Zhizhuo Yang

AI Research Scientist · Ph.D. · Computing & Information Sciences

WRITER, 111 Maiden Ln, San Francisco, CA 94108

zhizhuo@writer.com | George614 | zhizhuo-yang | Google Scholar

Education

Rochester Institute of Technology *Rochester, NY, USA*

PhD Computing and Information Sciences *08/2017 - 12/2024*

- Dissertation: *Robust and Efficient Active Inference for Perception, Action, and Learning*
- Co-Advisor: Dr. Alexander G. Ororbia II
- Co-Advisor: Dr. Reynold Bailey

Beijing Institute of Technology *Haidian, Beijing, China*

MS Virtual Reality and Augmented Reality Technology *09/2014 - 06/2017*

- Advisor: Dr. Dongdong Weng

Beijing Jiaotong University *Haidian, Beijing, China*

BS Optical Information Science and Technology *09/2010 - 06/2014*

Professional Experience

02/2025 - present
AI Research Scientist at WRITER Owned LLM post-training end to end as the primary builder of WRITER's agentic RL stack on the Slime framework (data synthesis, verification and rewards, and integration with the production agent harness) for a 744B-parameter MoE model. Built the synthetic tool-use data pipeline whose data post-trained Palmyra x6, WRITER's shipped agentic model, and co-authored its technical report. Served as internship supervisor for research on efficient and self-correcting diffusion language models, with two papers under review. Earlier work covered self-evolving agents and parameter-efficient differential Transformers.

09/2023 - 12/2023
Part-time Student Researcher at Meta Continued the work from the previous internship. Gathered new specialized data and built a full stack machine learning pipeline to store, process, train, and evaluate our proposed model using cloud infrastructures such as notebooks and Linux clusters. Helped the team validate the research path with positive results (reduced mean tracking error by 51% and variance by 32%).

05/2023 - 09/2023
AI Research Scientist Intern at Meta Researched the end-to-end gradient-based EKF learning method for sensor fusion. Developed a neural network-based residual pose learning algorithm to improve tracking accuracy. Processed 1.35 TB of data using Python, C++ executable, and Bash scripts. Implemented the method in PyTorch and trained on cluster. Reduced mean tracking error by 39.76% and standard deviation by 12.68% respectively. Analyzed the raw data and results, generated plots and animation using pandas, scikit-learn, seaborn and plotly.

05/2020 - 08/2020
Research Intern at Facebook Developed and tested a real-time statistical filtering algorithm in C++ to improve Asynchronous TimeWarp technology used in Quest 2 by reducing visual artifacts. Learned to work with Android developer toolkit quickly and read about ARM NEON API. Improved team working efficiency by building efficiency tools in Python and writing informative Wiki pages in the company.

05/2019 - 08/2019
Software Engineering Intern at Facebook Developed a deep learning-based algorithm which extrapolates frames using adjacent frames and motion vectors to increase rendering frame rate and save GPU resources on VR headsets. Implemented it in Tensorflow and PyTorch and trained on clusters with Horovod and Slurm. Iterated on ideas and made rapid implementation by working closely with engineers and researchers from different teams.

Publications

Published

Nguyen, V. D., **Yang, Z.**, Buckley, C., Ororbia, A. 2025. SR-AIF: Solving Sparse-Reward Robotic Tasks from Pixels with Active Inference and World Models, *ICRA 2025*

Yang, Z., Diaz, G.J., Fajen, B.R., Bailey, R. and Ororbia, A.G., 2023. A neural active inference model of perceptual-motor learning. *Frontiers in Computational Neuroscience*, 17, p.1099593.

Kothari, R., **Yang, Z.**, Kanan, C., Bailey, R., Pelz, J. B., & Diaz, G. J. 2020. Gaze-in-wild: A dataset for studying eye and head coordination in everyday activities. *Scientific Reports*, 10(1),1-18.

Nair, N., Kothari, R., Chaudhary, A. K., **Yang, Z.**, Diaz, G. J., Pelz, J. B., & Bailey, R. J. 2020, September. RIT-Eyes: Rendering of near-eye images for eye-tracking applications. In *ACM Symposium on Applied Perception* 2020 (pp. 1-9).

Yang, Z., & Bailey, R. 2019, June. Towards a data-driven framework for realistic self-organized virtual humans: coordinated head and eye movements. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (pp. 1-3)

Kothari, R., **Yang, Z.**, Binaee, K., Bailey, R., Kanan, C., Pelz, J., & Diaz, G. 2018. Classification and Statistics of Gaze In World Events. *Journal of Vision*, 18(10), 376-376.

