Alvaro Velasquez, PhD

Curriculum Vitae

+1 (407) 399 6245
alvaro.velasquez@colorado.edu
Security Clearance Level: TS/SCI

Summary

My research focus is on developing data-efficient and generalizable foundation models through neuro-symbolic AI. At DARPA, I have secured and led $200 million in programs on efficient neurosymbolic AI, robust autonomy, and foundation models for scientific discovery. My research has received best paper awards from AAAI (0.1% success rate), the IEEE Computational Intelligence Society, and AFRL. I have authored over 100 peer-reviewed papers, 3 patents, and my work has been featured in Forbes and the Washington Times. I also serve as co-founder of the International Neuro-symbolic Systems (NeuS) conference and associate editor for IEEE Transactions on AI. My research has been funded by DARPA, the Air Force Office of Scientific Research, and Lockheed Martin.

Professional Experience

08/2025–present **Chief Executive Officer**, *Neurosymbolic Intelligence*
We are creating the first foundation model for logistics.

08/2022–present **Assistant Professor**, *University of Colorado Boulder*
Department of Computer Science

08/2022–08/2025 **Program Manager**, *Defense Advanced Research Projects Agency (DARPA)*
As PM for the Assured Neuro-Symbolic Learning Reasoning (ANSR), Transfer from Imprecise and the Abstract Models to Autonomous Technologies (TIAMAT), and Foundation Models for Scientific Discovery (FoundSci) programs, I have secured and lead over $200 million in fundamental and applied research investments. This entails close collaboration with academia and the private sector, as well as the oversight of 67 performer teams. These programs led to eight startups and technology transitions to various government agencies.

06/2018–07/2022 **AI Technical Lead**, *Air Force Research Laboratory*
Managed the $30+ million Machine Intelligence portfolio of investments for the Information Directorate of the Air Force Research Laboratory (AFRL/RI), provided technical oversight and guidance to existing programs, and guided the formation of new research programs.

Summer 2015-Spring 2018 **Graduate Research Fellow**, *National Science Foundation*
Formal Methods Based Algorithmic Synthesis of Emerging Computing Architectures.
NSF Field of Study: Formal Methods, Verification, and Programming Languages.

Summers 2014-2017 **Research Intern**, *Air Force Research Laboratory*
Trusted Systems Branch; Autonomy, Command and Control Branch

2014–2015 **Graduate Research Assistant**, *University of Central Florida*
Department of Computer Science

Fall 2014 **Consultant**, *Griffiss Institute*

Summer 2013 **Undergraduate Researcher**, *University of Central Florida*
Center for Research in Computer Vision

2010-2013 **Software engineer intern, tutor, gas station clerk, retail clerk, cook**

Awards

2025 IEEE Computational Intelligence Society (CIS) Spotlight Paper Award.

2025 IEEE Conference on Ubiquitous Intelligence (UIC) Best Paper Award.

2025 University of Central Florida Distinguished Alumnus Award.

2025 DARPA Meritorious Service Medal.

2023 AAAI Distinguished Paper Award.

2021 AFRL Fred L. Diamond (Best Paper) Award.

2021 AFRL Best Patent Award 2021.

2020 AFRL Oliver G. Tallman award for "contributions and accomplishments that have had a significant impact and enhanced the credibility of AFRL"

2019 University of Central Florida 30 Under 30 Award.

2015-2018 National Science Foundation Graduate Research Fellowship (NSF GRFP).

2015 University of Central Florida Dean's Fellowship.

Education

05/2018 **PhD in Computer Science**, *University of Central Florida*

05/2016 **Master's in Computer Science**, *University of Central Florida*
Canadian Summer School on Quantum Information, Fields Institute.

05/2014 **Bachelor's in Computer Science (with Honors)**, *University of Central Florida*
Honors Thesis (**Best Thesis Award**).
Minor in Mathematics.
Spring School on Monte Carlo Methods in Artificial Intelligence, Oregon State University.

Books

2024 [3] **Neuro-Symbolic Artificial Intelligence: Foundations and Applications.**
Alvaro Velasquez, *Shankar Sastry, Pradeep Ravikumar, and Houbing Song*
Book proposal accepted by Wiley-IEEE Press

[2] **Artificial Intelligence: From Simulation to Reality.**
Alvaro Velasquez, *Vishal Patel, Antonio Loquercio, and Alan Fern*
Book proposal accepted by Wiley-IEEE Press

[1] **AI for Cybersecurity: Research and Practice.**
Houbing Song, Elisa Bertino, ***Alvaro Velasquez****, Huihui Wang, Yan Shoshitaishvili, and Sumit Kumar Jha*
Book proposal accepted by Wiley-IEEE Press

Selected Publications (Or Accepted for Publication)

2026 [124] **On the Dataless Training of Neural Networks**
Alvaro Velasquez, *Susmit Jha, and Ismail Alkhouri*
40th AAAI Conference on Artificial Intelligence (**AAAI 2026**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[123] **Towards Persistent Noise-Tolerant Active Learning of Regular Languages with Class Query**
*Lekai Chen, Ashutosh Trivedi, and **Alvaro Velasquez***
The 14th International Conference on Learning Representations (**ICLR 2026**)
CORE A* (**best**) conference ranking

[122] **Coarse Adversarial Training with Label Grouping for Robust Classification**
*Ismail R. Alkhouri, Akram Heidarizadeh, **Alvaro Velasquez**, and George K. Atia*
50th IEEE International Conference on Acoustics, Speech and Signal Processing (**ICASSP 2025**)
QUALIS A1 (**best**) conference ranking

[121] **Average Reward Reinforcement Learning for Omega-Regular and Mean-Payoff Objectives**
*Milad Kazemi, Mateo Perez, Fabio Somenzi, Sadegh Soudjani, Ashutosh Trivedi, and **Alvaro Velasquez***
Journal of Artificial Intelligence Research (**JAIR**)
SCImago Q1 (**best**) journal ranking

[120] **Neuro-symbolic Agentic AI: Architectures, Integration Patterns, Applications, Open Challenges, and Future Research Directions**
*Safayat Bin Hakim, Muhammad Adil, **Alvaro Velasquez**, Houbing Herbert Song*
Computer Science Review
SCImago Q1 (**best**) journal ranking

2025 [119] **A Call for Built-In Biosecurity Safeguards for Generative AI Tools**
*Mengdi Wang, Zaixi Zhang, Amrit Singh Bedi, **Alvaro Velasquez**, Stephanie Guerra, Sheng Lin-Gibson, Le Cong, Yuanhao Qu, Souradip Chakraborty, Megan Blewett, Jian Ma, Eric Xing, George Church*
Nature Biotechnology
SciMago Q1 (**best**) journal ranking

[118] **Quadratic Differentiable Optimization for the Maximum Independent Set Problem.**
*Ismail Alkhouri, Cedric Le Denmat, Yingjie Li, Cunxi Yu, Jia Liu, Rongrong Wang, and **Alvaro Velasquez***
42nd International Conference on Machine Learning (**ICML 2025**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[117] **Finite-Time Global Optimality Convergence in Deep Neural Actor-Critic Methods for Decentralized Multi-Agent Reinforcement Learning.**
*Zhiyao Zhang, Myeung Suk Oh, FNU Hairi, Ziyue Luo, **Alvaro Velasquez**, and Jia Liu*
42nd International Conference on Machine Learning (**ICML 2025**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[116] **Immune: Improving Safety Against Jailbreaks in Multi-modal LLMs via Inference-Time Alignment**
*Soumya Suvra Ghosal, Souradip Chakraborty, Vaibhav Singh, Tianrui Guan, Mengdi Wang, Ahmad Beirami, Furong Huang, **Alvaro Velasquez**, Dinesh Manocha, Amrit Singh Bedi*
EEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR 2025**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[115] **TOGA: Temporally Grounded Open-Ended Video QA with Weak Supervision**
*Ayush Gupta, Anirban Roy, Rama Chellappa, Nathaniel D. Bastian, **Alvaro Velasquez**, and Susmit Jha*
International Conference on Computer Vision (**ICCV 2025**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[114] **Automatic Biomarker Discovery and Enrichment with BRAD**
*Joshua Pickard, Marc Choi, Natalie Oliven, Cooper Stansbury, Jillian Cwycyshyn, **Alvaro Velasquez**, and Indika Rajapakse*
Bioinformatics
SciMago Q1 (**best**) journal ranking

