-0709

LinkedIn: linkedin.com/in/jason-liang-3189b249

Scholar: scholar.google.com/citations?user=h8yYr7oAAAAJ

Github: github.com/jasonzliang

Research Summary

AI Research Scientist & Principal Architect dedicated to achieving Artificial General Intelligence (AGI) through the lens of nature-inspired computation and scalable reasoning. I specialize in the intersection of Evolutionary Optimization, Complexity Science, and Inference-Time Compute, moving beyond static scaling laws toward open-ended learning and recursive self-improvement.

My expertise lies in building agentic systems for deliberative System 2 reasoning and scaling neuroevolutionary paradigms for massive distributed optimization. I developed Caesar, an autonomous deep research agent that utilizes adversarial self-refinement and generative merging to achieve state-of-the-art creative synthesis. As the inventor of the CoDeepNEAT algorithm, I am eager to apply my experience in evolutionary, population-based optimization and architecture search to improve the next generation of AI agents and reasoning models.

Research Interests

- Novel evolutionary approaches to Artificial General Intelligence (AGI)
- Emergent intelligence and reasoning capabilities in LLMs and multi-agent systems
- Generalization across diverse tasks via meta-learning and open-ended optimization
- Neuroevolution applied to complex control and reinforcement learning problems
- Evolutionary adaptation and optimization of neural architectures
- Integration of subsymbolic and symbolic components in hybrid AI systems
- Scaling laws, inference-time compute, and emergent capabilities in foundation models

Technical Skills and Expertise

- **Core Research:** Inference-Time Compute Scaling, Agentic Reasoning (System 2), Associative Synthesis, Evolutionary Population-Based Optimization, Neural Architecture Search, Neuroevolution, Recursive Self-Improvement, Meta-Learning
- **Agent Infrastructure:** Graph-Augmented RAG, Stateful Agentic Memory, Multi-Agent Orchestration, FSM-Based Agent Coordination, Automated Evaluation (LLM-as-a-Judge), Evolutionary Prompt Optimization
- **Languages & Tooling:** Python (Advanced), Keras/TensorFlow, Frontier LLM APIs, Multi-Agent Frameworks (AutoGen/MetaGPT), Automated Web Parsing

Professional Experience

Cognizant AI Labs | Research Scientist & Principal Architect

San Francisco, CA | Dec 2018 – Present | Lead Researcher for Multi-Agent Architectures and Evolutionary Optimization.

1. Caesar: Autonomous Reasoning & Knowledge Synthesis

- **Automated R&D via Inference-Time Self-Improvement:** Architected "Caesar," an autonomous research agent similar to Gemini's "Deep Research" agent. Scaled test-time compute using an adversarial refinement loop that actively critiques internal drafts, generates orthogonal queries to target weaknesses, and consolidates findings via a generative merge to produce state-of-the-art creative artifacts.
- **Graph-Based Associative Reasoning:** Replaced static RAG retrieval with a dynamic "Perceive-Think-Act" loop backed by a self-organizing knowledge graph. By evaluating information against local graph topologies, the agent performs associative reasoning to discover non-obvious, cross-disciplinary connections—fundamentally overcoming the "navigational amnesia" of standard linear agents.
- **Open-Ended Exploration Policy:** Designed an active information foraging policy that breaks free from static scaling laws. The system continuously monitors memory and graph structures to detect stagnation, autonomously executing strategic backtracking to ensure open-ended learning and maximum information gain over long horizons.

2. Multi-Agent Systems & Large Scale Evolutionary Optimization

- **Multi-Agent Code Generation:** Developed a **Finite State Machine (FSM)** framework to coordinate teams of specialized LLM agents. Applied evolutionary algorithms to optimize the *composition* of these teams, achieving high performance on complex reasoning benchmarks like SciCode and HumanEval+.
- **Hierarchical Expert Agents:** Designed a dynamic expert-agent generation pipeline that decomposes hard coding problems and routes sub-tasks across specialized LLM roles, enabling scalable orchestration of multi-step reasoning.
- **Evolutionary Population-Based Training (EPBT):** Pioneered a population-based RL approach that dynamically evolves hyperparameters and loss functions during training. This method significantly reduces the compute required to tune large vision and language models compared to standard grid search.
- **Evolutionary Prompt Optimization:** Pioneered evolutionary techniques for LLM prompt optimization, improving performance on challenging code generation benchmarks.
- **Code Repository Analysis:** Created "Code Archaeologist," an agentic system capable of reasoning over repository histories (via Git/LFS) to detect architectural patterns and technical debt.
- **Scalable Model Training:** Scaled distributed ML systems to produce state-of-the-art neural network models for computer vision and language tasks.