Yang, Z., & Weng, D. 2016, December. Passive haptics based MR system for geography teaching. In *Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry*-Volume 1 (pp. 23-29)

Yang, Z., Weng, D., Zhang, Z., Li, Y., & Liu, Y. 2016, September. Perceptual issues of a passive haptics feedback based MR system. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)* (pp. 310-317). *IEEE*.

Technical Reports

Du, P., Kamble, K., Vasudev, R., **Yang, Z.**, Nadimpally, R., Krishna, A., Alshikh, W., & Bikel, D. M. 2026. Palmyra x6 Technical Report: An Agentic, Tool-Use Model Post-Trained via Anchored Supervised Fine-Tuning. *arXiv preprint arXiv:2608.16620*.

Under Review

Denoising Blocks, Not Tokens: Efficient Compressed Continuous Diffusion with Branching Token Realization, *ICLR 2027*

When Are Errors Evidence? Deciding Whether to Revise, What to Keep, and How to Commit in Diffusion Language Models, *ACL ARR October 2026*

Teaching Experience

Fall 2022	**CSCI-635 Intro to Machine Learning**, Graduate Teaching Assistant	*Python*
Fall 2021	**CSCI-635 Intro to Machine Learning**, Graduate Teaching Assistant	*Python*
Spring 2020	**CSCI-620 Intro to Big Data**, Graduate Teaching Assistant	*JAVA, MongoDB*
Spring 2018	**CSCI-620 Intro to Big Data**, Graduate Teaching Assistant	*MySQL, R*

Mentoring

Summer 2026	**Xinsong Feng** 🎓, Summer intern, WRITER
Spring 2021	**Timothy Johnson**, Undergraduate independent study, Rochester Institute of Technology
2019-2020	**Nitinraj Rajan Nair**, Graduate capstone project, Rochester Institute of Technology
2015-2017	**Haiyan Jiang, Mo Su, Yue Li, Xingyao Yu and Yaqiong Xue**, Subsequent master students in the same research group, Beijing Institute of Technology

Research Experience

Efficient compressed continuous diffusion language models

Collaborators: Xinsong Feng, Peng Du, Daniel M. Bikel, Jiayun Wang, Haipeng Chen *06/2026-present*

- Served as internship supervisor and contributed to all technical discussions on method, experimental design, and analysis
- Investigated Branching Latent Diffusion (BLD), which compresses a 1024-token sequence into 64 block latents (a 16× reduction) and decodes each latent with a local autoregressive branch, running all branches in parallel
- Generated latents in groups, each conditioned on previously generated latents, to address the difficulty of jointly generating strongly compressed latents
- Against a similarly sized continuous diffusion baseline, BLD cut generation FLOPs by over 80× and raised throughput by over 6×; against an autoregressive baseline, it reached over 6× higher throughput and over 4× lower latency

Error correction and selective revision in diffusion language models

Collaborators: Xinsong Feng, Jinjin He, Peng Du, Daniel M. Bikel, Jiayun Wang *06/2026-present*

- Served as internship supervisor and contributed to all technical discussions on method, experimental design, and analysis
- Studied whether suspicious tokens from parallel decoding should be erased and regenerated or kept as context for revision, showing that the better choice depends on where the evidence for the intended answer resides
- Showed that erasing and regenerating suspicious tokens repairs 75% of incorrect ParallelBench drafts (vs. 39% when they stay visible) when the prompt determines the answer, while preserving tokens that carry essential evidence recovers 74% of intended answers (vs. none when erased)
- Showed that correction should be selective: correcting every output drops ParallelBench accuracy from 94.0% to 63.7%, while a learned judge reaches 96.6% with 9.9 forward passes, comparable to token-by-token decoding with 32 passes

LLM post-training and agentic RL

Collaborators: Peng Du, Zhenyu Zhao, Kiran Kamble, Arjun Krishna *09/2025-09/2026*