[113] **Models for test cost minimization in database migration.**
***Alvaro Velasquez**, Piotr Wojciechowski, K. Subramani, and Matthew Williamson*
INFORMS Journal on Computing
SCImago Q1 (**best**) journal ranking

[112] **Neuro-Symbolic AI as an Antithesis to Scaling Laws**
***Alvaro Velasquez**, Neel P. Bhatt, Ufuk Topcu, Zhangyang Wang, Katia Sycara, Simon Stepputtis, Sandeep Neema, and Gautam Vallabha*
Proceedings of the National Academy of Sciences (**PNAS**) Nexus Journal
SciMago Q1 (**best**) journal ranking

[111] **Exploring Cycle Cover Variants: A Dataless Neural Networks Approach**
*Sangram Jena, K. Subramani, and **Alvaro Velasquez***
Neurocomputing
SciMago Q1 (**best**) journal ranking

[110] **Consensus-based Decentralized Multi-agent Reinforcement Learning for Random Access Network Optimization.**
*Hairi Hairi, Myeung Suk Oh, Zhiyao Zhang, **Alvaro Velasquez**, and Jia Liu*
26th ACM International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (**MobiHoc 2025**)
QUALIS A1 (**best**) conference ranking

[109] **Hybrid Offline Passive Grammatical Inference and Online Planning for Non-Markovian Tasks.**
*Mahyar Alinejad, **Alvaro Velasquez**, Yue Wang, and George Atia*
50th IEEE International Conference on Acoustics, Speech and Signal Processing (**ICASSP 2025**)
QUALIS A1 (**best**) conference ranking

2024 [108] **Assume-Guarantee Reinforcement Learning.**
*Milad Kazemi, Mateo Perez, Fabio Somenzi, Sadegh Soudjani, Ashutosh Trivedi, and **Alvaro Velasquez***
38th AAAI Conference on Artificial Intelligence (**AAAI 2024**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[107] **On the Design of Novel Attention Mechanism for Enhanced Efficiency of Transformers.**
*Sumit Kumar Jha, Susmit Jha, Rickard Ewetz, and **Alvaro Velasquez***
61st ACM/IEEE Design Automation Conference (**DAC 2024**)
QUALIS A1 (**best**) conference ranking

[106] **Robust Average-Reward Reinforcement Learning.**
*Yue Wang, **Alvaro Velasquez**, George Atia, Ashley Prater-Bennette, and Shaofeng Zou*
Journal of Artificial Intelligence Research (**JAIR**)
SCImago Q1 (**best**) journal ranking

[105] **Exploring The Predictive Capabilities of AlphaFold Using Adversarial Protein Sequences.**
*Ismail Alkhouri, Sumit Jha, Andre Beckus, George Atia, Rickard Ewetz, Susmit Jha, and **Alvaro Velasquez***
IEEE Transactions on Artificial Intelligence
SCImago Q1 (**best**) journal ranking

[104] **A Survey on Verification and Validation, Testing and Evaluations of Neurosymbolic Artificial Intelligence.**
*Justus Renkhoff, Ke Feng, Marc Meier-Dornberg, **Alvaro Velasquez**, and Houbing Song*
IEEE Transactions on Artificial Intelligence
SCImago Q1 (**best**) journal ranking

[103] **A Survey on Symbolic Knowledge Distillation of Large Language Models.**
*Kamal Acharya, **Alvaro Velasquez**, and Houbing Song*
IEEE Transactions on Artificial Intelligence
SCImago Q1 (**best**) journal ranking

[102] **Designing Dataless Neural Networks for Kidney Exchange Variants.**
*Sangram Kishor, K. Subramani, and **Alvaro Velasquez***
Neural Computing and Applications
SCImago Q1 (**best**) journal ranking

[101] **LgTS: Dynamic Task Sampling using LLM-generated sub-goals for Reinforcement Learning Agents.**
*Yash Shukla, Wenchang Gao, Vasanth Sarathy, **Alvaro Velasquez**, Robert Wright, and Jivko Sinapo*
23rd International Conference on Autonomous Agents and Multi-Agent Systems (**AAMAS 2024**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[100] **SayNav: Grounding Large Language Models for Dynamic Planning to Navigation in New Environments.**
*Abhinav Rajvanshi, Karan Sikka, Xiao Lin, Bhoram Lee, Han-Pang Chiu, and **Alvaro Velasquez***
The 34th International Conference on Automated Planning and Scheduling (**ICAPS 2024**)
CORE A* (**best**) conference ranking

[99] **Logical Specifications-guided Dynamic Task Sampling for Reinforcement Learning Agents.**
*Yash Shukla, Wenchang Gao, Vasanth Sarathy, **Alvaro Velasquez**, Robert Wright, and Jivko Sinapo*
The 34th International Conference on Automated Planning and Scheduling (**ICAPS 2024**)
CORE A* (**best**) conference ranking

[98] **Arc-Dependent Networks: Theoretical Insights and a Computational Study.**
***Alvaro Velasquez**, Piotr Wojciechowski, K. Subramani, and Matthew Williamson*
Annals of Operations Research
SCImago Q1 (**best**) journal ranking

2023 [97] **Robust Average-Reward Markov Decision Processes.**
Yue Wang, ***Alvaro Velasquez***, *George Atia*, *Ashley Prater-Bennette*, *and Shaofeng Zou*
37th AAAI Conference on Artificial Intelligence (**AAAI 2023**)
Distinguished Paper Award ($12/8777 = 0.136\%(0.7\%)$ of (accepted) submissions)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[96] **Model-Free Robust Average-Reward Reinforcement Learning.**
Yue Wang, ***Alvaro Velasquez***, *George Atia*, *Ashley Prater-Bennette*, *and Shaofeng Zou*
40th International Conference on Machine Learning (**ICML 2023**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[95] **Automaton-Guided Curriculum Generation for Reinforcement Learning Agents.**
Yash Shukla, *Robert Wright*, *Jivko Sinapov*, *Abhishek Kulkarni*, *and* ***Alvaro Velasquez***
33rd International Conference on Automated Planning and Scheduling (**ICAPS 2023**)
CORE A* (**best**) conference ranking

[94] **Neurosymbolic Reinforcement Learning and Planning: A Survey.**
Kamal Acharya, *Waleed Raza*, *Carlos Dourado*, ***Alvaro Velasquez***, *and Houbing Song*
IEEE Transactions on Artificial Intelligence
SCImago Q1 (**best**) journal ranking

[93] **LTL-Based Non-Markovian Inverse Reinforcement Learning.**
Mohammad Afzal, *Ashutosh Gupta*, *Ashutosh Trivedi*, *Krishna A*, *and Sankalp Gambhir*, *and* ***Alvaro Velasquez***
22nd International Conference on Autonomous Agents and Multi-Agent Systems (**AAMAS 2023**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[92] **Resilient Distributed Optimization.**
Jingxuan Zhu, *Yixuan Lin*, ***Alvaro Velasquez***, *and Ji Liu*
American Control Conference (**ACC 2023**)
QUALIS A1 (**best**) conference ranking

[91] **An investigation of the background potential in quantum constrictions using scanning gate microscopy and a swarming algorithm.**
Carlo R. da Cunha, *Nobuyuki Aoki*, *David K. Ferry*, ***Alvaro Velasquez***, *and Yu Zhang*
Physica A: Statistical Mechanics and its Applications
SCImago Q1 (**best**) journal ranking

2022 [90] **A differentiable approach to the maximum independent set problem using dataless neural networks**
Ismail Alkhouri, *George Atia*, *and* ***Alvaro Velasquez***
Neural Networks journal
SCImago Q1 (**best**) journal ranking

