robert@valgo.ai

ROBERT J. MOSS

www.robert-moss.com

San Mateo, CA

Research Interests

Interested in the development of algorithms for safe planning in high-dimensional, long-horizon, uncertain environments modeled as partially observable Markov decision processes (POMDPs). Developing algorithms to efficiently validate safety-critical systems such as aircraft collision avoidance, autonomous vehicles, and sustainable energy production.

Table of Contents

Education

Doctor of Philosophy (Ph.D.) in Computer Science (Artificial Intelligence) [2021 – 2025]

Stanford University, Stanford, CA

Advisor: Mykel J. Kochenderfer | Thesis: ''Safe planning under uncertainty using surrogate models'' (Defense Video)

Master of Science (M.S.) in Computer Science (Artificial Intelligence) [2019 – 2021]

Stanford University, Stanford, CA

Advisor: Mykel J. Kochenderfer | Thesis: ''Algorithms for efficient validation of black-box systems'' [20]

Bachelor of Science (B.S.) in Computer Science, Minor in Physics [2010 – 2014]

Wentworth Institute of Technology, Boston, MA

Advisor: James G. O'Brien | Magna Cum Laude

Research

Stanford Intelligent Systems Laboratory (SISL), *Stanford University* [2019 – Present]

Constrained POMDP planning using learned safety surrogates: Developed algorithms for safe planning in high-dimensional, partially observable environments, applied to carbon capture and storage (CCS) [26, 27].

Belief-state planning for long-horizon POMDPs: Developed algorithm to plan in high-dimensional, long-horizon POMDPs that learns approximately optimal surrogates of the value function and policy to replace heuristics in expensive Monte Carlo tree searches, applied to critical mineral exploration [25].

Survey of black-box safety validation algorithms: Co-authored literature survey of algorithms for black-box safety validation of cyber-physical systems, analyzing and distilling over 100 publications [19].

Probabilistic risk assessment of autonomous vehicles: Led the development of an autonomous vehicle risk assessment framework using risk analysis methods from the financial industry, demonstrated on 2D driving simulators and the 3D driving simulator CARLA [18]. Developed an algorithm to predict when failures will occur to increase failure rate of an autonomous vehicle [17] and developed an algorithm to detect weaknesses in autonomous vehicles [13].

Emergency evacuations under compounding levels of uncertainty: Developed a sequential decision support tool for emergency evacuation procedures and analyzed the sensitivity to different level of uncertainty. Demonstrated on a case study of the 2021 Afghanistan evacuation [21].

Stress testing autonomous systems: Developed reinforcement learning algorithm to stress test aircraft trajectory predictions in a flight management system (FMS) [14, 16]. Supported the use of the developed methods for certification of airborne software [15].

Surrogate model-based optimization: Developed stochastic optimization algorithms based on the cross-entropy method that use surrogates to find rare failure events in computationally expensive systems [12].

Stanford Doerr School of Sustainability, *Stanford University* [2022 – Present]

Sequential optimization of geothermal energy production: Developing POMDP to optimize geothermal energy production based on a real-world case study dealing with large action spaces and small economical margins.

Safe carbon capture and storage (CCS): Supported the development of a POMDP for safe CO_2 subsurface storage using a neural network surrogate model of CO_2 plume migration dynamics.

Model-fidelity sensitivity analysis for POMDPs: Studied the effect of model-fidelity on planning for critical mineral exploration and developed a general POMDP model-fidelity analysis framework [31].

Sequential decision making for critical mineral exploration: Supported the development of a belief-state MDP model of critical mineral exploration using reward shaping for more efficient and informed decision making.

Xwing, *Concord, CA* [2022 – 2023]

Failure probability estimation: Developed method using Bayesian optimization to falsify safety-critical subsystems of an autonomous aircraft and to estimate the probability of system failure using importance sampling [24]. Applied to a neural network-based runway detection system in simulation.

Certification of machine learning avionics systems: Supported the development of formal and practical considerations for the certification of machine learning components in aircraft [23]. This research is currently being used to support FAA certification of an autonomous cargo aircraft.

NASA Ames Research Center, *Moffett Field, CA* [2020 – 2021]

Autonomous rover validation: Applied safety validation techniques to stress test the decision making system of the autonomous lunar rover used in the VIPER mission, searching for water deposits on the Moon.

MIT Lincoln Laboratory, *Lexington, MA* [2013 – 2019]

Traffic advisory optimization: Researched global optimization methods to tune the traffic advisories (TAs) of the next-generation aircraft collision avoidance system (ACAS Xa) and developed a tool to efficiently optimize the TA logic over aircraft near-collision encounters from US and European airspace.

Pilot response modeling: Developed modeling technique using Bayesian networks that predicts the probability that an onboard pilot will respond to collision avoidance resolution advisory (RA) maneuvers [6, 7].

Autonomous safety maneuvering: Analyzed the safety of ACAS Xu on various vertical rate and turn rate limits which led the FAA to require all UAS equipped with ACAS Xu to autonomously respond to collision avoidance advisories.

Large-scale safety validation: Validated ACAS Xa and ACAS Xu on billions of realistic collision avoidance trajectories using parallel processing resources from the Lincoln Laboratory Supercomputing Center to support FAA certification of ACAS X. Developed a tool for document generation of safety and operational suitability metrics.