- Trained a 744B-parameter MoE agent end to end with RL (GRPO with LoRA and truncated importance sampling) on the Slime framework across 8-node H100/H200 clusters, using the production agent harness as the environment; an initial small-scale run raised held-out task reward by 0.077 over the base policy
- Profiled the RL loop (generation took 61% of wall-clock time) and ablated the SGLang rollout engines: FP8 weights fit one engine per node and gave 2.6-4.1× the output of the previous two-node layout, and speculative decoding with data-parallel-aware routing added 36-67% tokens/s at production load; one-step off-policy asynchronous training cut per-step non-generation time from 170s to 24s
- Found and fixed four silent defects (LoRA resume dropping the adapter, half-model LoRA exports under pipeline parallelism, a host RAM leak in the interconnect plugin, and rewards zeroed by a missing dependency); measured that 72% of GRPO groups had zero reward variance and consumed 72% of trajectory time, which redirected effort to data curation
- Built the synthetic tool-use data pipeline whose data post-trained Palmyra x6, WRITER's shipped agentic model, which achieved the top BFCL Core score (0.785) and the top six-benchmark mean among compared models
- Generated questions with checkable reference claims in four real tool environments (financial filings, clinical records, document workspaces, and Python/SQL analysis), filtered them with a 3-model answerability panel, and graded student rollouts claim by claim, keeping questions solved 1-3 times in 4 to target the right difficulty band
- Created a data curation pipeline with step-wise supervision using rejection sampling on the WRITER Agent platform
- Scaled multi-node mixed-precision SFT and preference fine-tuning on agentic datasets
- Built a verifier library (atomic verifiers, grouped verifiers, and agent-as-judge), rubric-, verifier-, and judge-grounded rewards for non- and semi-verifiable agentic tasks, and an evaluation framework spanning internal and external, agentic and non-agentic benchmarks; connected Slime to internal verifiers and external environments such as MCP Atlas

Self-evolving agentic AI systems

Collaborators: Melisa Russak, Arjun Krishna *04/2025-09/2025*

- Explored methods, including Monte Carlo tree search for planning and learning from the trajectories most likely to succeed, that let LLM agents learn from their own mistakes
- Implemented self-correction mechanisms to improve agent task performance
- Used LoRA fine-tuning and model merging for continual learning, raising accuracy in agentic workflows

More efficient differential Transformer with low-rank adapters

Collaborators: Melisa Russak, Kiran Kamble, Dr. Krzysztof Święcicki, Maciej Pióro *02/2025-08/2025*

- Proposed a cumulative relevance score and used it, together with sparsity scores, to evaluate vanilla, differential, and improved differential Transformers across datasets
- Examined how low-rank adapters improve the training and parameter efficiency of the differential Transformer, targeting long-context performance, lower memory use, and lower training cost
- Integrated hyper-attention into the improved differential Transformer and benchmarked it against FlashAttention-2

Robust and efficient active inference

Advisors: Dr. Alexander G. Ororbia II and Dr. Reynold Bailey *11/2023-12/2024*

- Proposed energy-based sampling methods to improve sample efficiency and learning stability
- Developed contrastive prior preference learning method using current state space model with colleagues
- Our approach empirically showed better convergence speed and stability compared to state-of-the-art model based reinforcement learning algorithm on robotic tasks
- Proposed asynchronous inference-based sampling strategy for active inference

Memory augmented active inference

Advisors: Dr. Alexander G. Ororbia II and Dr. Reynold Bailey *09/2023-03/2024*

- Researched self-supervised learning methods such as contrastive learning and Siamese networks to learn better visual representations for an agent to solve tasks without the likelihood approximator
- Investigated efficient embedding methods related to Hopfield networks and energy transformers to do associative learning in the latent space for constructing and retrieving past experience from memory
- Developed unsupervised methods for clustering learned representations in the memory
- Prototyped and tested a continual learning agent on benchmarks including Minigrid and MetaWorld

Continual reinforcement learning

Advisors: Dr. Alexander G. Ororbia II and Dr. Reynold Bailey *06/2022-05/2023*

- Researched techniques such as Subjective Timescale Model and episodic memory related methods to construct better representation of past experience
- Investigated efficient embedding methods and unsupervised learning to develop better exploration techniques as well as boost transfer abilities of continual learning agents
- Prototyped and tested a continual learning agent on the Jelly Bean World benchmark

Empirical comparison and evaluation between deep active inference and deep reinforcement learning on cognitive tasks

Advisors: Dr. Alexander G. Ororbia II and Dr. Reynold Bailey *02/2021-05/2022*

- Developed an Active Inference (AIF) model that approximates the Expected Free Energy using Q-learning
- Compared the performance of variations of prior preference such as global prior, local prior and empirical prior learned from demonstration
- Systematically investigated and demonstrated the effects of components of our AIF framework on agent's performance
- Compared variants of our AIF framework with a popular reinforcement learning approach, i.e., Deep Q Network, on OpenAI gym environments
- Extended the benchmark tasks from OpenAI gym to more advanced benchmarks such as Animal-AI which focus on the cognitive capabilities of the intelligent agents

Generalizing active inference to continuous state-spaces: an application to perceptual-motor learning tasks