[89] **Controller Synthesis for Omega-Regular Specifications with Steady-State Constraints.**
Alvaro Velasquez, *Ismail Alkhouri*, *Andre Beckus*, *George Atia*, *and Ashutosh Trivedi*
21st International Conference on Autonomous Agents and Multi-Agent Systems (**AAMAS 2022**), invited to JAAMAS fast-track publication (**top $\%5$ of submissions**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[88] **Translating Omega-Regular Specifications to Average Objectives for Model-Free Reinforcement Learning.**
Milad Kazemi, Mateo Perez, Fabio Somenzi, Sadegh Soudjani, Ashutosh Trivedi, and ***Alvaro Velasquez***
21st International Conference on Autonomous Agents and Multi-Agent Systems (**AAMAS 2022**), invited to JAAMAS fast-track publication (**top $\%5$ of submissions**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[87] **Shaping Noise for Robust Attributions in Neural Stochastic Differential Equations.**
Sumit K. Jha, Rickard Ewetz, ***Alvaro Velasquez****, Arvind Ramanathan, Susmit Jha*
36th AAAI Conference on Artificial Intelligence (**AAAI 2022**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[86] **ExplainIt!: A Tool for Computing Robust Attributions of DNNs**
Sumit Jha, ***Alvaro Velasquez****, Rickard Ewetz, Laura Pullum, Susmit Jha*
Thirty-First International Joint Conference on Artificial Intelligence (**IJCAI 2022**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[85] **Steady-State Planning in Expected Reward Multichain MDPs.**
George Atia, Andre Beckus, Ismail Alkhouri, and ***Alvaro Velasquez***
Journal of Artificial Intelligence Research (**JAIR**)
SCImago Q1 (**best**) journal ranking

[84] **Multi-Agent Tree Search with Dynamic Reward Shaping.**
Alvaro Velasquez*, Brett Bissey, Lior Barak, Daniel Melcer, Andre Beckus, Ismail Alkhouri and George Atia*
32nd International Conference on Automated Planning and Scheduling (**ICAPS 2022**)
CORE A* (**best**) conference ranking

[83] **Inferring Probabilistic Reward Machines from Non-Markovian Reward Signals for Reinforcement Learning.**
Taylor Dohmen, Noah Topper, George Atia, Andre Beckus, Ashutosh Trivedi, and ***Alvaro Velasquez***
32nd International Conference on Automated Planning and Scheduling (**ICAPS 2022**)
CORE A* (**best**) conference ranking

[82] **Active Grammatical Inference for Non-Markovian Planning.**
Noah Topper, George Atia, Ashutosh Trivedi, and ***Alvaro Velasquez***
32nd International Conference on Automated Planning and Scheduling (**ICAPS 2022**)
CORE A* (**best**) conference ranking

[81] **Synthesis of Adversarial Samples in Two-Stage Classifiers.**
Ismail Alkhouri, ***Alvaro Velasquez****, and George Atia*
47th IEEE International Conference on Acoustics, Speech and Signal Processing (**ICASSP 2022**)
QUALIS A1 (**best**) conference ranking

[80] **On the Complexity of and Solutions to the Minimum Stopping and Trapping Set Problems.**
Alvaro Velasquez*, K. Subramani, and Piotr Wojciechowski*
Theoretical Computer Science journal

2021 [79] **Dynamic Automaton-Guided Reward Shaping for Monte-Carlo Tree Search.**
Alvaro Velasquez, *Brett Bissey, Lior Barak, Ismail Alkhouri, Andre Beckus, and George Atia*
35th AAAI Conference on Artificial Intelligence (**AAAI 2021**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[78] **On Smoother Attributions using Neural Stochastic Differential Equations.**
Sumit K. Jha, Rickard Ewetz, ***Alvaro Velasquez****, and Susmit Jha*
30th International Joint Conference on Artificial Intelligence (**IJCAI 2021**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[77] **Automated Synthesis of Quantum Circuits using Symbolic Abstractions and Decision Procedures.**
Alvaro Velasquez*, Sumit K. Jha, Rickard Ewetz, and Susmit Jha*
2021 IEEE International Symposium on Circuits and Systems (**ISCAS 2021**)
QUALIS A1 (**best**) conference ranking

2020 [76] **Steady-State Policy Synthesis for Multichain Markov Decision Processes.**
George Atia, Andre Beckus, Ismail Alkhouri, and ***Alvaro Velasquez***
29th International Joint Conference on Artificial Intelligence (**IJCAI 2020**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

2019 [75] **Steady-state policy synthesis for verifiable control.**
Alvaro Velasquez
28th International Joint Conference on Artificial Intelligence (**IJCAI 2019**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

Remaining Peer-Reviewed Publications

2026 [74] **A Scalable Lift-and-Project Differentiable Approach For the Maximum Cut Problem.**
Ismail Alkhouri, Mian Wu, Cunxi Yu, Jia Liu, Rongrong Wang, ***Alvaro Velasquez***
The 29th International Conference on Artificial Intelligence and Statistics (**AISTATS 2026**).

[73] **Lightweight satisfiability solving using dataless neural networks.**
Andrew Gautier, Piotr Wojciechowski, Sangram K. Jena, K. Subramani, ***Alvaro Velasquez***
International Symposium for Artificial Intelligence and Mathematics (**ISAIM 2026**).

2025 [72] **From Abstraction to Reality: DARPA's Vision for Robust Sim-to-Real Autonomy.**
Erfaun Noorani, Zachary Serlin, Ben Price, and ***Alvaro Velasquez***
AAAI AI Magazine.

[71] **Gaps in Generalization: Frontier Problems for Neurosymbolic AI.**
Alvaro Velasquez
3rd International Conference on Bridging the Gap Between AI and Reality (**AISoLA 2025**).

[70] **Advancing Discrete Optimization: Novel Approaches with Dataless Neural Networks.**
Sangram Jena, K. Subramani, and ***Alvaro Velasquez***
Journal of Combinatorial Optimization.

[69] **Zero-Shot Detection of Out-of-Context Objects Using Foundation Models.**
Anirban Roy, Adam Cobb, Ramneet Kaur, Sumit Jha, Nathaniel Bastian, Robert H. Thomson, Iain Cruikshank, ***Alvaro Velasquez****, and Susmit Jha*
IEEE/CVF Winter Conference on Applications of Computer Vision (**WACV 2025**).

[68] **Explainable AI for Digital Twins via Neuro-Symbolic Rule Extraction.**
*Safayat Bin Hakim, Muhammad Adil, **Alvaro Velasquez**, and Houbing H Song*
22nd IEEE International Conference on Ubiquitous Intelligence and Computing (**UIC 2025**).

[67] **Bidirectional End-to-End Framework for Transfer from Abstract Models in Non-Markovian Reinforcement Learning.**
*Mahyar Alinejad, Precious Nwaorgu, Chinwendu Enyioha, Yue Wang, **Alvaro Velasquez**, and George Atia*
2nd International Conference on Neurosymbolic Systems (**NeuS 2025**).

[66] **SymRAG: Efficient Neuro-Symbolic Retrieval Through Adaptive Query Routing.**
*Safayat Bin Hakim, Muhammad Adil, Houbing Herbert Song, and **Alvaro Velasquez***
19th International Conference on Neurosymbolic Learning and Reasoning (**NeSy 2025**).

[65] **KGAccel: A Domain-Specific Reconfigurable Accelerator for Knowledge Graph Reasoning.**
*Hanning Chen, Ali Zakeri, Yang Ni, Fei Wen, Behnam Khaleghi, Hugo Latapie, **Alvaro Velasquez**, and Mohsen Imani*
The 2nd International Conference on Neurosymbolic Systems (**NeuS 2025**).

2024 [64] **Controller Synthesis for Linear Temporal Logic and Steady-State Specifications.**
***Alvaro Velasquez**, Ismail Alkhouri, Andre Beckus, Ashutosh Trivedi, and George Atia*
Journal of Autonomous Agents and Multi-Agent Systems.

[63] **The Hexatope and Octatope Abstract Domains for Neural Network Verification.**
*Stanley Bak, Taylor Dohmen, K. Subramani, Ashutosh Trivedi, **Alvaro Velasquez**, and Piotr Wojciechowski*
Formal Methods in System Design.