Sentient Technologies | Research Scientist

San Francisco, CA | June 2017 – Nov 2018

- **Massive-Scale Neuroevolution:** Invented and scaled **CoDeepNEAT** (Evolutionary Neural Architecture Search) to run on clusters of hundreds of GPUs. Validated that evolutionary search scales with parallel compute resources, a critical property for pre-training large-scale ML systems.
- **Multi-Task Optimization:** Achieved state-of-the-art results on image captioning and object recognition by evolving neural topologies that share weights across tasks, proving that evolution can discover

efficient modular structures.
- **Production AutoML:** Designed and deployed automated machine learning systems targeted at real-world production use cases.

Sentient Technologies | Research Intern

San Francisco, CA | Dec 2015 – May 2017

- Researched early applications of evolutionary algorithms to deep neural networks, laying the groundwork for CoDeepNEAT.
- Helped build scalable ML infrastructure for neural network architecture optimization across distributed clusters.

The University of Texas at Austin | Research Assistant / Ph.D. Candidate

Austin, TX | Aug 2013 – Dec 2018

- **Solving Sparse Reward Problems:** Developed **MEA (Meta-Evolutionary Algorithm)**, a bilevel optimization framework. Successfully applied this to complex continuous control tasks (e.g., helicopter hovering), solving non-differentiable control problems where standard Gradient-based RL struggled to converge.
- **Neural Architecture Search Foundations:** Created the original **CoDeepNEAT** algorithm, combining evolutionary computation with hierarchical coevolution to automatically discover deep neural architectures.
- Published extensively in top venues including GECCO and Applied Soft Computing.

Open Source Robotics Foundation | Research Intern

Mountain View, CA | May 2015 – Aug 2015

- Contributed to the Gazebo 3D robotics simulator and authored a RoboCup plugin for 3D soccer simulation.

UC Berkeley | Research Assistant

Berkeley, CA | May 2012 – Aug 2013

- Developed a computer vision system for vision-based indoor localization from cellphone imagery, fused with LIDAR-derived point clouds.

Qualcomm | Intern

Berkeley, CA | Sep 2011 – Dec 2011

- Built benchmarking and performance measurement tools for AR applications (Layar SDK) on Android.

Research Projects

Multi-Agent LLM Systems (Current)

- Designed an **FSM-based framework** for orchestrating multi-agent software development workflows covering automated code generation, testing, and execution against benchmarks such as HumanEval+.
- Developed a hierarchical expert-team approach that decomposes and solves hard coding problems through specialized agent roles.

- Evaluated LLM reasoning performance on challenging benchmarks including SciCode.
- Explored evolutionary optimization of prompts and agent team composition for emergent collaboration.

Autonomous Knowledge Agents (Current)

- Built an autonomous web exploration agent ("Caesar") that constructs a graph-based knowledge base from its discoveries.
- Developed creative question/answer synthesis using iterative knowledge-graph querying and adversarial refinement.

Neural Architecture Search

- Developed **CoDeepNEAT**, a novel algorithm that automatically discovers optimal deep neural network architectures by evolving both macro and micro structure.
- Scaled these techniques to large distributed computing environments with hundreds of GPUs.

Evolutionary Bilevel Optimization

- Invented **MEA**, a meta-evolutionary bilevel optimization algorithm for complex control problems.
- Applied the method to challenging tasks including helicopter hovering and sparse-reward RL, optimizing both topology and weights simultaneously.

Population-Based Training

- Created **EPBT**, which evolves neural network hyperparameters and loss functions during training.
- Implemented scalable distributed infrastructure to run large-scale evolutionary optimization in parallel with gradient-based learning.

Awards & Honors

- **BEACON Research Grant** (2015) — Funded by the BEACON Center for the Study of Evolution in Action (NSF Science & Technology Center) for research on epigenetic mechanisms in neuroevolution and evolution of deep neural network architectures.

Professional Activities

- **Reviewer** — Genetic and Evolutionary Computation Conference (GECCO) | 2019 – Present

Education

Ph.D. in Computer Science (Neuroevolution & Deep Learning) *The University of Texas at Austin* | Sep 2013 – Dec 2018 *Advisor: Risto Miikkulainen* Dissertation: *Evolutionary Neural Architecture Search for Deep Learning*

B.S. in Electrical Engineering and Computer Science *University of California, Berkeley* | Sep 2009 – May 2013

Selected Publications & Patents

Full publication list available on Google Scholar

Featured Publications

Caesar: Deep Agentic Web Exploration for Creative Answer Synthesis (2026) *Jason Liang, Elliot Meyerson, Risto Miikkulainen*, Preprint, ResearchGate. Under review at ICML 2026.

- Autonomous R&D agent that achieves state-of-the-art knowledge synthesis by scaling inference-time compute through an adversarial refinement loop and a self-organizing knowledge graph.

Self-Transcendence: Achieving AGI via Chaotic Dynamics and Thermodynamic Attractors (2025) *Jason Liang*, TechRxiv Preprint.

- Position paper with framework for AGI based on non-equilibrium thermodynamics and self-organizing systems, moving beyond static objective functions.

Asynchronous Evolution of Deep Neural Network Architectures (2024) *Jason Liang, Hormoz Shahrzad, Risto Miikkulainen*, Applied Soft Computing, Vol. 152.

