

Jacob Dineen

AI, Reasoning & Cognition (ARC) Lab — ARIZONA STATE UNIVERSITY
jdineen@asu.edu | (480) 603-6994 | jacobdineen.com

RESEARCH INTERESTS

I work on reasoning and alignment in large language models (LLMs), with complementary interests in multi-agent reinforcement learning, controllability, and explainable AI (XAI).

EDUCATION

ARIZONA STATE UNIVERSITY — Ph.D. Artificial Intelligence [†] GPA: 4.00/4.00	2022–2027
Advisor: Ben Zhou [‡] ; Committee: Muhao Chen [§] , Chitta Baral [§] , Vivek Gupta [§]	
UNIVERSITY OF VIRGINIA — M.Sc. Computer Science GPA: 3.96/4.00	2019–2021
Advisor: Madhav Marathe [‡]	
SYRACUSE UNIVERSITY — M.S. Data Science GPA: 4.00/4.00	2017–2018
GRAND CANYON UNIVERSITY — B.S. Finance & Economics GPA: 3.65/4.00	2012–2015

[†] Expected completion. [‡] Advisor. [§] Committee member.

RESEARCH EXPERIENCE

Graduate Research Assistant, ARC Lab (ARIZONA STATE UNIVERSITY)	2024–Present
<ul style="list-style-type: none"> Research on reasoning, alignment, controllability, and multi-agent RL in LLMs. 	
Graduate Research Assistant, SEFCOM (ARIZONA STATE UNIVERSITY)	2022–2024
<ul style="list-style-type: none"> Research at the intersection of AI and cybersecurity. 	
Applied Research, CAPITAL ONE — Center for Machine Learning (C4ML)	2020–2021
<ul style="list-style-type: none"> Built reinforcement learning and agent-based simulations for organizational dynamics. 	
Research Assistant, BIOCOMPLEXITY INSTITUTE (UNIVERSITY OF VIRGINIA)	2019–2020
<ul style="list-style-type: none"> Worked on graph dynamical systems, cooperative game theory, and behavioral modeling. 	

PROFESSIONAL EXPERIENCE

Research Engineering Intern, PARETO AI	2025–2025
Machine Learning Engineer, SPRING OAKS CAPITAL	2022–2025
Data Scientist, CAPITAL ONE	2021–2022
Ph.D. Internships (3×), CAPITAL ONE	2020–2021
Analyst & Business Intelligence, REAL WORLD MARKETING	2016–2019
Data Scientist, BUFFALO CHECK LLC	2015–2019
Optimization Analyst, VOLTARI	2012–2015

PUBLICATIONS

Google Scholar Profile

* Equal Contribution, + Corresponding Author / Mentor

Peer-Reviewed Conference Proceedings (C)

- C1. **Dineen, Jacob⁺**, Rrv, Aswin, Liu, Qin, Xu, Zhikun, Ye, Xiao, Shen, Ming, Li, Zhaonan, Lu, Shijie, Baral, Chitta, Chen, Muhao, & Zhou, Ben (2025). *QA-LIGN: Aligning LLMs through Constitutionally Decomposed QA*. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 20619–20642.
- C2. RRV, Aswin, **Dineen, Jacob**, Handa, Divij, Uddin, Md Nayem, Parmar, Mihir, Baral, Chitta, & Zhou, Ben (2025). *ThinkTuning: Instilling Cognitive Reflections without Distillation*. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 31236–31250.
- C3. Xu, Zhikun, Shen, Ming, **Dineen, Jacob**, Li, Zhaonan, Ye, Xiao, Lu, Shijie, RRV, Aswin, Baral, Chitta, & Zhou, Ben (2025). *ToW: Thoughts of Words Improve Reasoning in Large Language Models*. In *Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 3057–3075.
- C4. Srinivasan, Adarsh, **Dineen, Jacob⁺**, Afzal, Muhammad Umar, Sarfraz, Muhammad Uzair, Riaz, Irbaz Bin, Zhou, Ben (2026). *RECAP: Transparent Inference-Time Emotion Alignment for Medical Dialogue Systems*. In *Proceedings of the 8th Clinical Natural Language Processing Workshop (Oral)*.
- C5. **Dineen, Jacob⁺**, Kridel, Donald J., Dolk, Daniel R., & Castillo, David G. (2023). *Unified Explanations in Machine Learning Models: A Perturbation Approach*. In *Proceedings of the 56th Hawaii International Conference on System Sciences (HICSS-56)*, pp. 795–804.
- C6. **Dineen, Jacob⁺**, Haque, A. S. M. Ahsan-Ul, & Bielskas, Matthew (2021a). *Formal Methods for an Iterated Volunteer's Dilemma*. In *Proceedings of the 14th International Conference on Social, Cultural, and Behavioral Modeling (SBP-BRiMS 2021)*, pp. 81–90.
- C7. **Dineen, Jacob⁺**, Haque, A. S. M. Ahsan-Ul, & Bielskas, Matthew (2021b). *Reinforcement Learning for Data Poisoning on Graph Neural Networks*. In *Proceedings of the 14th International Conference on Social, Cultural, and Behavioral Modeling (SBP-BRiMS 2021)*, pp. 141–150. .
- C8. Dolk, Daniel R., Kridel, Donald J., **Dineen, Jacob⁺**, & Castillo, David G. (2020). *Model Interpretation and Explainability towards Creating Transparency in Prediction Models*. In *Proceedings of the 53rd Hawaii International Conference on System Sciences (HICSS-53)*.