[62] **A differential approach for several NP-hard optimization problems.**
*Sangram Kishor, K. Subramani, and **Alvaro Velasquez***
The 18th International Symposium for Artificial Intelligence and Mathematics (**ISAIM 2024**).

[61] **An Exploration of Optimizing Kidney Exchanges with Graph Machine Learning.**
*Calvin Nau, Prashanth Sankarian, Moises Sudit, Payam Khazaelpour, Katie McConky, **Alvaro Velasquez**, and Liise Kayler*
The 14th IEEE Conference on Cognitive and Computational Aspects of Situation Management (**CogSIMA 2024**).

[60] **Byzantine-Resilient Decentralized Multi-Armed Bandits.**
*Jingxuan Zhu, Alec Koppel, **Alvaro Velasquez**, and Ji Liu*
Transactions on Machine Learning Research (**TMLR**).

[59] **Neuro-symbolic Generative AI Assistant for System Design.**
*Susmit Jha, Sumit K. Jha, and **Alvaro Velasquez***
22nd ACM-IEEE International Symposium on Formal Methods and Models for System Design (**MEMOCODE 2024**).

[58] **FuseGAT: Advancing Graph Attention Networks for Robust Recommendations through Explicit Side Information Fusion.**
Dong Hyun Jeon, Wenbo Sun, Houbing Herbert Song, Dongfang Liu, ***Alvaro Velasquez****, Yixin Chloe Xie, and Shuteng Niu*
IEEE Conference on Big Data (**BigData 2024**).

[57] **Automated Synthesis of Hardware Designs using Symbolic Feedback and Grammar-Constrained Decoding in Large Language Models.**
Sumit K. Jha, Susmit Jha, Muhammad Rashedul Haq Rashed, Rickard Ewetz, and ***Alvaro Velasquez***
IEEE National Conference on Aerospace and Electronics (**NAECON 2024**).

[56] **On the Hardness of Decentralized Multi-Agent Policy Evaluation under Byzantine Attacks.**
Hairi, Minghong Fang, Zifan Zhang, textbfAlvaro Velasquez, and Jia Liu
The 22nd IEEE International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (**WiOpt 2024**).

[55] **Improving Air Mobility for Pre-Disaster Planning with Neural Network Accelerated Genetic Algorithm.**
Justus Renkhoff, Ke Feng, Marc Meier-Dornberg, ***Alvaro Velasquez****, and Houbing Song*
27th IEEE International Conference on Intelligent Transportation Systems (**ITSC 2024**).

[54] **Co-Synthesis of Code and Formal Models Using Large Language Models and Functors.**
Sumit K. Jha, Susmit Jha, Rickard Ewetz, and ***Alvaro Velasquez***
42nd IEEE Military Communications Conference (**MILCOM 2024**)

[53] **Solving Mystery Planning Problems Using Category Theory, Functors, and Large Language Models.**
Sumit K. Jha, Susmit Jha, Rickard Ewetz, and ***Alvaro Velasquez***
IEEE International Conference on Assured Autonomy (**ICAA 2024**)

[52] **TaskCLIP: Extend Large Vision-Language Model for Task Oriented Object Detection.**
Hanning Chen, Wenjun Huang, Yang Ni, Sanggeon Yun, Yezi Liu, Fei Wen, ***Alvaro Velasquez****, Hugo Latapie, and Mohsen Imani*
FOundation models Creators meet USers (**FOCUS Worskhop at ECCV 2024**).

2023 [51] **A Non-Targeted Attack Approach for the Coarse Mis-Classification Problem.**
Ismail Alkhouri, George Atia, and ***Alvaro Velasquez***
International Joint Conference on Neural Networks (**IJCNN 2023**)

[50] **NoiseCAM: Explainable AI for the Boundary Between Noise and Adversarial Attacks.**
Wenkai Tan, Justus Renkhoff, ***Alvaro Velasquez****, Ziyu Wang, Lusi Li, Jian Wang, Shuteng Niu, Fan Yang, Yongxin Liu, and Houbing Song*
IEEE International Conference on Fuzzy Systems (**FUZZ 2023**)

[49] **The Octatope Abstract Domain for Verification of Neural Networks.**
Taylor Dohmen, Stanley Bak, Ashutosh Trivedi, Piotr Wojciechowski K. Subramani, and ***Alvaro Velasquez***
25th International Symposium on Formal Methods (**FM 2023**)

[48] **Reachability in Choice Networks.**
Piotr Wojciechowski, K. Subramani, and ***Alvaro Velasquez***
Discrete Optimization journal.

[47] **A Resilient Distributed Algorithm for Solving Linear Equations.**
Jingxuan Zhu, ***Alvaro Velasquez***, *and Ji Liu*
62nd IEEE Conference on Decision and Control (**CDC 2023**)

[46] **Priority-based bin packing with subset constraints.**
Bugra Caskurlu, K. Subramani, ***Alvaro Velasquez***, *and Piotr Wojciechowski*
Discrete Applied Mathematics journal.

[45] **Differentiable Discrete Optimization using Dataless Neural Networks.**
Sangram Kishor, K. Subramani, and ***Alvaro Velasquez***
16th Annual International Conference on Combinatorial Optimization and Applications (**COCOA 2023**).

[44] **Farkas bounds on Horn constraint systems.**
Piotr Wojciechowski, K. Subramani, and ***Alvaro Velasquez***
Theory of Computing Systems journal.

[43] **Dehallucinating Large Language Models Using Formal Methods Guided Iterative Prompting.**
Susmit Jha, Sumit K. Jha, Patrick Lincoln, Nathaniel D. Bastian, ***Alvaro Velasquez****, and Sandeep Neema*
IEEE International Conference on Assured Autonomy (**ICAA 2023**)

[42] **Safety Margins for Reinforcement Learning.**
Walt Woods, Alexander Grushin, ***Alvaro Velasquez****, and Simon Khan*
IEEE Conference on Artificial Intelligence (**CAI 2023**)

[41] **Counterexample Guided Inductive Synthesis Using Large Language Models and Satisfiability Solving.**
Sumit K. Jha, Susmit Jha, Patrick Lincoln, Nathaniel D. Bastian, ***Alvaro Velasquez****, Rickard Ewetz, and Sandeep Neema*
41st IEEE Military Communications Conference (**MILCOM 2023**)
Best paper award candidate

[40] **Neural SDEs for Robust and Explainable Analysis of Electromagnetic Unintended Radiated Emissions.**
Sumit K. Jha, Susmit Jha, Rickard Ewetz, and ***Alvaro Velasquez***
41st IEEE Military Communications Conference (**MILCOM 2023**)

[39] **The Utility of Feature Reuse: Transfer Learning in Data-Starved Regimes.**
Rashik Shadman, M.G. Sarwar Murshed, Edward Verenich, ***Alvaro Velasquez****, and Faraz Hussain*
IEEE International Conference on Computational Science and Computational Intelligence (**CSCI 2023**). Acceptance rate: 18%.