Advisors: Dr. Alexander G. Ororbia II and Dr. Reynold Bailey *09/2020-08/2022*

- Expanded a simple AIF model under discrete state-space by Buckley and made visualizations of the agent's behavior to teach myself and our research group about Active Inference
- Scaled AIF using deep neural networks and variational inference to tackle continuous state-space problems including classic control problems in OpenAI gym
- Devised hypothesis tests and experiments for investigating human perception-motor task with deep AIF framework based on an interception task
- Proposed a novel prior function that maps original observations to a new low-dimensional and interpretable space where prior preference can be specified as a distribution

Predictive processing framework for human perception & action

Advisors: Dr. Alexander G. Ororbia II and Dr. Reynold Bailey *02/2020-06/2020*

- Attempted modeling human behaviors in visuomotor tasks with predictive coding methods to answer cognitive questions on human perception, action and decision making. Specifically, investigated a virtual ball-catching experiment with eye and head tracking data using pandas
- Proposed hypothesis tests for information sufficiency of human's catching behavior under different frame of references
- Visualized the human subject data from previously a published study in a virtual environment using Unity3D engine
- Implemented a particular flavor of predictive coding framework called PC/BC-DIM by Spratling with TensorFlow2 and tested on the MNIST dataset

Rendering of near-eye images for eye-tracking applications

Advisor: Dr. Reynold Bailey *09/2019-03/2020*

- Collaborated with a cross-laboratory team to build a computer graphics system using blender for synthesizing a photo-realistic eye-tracking video dataset, which is for training image-based eye tracking algorithms
- Evaluated and analyzed trained machine learning models

Data-driven framework for realistic self-organized virtual avatars

Advisors: Dr. Reynold Bailey and Dr. Alexander G. Ororbia II *03/2019-12/2019*

- Led a team to develop algorithms including Mixture Density Networks, Recurrent Neural Networks and Mode-Adaptive Neural Networks in PyTorch for computationally modeling coordinated human eye-head movements from synchronized motion capture and eye tracking data in order to generate realistic virtual avatar animation
- Proposed a research pathway to move from script-based eye and head movement animation to self-organized and task-driven animation using reinforcement learning
- Mentored a master student in computer science on conducting research and experiments

Event-level error metric for evaluating gaze events classification

Advisor: Dr. Reynold Bailey *06/2018-12/2018*

- Devised a window-based event-level error metric which reports categorical misclassifications (by confusion matrix) and temporal misalignment (with timing offsets and overlap ratio) of gaze events. This metric is generally applicable to evaluate classification performance for any machine learning methods at event-level
- Implemented proposed methods in both Matlab and Python and compared to other existing metrics

Event detection of human eye and head movements

Advisor: Dr. Reynold Bailey *01/2018-05/2018*

- Collected data with a collaborator using Pupil Lab's eye tracker and an IMU while participants performed daily tasks such as walking, visual searching and ball-catching
- Worked with a collaborator on designing and training the gaze event classifier with Recurrent Neural Networks in PyTorch

Enhancing password recollection performance using augmented reality with the method of Loci

Advisor: Dr. Reynold Bailey *10/2017-03/2018*

- Used Microsoft HoloLens to assist the famous mnemonics, the method of loci, to help user with password recollection. Voice commands and spatial mapping features of the headset were used in the project to record the 3D positions of the digits placed into the environment
- Designed experimental setup and conducted data collection. Analyzed performance of human subjects under 3 conditions with ANOVA

Gaze-In-World (GIW) event classification

Advisor: Dr. Reynold Bailey *08/2017-12/2017*

- Trained a parallel bidirectional (Forward-Backward Window) LSTM neural network model using PyTorch with eye-tracking data such as gaze velocity, head movement data as input for GIW classification
- Developed a CNN-LSTM neural model for GIW event classification with a collaborator. Eye images and head velocity data served as input data. A CNN model was used to extract features directly from eye images, which were then fed into the LSTM model along with head velocity data

Room-scale MR system based on passive haptics

Advisor: Dr. Dongdong Weng *10/2016-03/2017*

- Built a room-scale Mixed Reality system using Unity3D engine with passive haptic feedback from real objects. Adopted PhaseSpace active optical tracking system as the tracking subsystem for objects in the scene and the user's head. Used Perception Neuron from Noitom as tracking device for user's arms and hands

Passive haptics-based MR system for geography teaching

Advisor: Dr. Dongdong Weng *07/2016-10/2016*

- Validated the feasibility of applying MR system in geography teaching based on passive haptics with user studies
- Integrated key SDKs, including Vuforia, Leap Motion, Oculus in Unity3D, calibrated tracking camera and LeapMotion camera, established hardware environment and programmed user study software based on passive haptics
- Recruited and selected human subjects, designed and conducted 3 experiments under different learning conditions, collected questionnaire data and recorded completion time and quiz accuracy
- Analyzed experiment results with a Friedman Test