- Techniques for asynchronous population-based optimization at scale, essential for efficient training on massive distributed clusters.

Regularized Evolutionary Population-Based Training (2021) *Jason Liang, Santiago Gonzalez, Hormoz Shahrzad, Risto Miikkulainen*, GECCO 2021.

- Method to dynamically optimize loss functions during training, improving convergence speed and generalization.

Discovering Gated Recurrent Neural Network Architectures (2020) *Aditya Rawal, Jason Liang, Risto Miikkulainen*, Deep Neural Evolution.

- Evolutionary discovery of novel gated RNN cells that outperform hand-designed baselines like LSTM.

Evolutionary Neural AutoML for Deep Learning (2019) *Jason Liang, Elliot Meyerson, Babak Hodjat, et al.*, GECCO 2019.

- Evolutionary optimization of novel neural network architectures for multi-tasking learning.

Evolutionary Architecture Search for Deep Multitask Networks (2018) *Jason Liang, Elliot Meyerson, Risto Miikkulainen*, GECCO 2018.

- Novel evolutionary algorithm (CoDeepNEAT) that automates the discovery of deep neural architectures, outperforming human-designed baselines.

Evolving Deep Neural Networks (2018) *Risto Miikkulainen, Jason Liang, Elliot Meyerson, et al.*, Artificial Intelligence in the Age of Neural Networks and Brain Computing.

- Foundational text establishing the synergy between Evolutionary Strategies (ES) and Deep Learning as a viable path toward scalable AI.

Evolution Is the New Deep Learning (2018) *Risto Miikkulainen, Babak Hodjat, Xin Qiu, Jason Liang, Elliot Meyerson, Aditya Rawal, Hormoz Shahrzad*, ACM SIGEVOlution, Vol. 11, No. 1.

- Survey article articulating evolutionary computation as a scalable complement to gradient-based deep learning.

Evolutionary Bilevel Optimization for Complex Control Tasks (2015) *Jason Zhi Liang, Risto Miikkulainen*, GECCO 2015.

- Solving complex control tasks using evolutionary bilevel optimization, handling sparse reward signals where traditional RL failed.

Additional Publications

- **The Next Frontier in AI: Multi-Agent Systems** (2025) — *Babak Hodjat, Deepak Singh, Jason Liang, Xin Qiu*, NVIDIA GTC 2025 Conference Talk.
- **Training Stacked Denoising Autoencoders for Representation Learning** (2021) — *Jason Liang, Keith Kelly*, arXiv:2102.08012.
- **Evolutionary Neural Architecture Search for Deep Learning** (2018) — *Jason Zhi Liang*, Ph.D. Thesis, The University of Texas at Austin.
- **UT Austin Villa: RoboCup 2015 3D Simulation League Competition and Technical Challenges Champions** (2016) — *Patrick MacAlpine, Josiah Hanna, Jason Liang, Peter Stone*, RoboCup-2015: Robot Soccer World Cup XIX.
- **UT Austin Villa: A Machine Learning Approach for Kicking and Passing** (2015) — *Patrick MacAlpine, Josiah Hanna, Jason Liang, Peter Stone*, RoboCup 2015 Robot Soccer World Cup XIX.
- **Image-Based Positioning of Mobile Devices in Indoor Environments** (2015) — *Jason Zhi Liang, Nicholas Corso, Eric Turner, Avideh Zakhor*, Multimodal Location Estimation of Videos and Images.
- **Evolutionary Bilevel Optimization for Complex Control Problems and Blackbox Function Optimization** (2015) — *Jason Zhi Liang*, Masters Thesis, The University of Texas at Austin.
- **Image-Based Localization in Indoor Environments** (2013) — *Jason Zhi Liang, Nicholas Corso, Eric Turner, Avideh Zakhor*, 4th Int'l Conf. on Computing for Geospatial Research and Application.
- **Reduced-Complexity Data Acquisition System for Image-Based Localization in Indoor Environments** (2013) — *Jason Zhi Liang, Nicholas Corso, Eric Turner, Avideh Zakhor*, Int'l Conf. on Indoor Positioning and Indoor Navigation.

Patents

- **US Patent 11,507,844:** Asynchronous Evaluation Strategy for Evolution of Deep Neural Networks (2022).
- **US Patent 11,030,529:** Evolution of Architectures for Multitask Neural Networks (2021).
- **US Patent 11,003,994:** Evolutionary Architectures for Evolution of Deep Neural Networks (2021).
- **US Patent App. 17/389,961:** System and Method for Regularized Evolutionary Population-Based Training (2023).
- **US Patent App. 17/064,706:** Method and System for Sharing Meta-Learning Method(s) Among Multiple Private Data Sets (2022).
- **US Patent App. 16/671,274:** Multiobjective Coevolution of Deep Neural Network Architectures (2020).
- **US Patent App. 15/794,913:** Cooperative Evolution of Deep Neural Network Structures (2018).
- **US Patent App. 15/794,905:** Evolution of Deep Neural Network Structures (2018).