2022 [38] **Resilient Constrained Consensus over Complete Graphs via Feasibility Redundancy.**
Jingxuan Zhu, Yixuan Lin, ***Alvaro Velasquez****, and Ji Liu*
American Control Conference (**ACC 2022**)
QUALIS A1 (**best**) conference ranking

[37] **Domain Wall Leaky Integrate-and-Fire Neurons with Shape-Based Configurable Activation Functions.**
Wesley H. Brigner, Naimul Hassan, Xuan Hu, Christopher H. Bennett, Felipe Garcia-Sanchez, Can Cui, ***Alvaro Velasquez****, Matthew J. Marinella, Jean Anne C. Incorvia, and Joseph S. Friedman*
IEEE Transactions on Electron Devices

[36] **Optimal Controller Synthesis from Steady-State Distributions.**
Alvaro Velasquez, *Ismail Alkhouri, K. Subramani, Piotr Wojciechowski, and George Atia*
Journal of Automated Reasoning

[35] **On the Coarse Robustness of Classifiers.**
Ismail Alkhouri, ***Alvaro Velasquez****, Stanley Bak, and George Atia*
56th Asilomar Conference on Signals, Systems, and Computers

[34] **Reinforced Contrastive Graph Neural Networks (RCGNN) for Anomaly Detection.**
Zenan Sun, Jingyi Su, Donghyun Jeon, ***Alvaro Velasquez****, Houbing Song, and Shuteng Niu*
IEEE International Performance, Computing, and Communications Conference (**IPCCC 2022**)

[33] **Exploring Adversarial Attacks on Neural Networks: An Explainable Approach.**
Justus Renkhoff, Wenkai Tan, ***Alvaro Velasquez****, William Yichen Wang, Yongxin Liu, Jian Wang, Shuteng Niu, Lejla Begic Fazlic, Guido Dartmann, and Houbing Song*
IEEE International Performance, Computing, and Communications Conference (**IPCCC 2022**)

[32] **New results in Priority-Based bin packing.**
K. Subramani, Piotr Wojciechowski, and ***Alvaro Velasquez***
7th International Symposium on Algorithmic Aspects of Cloud Computing (**ALGOCLOUD 2022**)

[31] **Data-Driven Robust Multi-Agent Reinforcement Learning.**
Shaofeng Zou and ***Alvaro Velasquez***
32nd IEEE International Workshop on Machine Learning for Signal Processing (**MLSP 2022**)

[30] **An Approach to the Maximum Independent Set Problem Using Graph-Based Neural Network Structures.**
Ismail Alkhouri, George Atia, and ***Alvaro Velasquez***
32nd IEEE International Workshop on Machine Learning for Signal Processing (**MLSP 2022**)

[29] **The Minimum Value State Problem in Actor-Critic Networks.**
Alvaro Velasquez*, Ismail Alkhouri, Brett Bissey, Lior Barak, and George Atia*
32nd IEEE International Workshop on Machine Learning for Signal Processing (**MLSP 2022**)

[28] **BOSS: Bidirectional One-Shot Synthesis of Adversarial Examples.**
Ismail Alkhouri, ***Alvaro Velasquez****, and George Atia*
32nd IEEE International Workshop on Machine Learning for Signal Processing (**MLSP 2022**)

[27] **On the Approximability of Path and Cycle Problems in Arc-Dependent Networks.**
Piotr Wojciechowski, K. Subramani, ***Alvaro Velasquez****, and Matthew Williamson*
8th Annual International Conference on Algorithms and Discrete Applied Mathematics (**CALDAM 2022**)

[26] **Reachability Problems in Interval-Constrained and Cardinality-Constrained Graphs.**
Alvaro Velasquez*, Piotr Wojciechowski, and K. Subramani*
Journal of Discrete Mathematics, Algorithms and Applications (**DMAA**)

[25] **Analyzing the Reachability Problem in Choice Networks.**
Piotr Wojciechowski, K. Subramani, and ***Alvaro Velasquez***
19th International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (**CPAIOR 2022**)

2021 [24] **Consensus-Based Value Iteration for Multiagent Cooperative Control.**
Jing Wang, Elias Wilson, and ***Alvaro Velasquez***
60th IEEE Conference on Decision and Control (**CDC 2021**)

[23] **Mitigating the Class Overlap Problem in Discriminative Localization: COVID-19 and Pneumonia Case Study.**
Edward Verenich, M. G. S. Murshed, Nazar Khan, ***Alvaro Velasquez****, and Faraz Hussain*
Book chapter in **Explainable AI Within the Digital Transformation and Cyber Physical Systems**.

[22] **Adversarial Perturbation Attacks on Nested Dichotomies Classification Systems.**
Ismail Alkhouri, ***Alvaro Velasquez****, and George Atia*
31st IEEE International Workshop on Machine Learning for Signal Processing (**MLSP 2021**)

[21] **Algorithmic Analysis of Priority-Based Bin Packing.**
Piotr Wojciechowski, K. Subramani, ***Alvaro Velasquez****, and Bugra Caskurlu*
7th Annual International Conference on Algorithms and Discrete Applied Mathematics (**CALDAM 2021**)

[20] **On the Copy Complexity of Width-3 Horn Constraint Systems.**
K. Subramani, Piotr Wojciechowski, and ***Alvaro Velasquez***
13th International Symposium on Frontiers of Combining Systems (**FroCoS 2021**)

2020 [19] **Verification-Guided Tree Search.**
Alvaro Velasquez *and Daniel Melcer*
19th International Conference on Autonomous Agents and Multi-Agent Systems (**AAMAS 2020**)
CORE A* (**best**) and QUALIS A1 (**best**) conference ranking

[18] **Plasticity-Enhanced Domain-Wall MTJ Neural Networks for Energy-Efficient Online Learning.**
Christopher H. Bennett, T. Patrick Xiao, Can Cui, Naimul Hassan, Otitoaleke G. Akinola, Jean Anne C. Incorvia, ***Alvaro Velasquez****, Joseph S. Friedman, and Matthew J. Marinella*
IEEE International Symposium on Circuits and Systems (**ISCAS 2020**)
QUALIS A1 (**best**) conference ranking

[17] **Improving Explainability of Image Classification in Scenarios with Class Overlap: Application to COVID-19 and Pneumonia.**
Edward Verenich, ***Alvaro Velasquez****, Nazar Khan, Faraz Hussain*
19th IEEE International Conference on Machine Learning and Applications (**ICMLA 2020**)

[16] **FlexServe: Deployment of PyTorch Models as Flexible REST Endpoints.**
Edward Verenich, ***Alvaro Velasquez****, M. G. Sarwar Murshed, and Faraz Hussain*
USENIX Conference on Operational Machine Learning (**OpML 2020**)

2019 [15] **Spatially Efficient In-Memory Addition Through Destructive and Non-Destructive Operations.**
Alvaro Velasquez *and Benjamin Shaia*
IEEE International Symposium on Circuits and Systems (**ISCAS 2019**)
QUALIS A1 (**best**) conference ranking

[14] **On the Susceptibility of Deep Neural Networks to Natural Perturbations.**
Mesut Ozdag, Sunny Raj, Steven L. Fernandes, ***Alvaro Velasquez****, Laura Pullum, and Sumit K. Jha*
AISafety Workshop at IJCAI 2019

2018 [13] **In-memory computing using paths-based logic and heterogeneous components.**
Alvaro Velasquez *and Sumit K. Jha*
Design, Automation Test in Europe Conference Exhibition (**DATE 2018**)
QUALIS A1 (**best**) conference ranking

[12] **3D Crosspoint Memory as a Parallel Architecture for Computing Network Reachability.**
Alvaro Velasquez *and Sumit K. Jha*
36th IEEE International Conference on Computer Design (**ICCD 2018**)

[11] **Brief Announcement: Parallel Transitive Closure Within 3D Crosspoint Memory.**
Alvaro Velasquez *and Sumit K. Jha*
30th on Symposium on Parallelism in Algorithms and Architectures (**SPAA 2018**)

[10] **Minimization of Testing Costs in Capacity-Constrained Database Migration.**
K. Subramani, Bugra Caskurlu, and ***Alvaro Velasquez***
4th International Symposium on Algorithmic Aspects of Cloud Computing (**ALGOCLOUD 2018**)

[9] **Finding Minimum Stopping and Trapping Sets: An Integer Linear Programming Approach.**
Alvaro Velasquez*, K. Subramani, and Steven L. Drager*
International Symposium on Combinatorial Optimization (**ISCO 2018**)

2017 [8] **Computation of Boolean matrix chain products in 3D ReRAM.**
Alvaro Velasquez *and Sumit K. Jha*
IEEE International Symposium on Circuits and Systems (**ISCAS 2017**)
QUALIS A1 (**best**) conference ranking.

2016 [7] **Parallel boolean matrix multiplication in linear time using rectifying memristors.**
Alvaro Velasquez *and Sumit K. Jha*
IEEE International Symposium on Circuits and Systems (**ISCAS 2016**)
QUALIS A1 (**best**) conference ranking.

[6] **Flow-based computing on nanoscale crossbars: Design and implementation of full adders.**
Zahirudden Alamgir, Karsten Beckmann, Nathaniel C. Cady, ***Alvaro Velasquez****, and Sumit K. Jha*
IEEE International Symposium on Circuits and Systems (**ISCAS 2016**)
QUALIS A1 (**best**) conference ranking.