Working Papers / Under Review (W)

- W1. Li, Zhaonan, Lu, Shijie, Wang, Fei, **Dineen, Jacob**, Ye, Xiao, Xu, Zhikun, Liu, Siyi, Cho, Young Min, Li, Bangzheng, Chang, Daniel, Nguyen, Kenny, Yang, Qizheng, Chen, Muhao, Zhou, Ben (2025). *Unbiased Visual Reasoning with Controlled Visual Inputs*. Pending COLM 2026.
- W2. Liu, Qin, **Dineen, Jacob**, Huang, Yuxi, Zhang, Sheng, Poon, Hoifung, Zhou, Ben, Chen, Muhao (2025). *ArenaBench: Automatic Benchmark Evolution via Multi-Model Competitive Evaluation*. Pending COLM 2026.

- W3. Ye, Xiao, Li, Zhaonan, **Dineen, Jacob**, Xu, Zhikun, Lu, Shijie, Shen, Ming, Shrivastava, Shaswat, Ahuja, Avneet, Zhou, Ben (2025).
CC-LEARN: Cohort-Based Consistency Learning. Pending COLM 2026.
- W4. Shen, Ming, Xu, Zhikun, **Dineen, Jacob**, Ye, Xiao, Zhou, Ben (2025). *Training Language Models with Context-Free Reasoning Bottlenecks*. Pending COLM 2026.
- W5. RRV, Aswin, **Dineen, Jacob**, Handa, Divij, Parmar, Mihir, Zhou, Ben, Baral, Chitta, Mishra, Swaroop (2026). *Mid-training with Self-Generated Data Improves Reinforcement Learning in Language Models*. Pending ICML 2026.
- W6. **Dineen, Jacob**, RRV, Aswin, Xu, Zhikun, Zhou, Ben(2026). *Vocabulary Dropout for Curriculum Diversity in LLM Co-Evolution*. Pending COLM 2026.
- W7. Xu, Zhikun, Feng, Yu, **Dineen, Jacob**, Zhou, Ben (2026). *Skill Reuse as Compression in Agentic RL*. Pending COLM 2026.
- W8. Li, Zhaonan, Chickering, Kyle R., Li, Bangzheng, **Dineen, Jacob**, Ye, Xiao, Xu, Zhikun, Lu, Shijie, Huang, Yuxi, Shen, Ming, Nguyen, Bach, Pavuluri, Jaya Adithya, Nguyen, Mau Son, Chavan, Sanika, Le, Ngoc Minh Thu, Chen, Muhao, Zhou, Ben (2025).
VisAnalog: A Diagnostic Suite for Visual Concept Transfer on Natural Images. Pending VisCon 2026.
- W9. Ye, Xiao, **Dineen, Jacob**, Li, Zhaonan, Xu, Zhikun, Chen, Weiyu, Lu, Shijie, Huang, Yuxi, Shen, Ming, Tran, Phu, Yum, Ji-Eun Irene, Khan, Muhammad Ali, Afzal, Muhammad Umar, Riaz, Irbaz Bin, Zhou, Ben (2025). *Evaluating Medical LLMs by Levels of Autonomy: A Survey*.

SKILLS

OS: Linux (Ubuntu), macOS, Windows

Languages: Python, Rust, C, C++, Java, JavaScript, R, Bash, PRISM, x86-64

Databases: MySQL, PostgreSQL, MongoDB, Snowflake, Redis, SQLite, Redshift

Markup: \LaTeX , HTML

ML Libraries: PyTorch, TensorFlow, Keras, JAX, Numpy, Pandas, Polars, Dask, HuggingFace, TRL, vLLM, VeRL, PySpark, NetworkX, DGL, Torch Geometric, BoTorch, SnowparkML

Tools: Git, VSCode, Docker, Kubernetes, Helm, Airflow, AWS (ECR/S3/EKS/CodeBuild), Databricks, OR-Tools, Sigma, Streamlit

SERVICE

Conference Reviewer: HICSS, SBP-BRiMS, NAACL, EMNLP, EACL

Teaching / Tutoring:

- CGCC Calculus & Linear Algebra Tutor (2019)
- Teaching Assistant — ASU CSE 365 / pwn.college (Security)

MISCELLANEOUS

Expert AI Trainer, PARETO AI (2024–Present)

Independent Contractor, MERCOR INTELLIGENCE (2025–Present)

Cybersecurity: Pwn.college green belt (binary exploitation) — pwn.college, user: jdin

GRADUATE COURSEWORK

ARIZONA STATE UNIVERSITY — PhD Computer Science:

Knowledge Representation; Computer Systems Security; Software Security; Planning and Learning Methods in AI; Algorithms; Research/Dissertation Hours.

UNIVERSITY OF VIRGINIA — MSc Computer Science:

Algorithms; Machine Learning; Computer Vision; Formal Methods; Reinforcement Learning; Graph Mining; Learning Theory (Game Theory); Cloud Computing; Research Hours.

SYRACUSE UNIVERSITY — MS Data Science:

Data Analysis and Decision Making; Business Analytics; Financial Analytics; Marketing Analytics; Advanced Information Systems; Data Science; Data Warehousing; Text Mining; Scripting for Data Analysis; Information Policy.