Evaluating vehicle ergonomics with driving simulation in Virtual Reality

Advisor: Dr. Dongdong Weng *07/2015-12/2015*

- Communicated with clients to analyze their requirements, designed software architecture, drew class diagrams, wrote shooting scripts
- Measured actual vehicle space model on site and built 3D modeling with an artist
- Built the virtual scene and implemented event scripts with Unity3D
- Integrated PhaseSpace optical tracking device, head mounted display, joysticks and other hardware on site, debugged the network environment and tested software

Research on motion sickness in Virtual Reality systems for automobile platform

Advisor: Dr. Dongdong Weng *04/2015-06/2015*

- Measured and recorded real-time automobile rotation data using MTx inertia measuring unit from Xsens. Collected real-time velocity through vehicle data recorder and transmitted the information to virtual reality system on a laptop
- Built a virtual environment using Unity3D, debugged and integrated with hardware
- Presented the subjects with same or different motion cues in the VR system as in the real world, while collecting statistics of subjects' dizziness through electrocardio data. Subjective questionnaires were also administered after the test

Passive haptics-based Virtual Reality system for assembly tasks

Advisor: Dr. Dongdong Weng *02/2015-03/2015*

- Conducted stereo camera calibration between Leap Motion's infrared camera and an RGB tracking camera. Measured position relations of different hardware systems and subject's pupil distance to set the relations in software system accurately
- Designed and completed experiments to verify the level of consistency between vision and tactile sensory information provided by the system after calibration
- Designed a virtual assembly application and verified usability of proposed system

Outreach & Professional Development

Service and Outreach

2018	**National Technical Institute for the Deaf at RIT**, Note taker
09/2015	**International Symposium on virtual reality, augmented reality & visual computing application technology**, Volunteer
08/2011	**Beijing Northking Electronic Technology Development Co, Ltd**, Intern
2010 - 2011	**Undergraduate computer science association**, Committee member
2010 - 2011	**College student union**, Committee member
2010 - 2012	**College biking club**, Committee member

Development

1st International Workshop on Active Inference The first IWAI discusses current trends, novel results, real-world applications, to what extent active inference can be used in modern machine learning settings, such as deep learning, and how it can be unified with the latest psychological and neurological insights. I learned a lot about how people from different field perceive, develop and apply Active Inference, as well as richer mathematical techniques and research contexts beyond computing. I consider it to be professional development because of the inspiration and techniques I gained from this workshop given its multidisciplinary nature.

Peer Review

I reviewed for Conference on Lifelong Learning Agents (CoLLAs) 2022.

I reviewed for Journal Neural Networks by Elsevier 2023.

I reviewed for AAAI 2026.

I reviewed for Springer Nature.

Certifications

April 2024	**Quantization Fundamentals**, DeepLearning.AI & HuggingFace
April 2024	**How Diffusion Models Work**, DeepLearning.AI
March 2024	**Learning Kubernetes**, LinkedIn
January 2024	**Generative AI with Large Language Models**, DeepLearning.AI, Amazon Web Services
January 2024	**MLOps Essentials: Model Development and Integration**, LinkedIn
March 2017	**Machine Learning**, Coursera

Coursework

Rochester Institute of Technology *09/2017-05/2019*

- Deep learning for vision
- Cyberinfrastructure
- Statistical machine learning and data mining
- Bio-inspired intelligent systems
- Foundation of software engineering
- Quantitative Foundation
- Foundation of computer graphics

Technical Skills

Research areas LLM post-training, preference fine-tuning, agentic RL, RL systems and inference optimization, agentic data synthesis, agent evaluation, distributed training, diffusion language models, continual learning, active inference, world models, multimodal perception, sensor fusion, VR/AR systems

ML frameworks & infrastructure PyTorch, Megatron-LM, SGLang, Slime, JAX, TensorFlow, Keras, scikit-learn, pandas, Horovod, SLURM, AWS SageMaker

Other libraries & tools OpenCV, OpenGL, Plotly, SQL, MySQL, Git, Bash, LaTeX, Cursor, Claude Code

XR hardware & software Unity3D, Microsoft HoloLens, Oculus Rift, Intel RealSense, Leap Motion, Noitom Axis Neuron, PhaseSpace, Xsens MT Manager

Programming languages Python, C++; also MATLAB, Julia, R, C#, Java