[5] **The Cardinality-Constrained Paths Problem: Multicast Data Routing in Heterogeneous Communication Networks.**
Alvaro Velasquez*, Piotr Wojciechowski, K. Subramani, Steven L. Drager, and Sumit K. Jha*
15th IEEE International Symposium on Network Computing and Applications (**NCA 2016**)

2015 [4] **Fault-tolerant in-memory crossbar computing using quantified constraint solving.**
Alvaro Velasquez *and Sumit K. Jha*
33rd IEEE International Conference on Computer Design (**ICCD 2015**)

[3] **Automated synthesis of crossbars for nanoscale computing using formal methods.**
Alvaro Velasquez *and Sumit K. Jha*
2015 IEEE/ACM International Symposium on Nanoscale Architectures (**NANOARCH 2015**)

2014 [2] **Parallel computing using memristive crossbar networks: Nullifying the processor-memory bottleneck.**
Alvaro Velasquez *and Sumit K. Jha*
9th International Design and Test Symposium (**IDT 2014**)

[1] **Putting humpty-dumpty together: Mining causal mechanistic biochemical models from big data.**
Faraz Hussain, ***Alvaro Velasquez****, Emily Sassano, and Sumit K. Jha*
IEEE 4th International Conference on Computational Advances in Bio and Medical Sciences (**ICCABS 2014**)

Patents

[3] **3D Crossbar Architecture for Fast, Efficient In-Memory Computing of Graph Transitive Closure.**
Alvaro Velasquez *and Sumit K. Jha*
Patent Number: US11538989B2

[2] **Computation of Boolean Formulas Using Sneak Paths in Crossbar Computing.**
Sumit K. Jha, Dilia E. Rodriguez, Joseph E. van Nostrand, and ***Alvaro Velasquez***
Patent Number: US9319047B2

[1] **Spintronic Computing Architecture and Method.**
Friedman, Joseph S., Wesley H. Brigner, Naimul Hassan, Xuan Hu, and ***Alvaro Velasquez***
Patent Number: US11977970B2

Program Management and Leadership

FY24 - FY27 **Transfer from Imprecise and Abstract Models to Autonomous Technologies (TIAMAT).**
Amount: $51,100,000

FY23 - FY27 **Assured Neuro-Symbolic learning and Reasoning (ANSR).**
Amount: $50,000,000

FY20 - FY24 **Guaranteeing AI Robustness against Deception (GARD).**
Amount: $59,000,000

FY24 - FY25 **Foundation Models for Scientific Discovery (FoundSci).**
Amount: $5,000,000

FY23 - FY24 **Geometries of Learning (GoL).**
Amount: $7,100,000

FY23 - FY24 **Reverse-Engineering of Deceptions (RED).**
Amount: $9,800,000

FY24 - FY26 **University Partnership in AI.**
Amount: $9,000,000

FY25 - FY27 **Dataless Neural Networks.**
Amount: $2,000,000

FY25 - FY27 **Countering Bias in AI (STTRs).**
Amount: $14,400,000

FY24 - FY26 **Young Faculty Awards.**
Amount: $1,500,000

FY25 - FY26 **Transition Council Funding.**
Amount: $2,800,000

FY21 - FY25 **Various SBIRs, STTRs, and seedlings. More information available upon request.**

External Research Funding

FY26 - FY27 **Simulated Alignment: AI Alignment as a Sequential Decision-Making Problem.**
PI: ***Alvaro Velasquez***
DARPA
BAA Topic: Information Innovation Office
Amount: $300,000

FY26 - FY27 **Adversarial Self-Alignment for Robust Mission Adherence in Generative AI.**
*PI: Amrit Bedi, co-PI:****Alvaro Velasquez***
DARPA
BAA Topic: Information Innovation Office
Amount: $340,000, co-PI share: $125,000

FY20 - FY22 **Verifiable Interactive Planning.**
PI: ***Alvaro Velasquez***
Air Force Office of Scientific Research (AFOSR)
BAA Topic: Optimization and Discrete Mathematics
Amount: $420,000

FY21 **Explainable Autonomy via Learned Automaton Representations.**
PI: ***Alvaro Velasquez***
Air Force Research Laboratory
Entrepreneurial Research Funds
Amount: $40,000

FY20 **Verifiable Reinforcement Learning with Testing for Assured Autonomy.**
PI: ***Alvaro Velasquez***
Air Force Office of Scientific Research (AFOSR)
BAA Topic: Agile Science of Test and Evaluation
Amount: $50,000

FY19 **Verifiable Reinforcement Learning for Trustworthy Autonomy.**
PI: ***Alvaro Velasquez***
Air Force Research Laboratory
Entrepreneurial Research Funds
Amount: $38,000

Speaking Engagements

Keynote Talks

2025 [14] **Challenges and Opportunities for Military Autonomy in the Era of Foundation Models**
Autonomy for Defense Conference

[13] **Challenges and Opportunities for Foundation Models in Military Autonomy and Logistics**
NATO Research Symposium on AI Security and Assurance in Military Systems

[12] **Gaps in Generalization: A Case for Neurosymbolic AI**
AISoLA

[11] **The ABCs of Neurons and Graphs: Autonomy, Biology, and Creativity**
MIT Lincoln Laboratory Graph Exploitation Symposium (GraphEx)

[10] **Neurosymbolic AI in Autonomy, Biology, and Creativity**
AAAI Spring Symposium on Agentic AI for Science

[9] **Neurons, Symbols, and Scale: Challenges and Opportunities in Neuro-Symbolic Foundation Models**
IEEE Winter Conference on Applications of Computer Vision (WACV 2025)

2024 [8] **Frontiers in Adversarial AI: Autonomy and Biology**
3rd ICML Workshop on New Frontiers in Adversarial Machine Learning (AdvML-Frontiers)

[7] **Neurons, Symbols, and Semantics: Challenges and Opportunities in Verification of Neuro-Symbolic AI**
AISoLA

[6] **Foundation Models for Assured Autonomy: from Abstraction to Reality**
IEEE International Conference on Assured Autonomy (ICAA 2024)

[5] **Neuro-Symbolic AI in the Era of Large Foundation Models**
3rd Neuro-Symbolic Summer School

[4] **Autonomy: From Abstraction to Reality**
Automotive Research Center Program Review, University of Michigan

[3] **DoD Challenges and Opportunities in the Era of Foundation Models**
SPIE AI and ML for MDO

2022 [2] **Of Neurons and Symbols: A Bird's Eye View of Neuro-Symbolic Artificial Intelligence.**
8th International Workshop on Symbolic-Numeric Methods for Reachability Analysis

[1] **Neuro-Symbolic Autonomy and the Quest for Smart(er) Decision-Making.**
IEEE Cybermatics Congress 2022

Invited Talks

2026 [38] **???**
Distinguished Speaker Series at the University of Virginia. Host: Ferdinando Fioretto.

[37] **A Vision for Neurosymbolic AI and Efficient Generalization**
University of Colorado Boulder National Security Institute.

2025 [36] **Dataless Neural Networks: Training from Structure, Not from Data**
INFORMS Annual Meeting.

[35] **Efficient Neurosymbolic Autonomy in the Age of Large Models**
Duke University. Host: Hai Li.

[34] **The ABCs of Frontier Problems in AI: Autonomy, Biology, and Creativity.**
Department of Defense Leaders Building Leaders (LBL) Symposium.

2024 [33] **Challenges and opportunities for neuro-symbolic autonomy**
NDIST

[32] **Research Funding Opportunities in the Federal Government**
Florida Defense Science and Technology University Symposium

[31] **Foundation Models for Autonomy: From Simulation to Reality**
Uncrewed Systems Conference

[30] **The Cat-and-Mouse Game of Adversarial Artificial Intelligence**
National Security Agency (NSA)

[29] **Can We Train Foundation Models on Phones? A Neuro-Symbolic Perspective**
Intel Corporation

[28] **The Cat-and-Mouse Game of Adversarial Artificial Intelligence**
DEFCON

[27] **Bitter Lessons and Scaling Laws in the Era of Large Models**
Carnegie Mellon University, host: Pradeep Ravikumar

[26] **Neuro-Symbolic AI in the Era of Large Models**
NSF-Sponsored Workshop on Hardware-Software Co-Design for Neuro-Symbolic Computation

[25] **Refuting the Bitter Lesson: Neuro-Symbolic AI in Theory and Practice**
Intel

[24] **Refuting the Bitter Lesson: Neuro-Symbolic AI in Theory and Practice**
MSU CSME Colloquium

[23] **Frontiers in Neuro-Symbolic AI: Going Beyond Vision and Language**
UMD College Park campus, hosted by the Center for Machine Learning, led by Tom, and the IEEE Signal Processing Washington Chapter, led by Furong.

[22] **Frontiers in Neuro-Symbolic AI: Going Beyond Vision and Language**
NSA Data Science Seminar

2023 [21] **Foundation Models and the Transfer of Embodied Autonomy**
Johns Hopkins University

[20] **DoD Challenges to LPTMs**
R&E Trusted AI and Autonomy event

[19] **Foundation Models for Autonomy and Scientific Discovery.**
UMBC Distinguished Speaker Series

[18] **Neuro-Symbolic Transfer Learning for Embodied Autonomy.**
MSU Colloquium Series

[17] **Neuro-Symbolic Computing at the Edge**
ERI Summit.

[16] **Neuro-Symbolic Foundation Models for Autonomy.**
17th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy 2023)

[15] **Foundation Models for Sequential Decision-Making and Transfer.**
Ohio State University, host: Dr. Jia Liu.

[14] **Foundation Models for Autonomy.**
Vanderbilt University, host: Dr. Sandeep Neema.

[13] **Neuro-Symbolic Transfer and Optimization.**
University of Southern California. Host: Dr. Jyotirmoy V. Deshmukh.

[12] **Neurons and Symbols at the Edge.**
CMU and Department of Defense AI Fusion Workshop

2022 [11] **Neuro-Symbolic Learning and Reasoning: Challenges and Opportunities.**
4th Buffalo Day for 5G and Wireless Internet of Things

[10] **Neurons, Symbols, and the Transfer of Knowledge.**
SRI International. Host: Dr. Ajay Divakaran

[9] **Neuro-Symbolic Learning, Reasoning, and Knowledge Transfer.**
University of Texas at Austin. Host: Dr. Ufuk Topcu

[8] **Challenges and Opportunities in Neurosymbolic Artificial Intelligence.**
18th International Conference on Emerging Technologies for a Smarter World (**CEWIT 2022**)

[7] **Neurolinguistic Autonomy.**
University of Colorado Boulder. Host: Ashutosh Trivedi

2021 [6] **Neurolinguistic Decision-Making.**
University of Texas at San Antonio. Host: Sushil Prasad

[5] **Artificial Intelligence and Machine Learning for Joint All-Domain Command, Control and Multi-Domain Operations.**
AAAI 2020 Fall Symposium track on AI in Government and the Public Sector

[4] **Automaton-Guided Learning.**
West Virginia University. Host: Sarika Khushalani Solanki

2020 [3] **Unintended Behavior, Trust, and Testing in Autonomous Systems.**
Air Force Summer Study Group on Unintended Consequences of Autonomous Behaviors

2019 [2] **Disruptive Effects of Artificial Intelligence and Machine Learning: Challenges and Opportunities.**
Griffiss Institute and Cyber Research Institute Dining and Development Series

2018 [1] **Steady-State Policy Synthesis for Verifiable Control.**
IEEE Mohawk Valley Section. Host: Lee Seversky

Panels

2025 [36] **Industry/Research and Development Perspectives and Priorities**
University of California Berkeley Technology and Industrial Policy Center (TCIP) Retreat.

[35] **Is symbolic reasoning necessary for robust generalization in AI?**
International Neurosymbolic Systems Conference (**NeuS 2025**).

[34] **AI, Formal Methods, and Mathematical Reasoning (AIMing)**
National Science Foundation Panel.

[33] **The Nexus of Biosecurity, AI, and Modern Conflict: A Tabletop Exercise**
Summit on Modern Conflict and Emerging Threats, Vanderbilt University.

[32] **AI for Safety-Critical Systems: Are We Ready?**
59th Annual Conference on Information Science and Systems (**CISS 2025**).

[31] **Trustworthy AI Implications in Cybersecurity.**
Northeastern University.

[30] **Fireside Chat with Dean.**
University of Central Florida Distinguished Alumnus Award Ceremony.

[29] **Certifiability of Aviation Systems with ML Components.**
MOSA Industry and Government Summit and Expo.

[28] **Physical AI: Moving From Bits to Atoms.**
CES 2025

[27] **Certification of AI-Enabled Systems.**
AIAA SciTech Forum

2024 [26] **The Future of Verification for Neuro-Symbolic AI.**
International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (AISoLA)

[25] **Resilient AI/ML in the Cyber and Edge Domains: Challenges and Opportunities.**
43rd IEEE Military Communications Conference (MILCOM)

[24] **SIM to REAL: The Metaverse Meets Reality**
AI and Automation Summit at Indianapolis Motor Speedway

[23] **Neuromorphic Systems and U.S. Government**
International Conference on Neuromorphic Systems (ICONS)

[22] **Physics-Based AI for Trusted Autonomy**
Stanford University

[21] **What will bring AI towards AGI?**
IEEE World Congress on Computational Intelligence (WCCI)

[20] **DOD AI/ML and Generative AI initiatives.**
Amazon Web Services (AWS) Defense Day.

[19] **Adversarial Steering of Autonomy.**
Army Intelligence Information Center (AI2C) Technical Exchange Meeting.

[18] **Government AI Safety Programs.**
Department of Transportation AI Assurance Workshop

[17] **Neurosymbolic Computation in Autonomous Cyber-Physical Systems**
NSF-Sponsored Workshop on Hardware-Software Co-Design for Neuro-Symbolic Computation

[16] **MIDAS AI for Scientific Discovery.**
First Conference on Scientific Foundation Models (SciFM)

[15] **Digital Twins and the Omniverse Roundtable.**
NVIDIA GTC (private event)

[14] **AI Panel.**
Creative Disruptors in the Desert

[13] **Funding Opportunities with DARPA.**
CRA Mentoring Workshop

[12] **Adversarial AI: Bridging Research With Operations.**
Advantage DoD 2024: Defense Data and AI Symposium

[11] **Jefferson Talks: AI and the Future of Defense.**
Advantage DoD 2024: Defense Data and AI Symposium (private event)

[10] **Sim-to-real autonomy in racing.**
CES 2024

2023 [9] **Generative AI and LLMs for Healthcare.**
ASU-Mayo Workshop on Generative AI and Large Language Models: Opportunities and Challenges for Healthcare and Medicine

[8] **The Role of Federal Agencies in AI.**
UM MIDAS Summit

[7] **Data Science Student Consortium**
MSU

[6] **Autonomy Challenges and Opportunities for the Special Forces.**
34th Annual NDIA SO/LIC Symposium

[5] **Artificial Intelligence: Projected S&T Advances in 2050 and Beyond**, *Army Virtual Workshop*

[4] **Novel Architectures for Neuro-Symbolic Computation**, *Electronics Resurgence Initiative (ERI) Summit*

[3] **Engineering Safe AI-Based Systems in Aerospace**, *2023 AIAA SciTech Forum*

[2] **Department of Defense Integrated AI and Autonomy Science and Technology Roadmap**, *AI Fusion Workshop at Carnegie Mellon University*

2019 [1] **Safe Autonomous Vehicles**, *AISafety 2019*

Tutorials

12/2025 **FEANS: Foundations and Emerging Applications of Neurosymbolic AI.**
39th Conference on Neural Information Processing Systems (**NeurIPS 2025**)

02/2023 **Advances in Neuro-Symbolic Reasoning.**
37th AAAI Conference on Artificial Intelligence (**AAAI 2023**)

Service

2025 ACM Emerging Interest Group (EIG) on Trustworthy and Responsible Systems (EIGTRUST).

2025 Co-founder for the NeurIPS workshop on Biosecurity Safeguards for Generative AI.

2025 Chair of Disruptive Ideas Track, Neuro-symbolic Systems (NeuS) conference.

2024 Founder and steering committee member for the Neuro-symbolic Systems (NeuS) conference.

2024 Member of the advisory board for the NASA University Leadership Initiative (ULI).

2023 Organizer for the workshop on Novel Architectures for Neuro-Symbolic Computation at the Electronics Resurgence Initiative (ERI) 2.0 Summit.

2023 Session lead for DARPA AI Forward.

2022 Session lead for Security in AI at the NSF Secure and Trustworthy CyberSpace (SaTC) Principal Investigators' Meeting.

Editorial Boards

2021–Present Associate Editor, IEEE Transactions on Artificial Intelligence.

2020 Editor, Air Force Research Laboratory Autonomy, Command and Control magazine. Copies available upon request for Department of Defense employees and U.S. DoD contractors only (distribution D).

Program Committees

2023 PC member for AAAI 2023, IJCAI 2023, AISTATS 2022, AAAI 2021, DAC 2020.

2020–2021 NATO exploratory group IST-ET-112 on Machine Learning Ecosystem for the Rapid Research, Development, and Deployment of Artificial Intelligence and Machine Learning Capabilities.

2020–2021 Air Force Research Laboratory Trusted AI challenge series.

2020–2021 Air Force Research Laboratory Innovare Aspire series.

2021 Technical program committee member, National Security Innovation Network (NSIN) Propel Series, a startup accelerator program targeting Department of Defense problems.

Reviewership

2019–2024 Reviewed over 300 proposals for the following DARPA and ARPA-H programs: AIQ, AIR, CASTLE, Vail, GARD, LwLL, ICS, CARCOSA, DIGIHEALS, UPGRADE, MOCHA, FoundSci, TIAMAT, ANSR, SciFy, E-BOSS, INGOTS, TRACTOR, PWND^2, ITM, SABER, BEST, the AI Forward workshops, the Information Innovation Office (I2O) office-wide BAA, the Tactical Technology Office (TTO) office-wide BAA, the AIxBTO bio AI intiative, STTRs, and the Young Faculty Award (YFA) program.

2022–2025 Reviewer for CVPR 2026, NeurIPS 2025, CVPR 2025, ICCV 2025, ICLR 2024, AAAI 2024, NeurIPS 2023, ICML 2023, NeurIPS 2022, ICASSP 2022, ICML 2022.

2020, 2021 Evaluation committee member for the MIT Lincoln Laboratories Autonomous Systems line of programs.

2019–2021 Proposal reviewer for the following Air Force Office of Scientific Research (AFOSR) programs: Agile Science of Test and Evaluation, Optimization and Discrete Mathematics, Cognitive and Computational Neuroscience, and Trusted Autonomy and Cyber.

2021 DoD National Defense Science and Engineering Graduate (NDSEG) fellowship evaluation committee.

Teaching and Mentoring

Teaching

Fall 2025 **CSCI 7000: Neurosymbolic AI**, University of Colorado Boulder. Student evaluation score: 4.6/5

Spring 2022 Research advisor for the course **DS198: Data Science Discovery Project**, University of California at Berkeley.

Spring 2022 Research advisor for the course **ENGIE 4800: Data Science Capstone and Ethics**, Columbia University

Fall 2021 Research advisor for the course **DS198: Data Science Discovery Project**, University of California at Berkeley.

Fall 2021 Research advisor for the course **CSE 544A Software engineering for External Clients**, Washington University in St. Louis.

Doctoral Advisees

Lekai Chen

Patrick Cooper

Alireza Nadali

Antony Zhao

Lucas White

Michael Buchanan

Doctoral Committees

2025 **Reinforcement Learning for Formal Specifications.**
Amin Falah. Advisor: Ashutosh Trivedi, University of Colorado Boulder.

2025 **Counterfactual Regret Minimization for Space Domain Awareness.**
Tyler Becker. Advisor: Zachary Sunberg, University of Colorado Boulder.

2024 **Neurosymbolic Video Understanding.**
Minkyu Choi. Advisor: Sandeep Chinchali, University of Texas Austin.

2024 **Reinforcement Learning for Combinatorial Optimization.**
Gabriel Maliakal. Advisor: Saiprasad Ravishankar, Michigan State University.

2024 **Correct By Design: Machine Learning in Safety Critical Devices.**
John Komp. Advisor: Ashutosh Trivedi, University of Colorado Boulder.

2024 **Enhancing Demand Modeling for Advanced Air Mobility Using Data-Driven Learning and Neurosymbolic AI.**
Kamal Acharya. Advisor: Houbing Song, University of Maryland Baltimore.

2023 **Neurosymbolic AI in Autonomous Systems: Towards Safer and More Transparent Deep Learning.**
Justus Renkhoff. Advisor: Houbing Song, University of Maryland Baltimore.

2022 **Bidirectional Synthesis in Neural Networks with Applications to Adversarial Attacks and Combinatorial Optimization.**
Ismail Alkhouri. Advisor: George Atia, University of Central Florida.

2021 **Improving Operational Machine Learning in the Pulmonary Radiology Domain Towards Clinical Adoption.**
Edward Verenich. Advisor: Faraz Hussein, Clarkson University.

Media

2025 **Physical AI: Moving From Bits to Atoms.**
CES, Link: https://www.ces.tech/videos/2025/january/physical-ai-moving-from-bits-to-atoms/

2025 **Built-in safeguards might stop AI from designing bioweapons.**
Kristel Tjandra
Science. https://www.science.org/content/article/built-safeguards-might-stop-ai-designing-bioweapons

2025 **Neurosymbolic AI could be leaner and smarter.**
Emma Marris
EurekAlert! https://www.eurekalert.org/news-releases/1084289

2025 **To 'democratize' AI, make it work more like a human brain.**
Daniel Strain
CU Boulder Today https://www.colorado.edu/today/2025/06/05/democratize-ai-make-it-work-more-human-brain

2024 **Autonomous Cars Race On The Las Vegas Speedway At CES 2025**
Sabbir Rangwala
Forbes, Link: https://www.forbes.com/sites/sabbirrangwala/2024/10/28/autonomous-cars-race-on-the-las-vegas-speedway-at-ces-2025/

2024 **Harnessing The Power Of AI (and Physics) In Movement Automation**
Sabbir Rangwala
Forbes, Link: https://www.forbes.com/sites/sabbirrangwala/2024/10/02/harnessing-the-power-of-ai-and-physicsin-movement-automation/

2024 **AI Driver Training from SIM to REAL.**
CES, Link: https://videos.ces.tech/detail/videos/cta-stage/video/6345006827112/ai-driver-training-from-sim-to-real?autoStart=true&sort_by=DISPLAY_NAME%$3AASC

2023 **Pentagon experiments find generative AI easy to exploit.**
Ryan Lovelace
The Washington Times, Link: https://www.washingtontimes.com/news/2023/nov/3/pentagon-experiments-find-generative-ai-easy-explo/

2023 **Inside DARPA's search for an "autonomous scientist" to support its researchers.**
NextGov, Link: https://www.nextgov.com/artificial-intelligence/2023/11/inside-darpas-search-autonomous-scientist-support-its-researchers/392139/

2022 **Episode 62: The Model (& Simulation) Researcher.**
DARPA Podcast, YouTube: https://www.youtube.com/watch?v=3-vK1WAJ0x8, iTunes: https://podcasts.apple.com/us/podcast/voices-from-darpa/id1163190520

2021 **Autonomy, Command and Control Portfolio and Strategy.**
Innovare Aspire Summit, https://youtu.be/AxLCaxCD1k8?t=6237

2021 **Verification of Autonomous Systems.**
AFRL Trusted AI Challenge Series 2: Grounding the Critical Path, https://youtu.be/-49zeVZQ4k8

2021 **Ask Me Anything on Transfer Learning for Control Policies and Hierarchical Planning.**
Innovare Aspire Series, https://youtu.be/-63cq5bVkMA