Algorithm specification languages: Studied the use of various programming languages to be used as a specification language for ACAS X to replace legacy pseudocode. Implemented ACAS X in the Julia programming language, enabling line-by-line evaluation of the published specification. Developed an automated tool to generate the 700+ page algorithm design document.

Wildfire resource allocation: Developed a wildfire simulator and decision support tool for wildfire incident commanders to optimize the placement of aerial and land water resources to efficiently extinguish wildfires [5, 2]. Implemented realistic wildfire dynamics models and heuristics for better decisions using Monte Carlo tree search.

Wentworth Institute of Technology, *Boston, MA* [2013 – 2014]

Galactic rotation curve modeling: Developed modeling and simulation tool to study the fit of gravitational theories to observational data of galactic rotation [1, 3, 4, 8]. Implemented the χ^2 test for statistical hypothesis testing.

Industry Experience

Valgo, *San Mateo, CA* [Nov. 2025 – Present]
Co-founder and CEO
Building risk quantification tools in simulation to provide insurance for physical AI systems (Y Combinator W26).

Xwing, *Concord, CA* [Jun. 2022 – Jun. 2023]
AI Safety and DAA Consultant, PhD Student Research Intern
Developed statistical validation method for failure probability estimation of safety-critical systems. Consulted on detect-and-avoid (DAA) algorithm integration and the certification of machine learning components for autonomous aircraft.

NASA Ames Research Center, *Moffett Field, CA* [Oct. 2020 – Jun. 2021]
Research Engineer
Developed interactive lunar rover traverse GUI for NASA's VIPER mission, searching for water deposits on the Moon.

MIT Lincoln Laboratory, *Lexington, MA* [Jan. 2013 – Aug. 2019]
Associate Staff
Worked on the ACAS Xa, ACAS Xu, and ACAS sXu teams (collision avoidance systems for commercial aircraft, unmanned aircraft, and small UAS, respectively). Worked on the development, analysis, optimization, and validation of ACAS X in support of FAA certification. Managed the software development team for ACAS Xu, written in the Julia programming language. Led data collection and analysis for a limited implementation of ACAS Xa on experimental aircraft flying in US airspace. Led data collection for the ACAS Xu unmanned aircraft flight test collaborating with NASA Armstrong. Co-chaired the RTCA ACAS Xu Threat Working Group, guiding discussion and analysis for changes to the threat logic. Developed tool to auto generated the algorithm design specification for ACAS Xa and Xu written in Julia. Built an aggregate data visualization tool using D3.js for the international aviation safety community, distributed through RTCA.

Harvard University, *Cambridge, MA* [May 2012 – Dec. 2012]
Technical Support (IT) Co-op
Built a diagnostic system for all of Harvard's computer labs and provided IT support for the Faculty of Arts and Sciences.

Awesome Products, LLC, *Cambridge, MA* [May 2012 – Apr. 2014]
Co-owner, Lead iOS Developer, Android Developer
Handled high-level programming design and development of Android and iPhone music making apps. Created *LottoPilot* for iOS to track lottery ticket spending to help people mitigate future purchases. Led discussion on business advice in meeting with Daymond John (investor on *Shark Tank*) and secured two funding rounds via the WIT Accelerate Program.

Jabez International, LLC, *Cambridge, MA* [Feb. 2012 – Sep. 2014]
Freelance Java Developer
Designed, illustrated, and coded a children's computerized board game.

Awards and Recognitions

Lance Stafford Larson Student Writing Award
For the paper ''Bayesian Safety Validation for Failure Probability Estimation of Black-Box Systems'' [28]
IEEE Computer Society, 2023

Best-of-Conference Award
For the paper ''Formal and Practical Elements for the Certification of Machine Learning Systems'' [23]
AIAA/IEEE Digital Avionics Systems Conference (DASC), 2023

Best-of-Track Award in Unmanned Aircraft Systems & Advanced Air Mobility (UAS/AAM)
For the paper ''Formal and Practical Elements for the Certification of Machine Learning Systems'' [23]
AIAA/IEEE Digital Avionics Systems Conference (DASC), 2023

Best-of-Session Award in Artificial Intelligence/Machine Learning (AI/ML)
For the paper ''Formal and Practical Elements for the Certification of Machine Learning Systems'' [23]
AIAA/IEEE Digital Avionics Systems Conference (DASC), 2023

R&D 100 Award for the Collision Avoidance System for Small Uncrewed Aircraft Systems (ACAS sXu)
Research and Development World, 2022

Christofer Stephenson Memorial Award
For the best CS master's thesis titled ''Algorithms for Efficient Validation of Black-Box Systems'' [20]
Stanford University, 2021

Centennial Teaching Assistant (TA) Award
For teaching excellence in CS238/AA228: Decision Making Under Uncertainty
Stanford University, 2021

Best-of-Session Award in Safe & Secure Technologies
For the paper ''Certification Considerations for Adaptive Stress Testing of Airborne Software'' [15]
AIAA/IEEE Digital Avionics Systems Conference (DASC), 2021

First Place Student Research Award
For the paper ''Adaptive Stress Testing of Trajectory Predictions in Flight Management Systems'' [14]
AIAA/IEEE Digital Avionics Systems Conference (DASC), 2020

Best-of-Session Award in Validation & Verification
For the paper ''Adaptive Stress Testing of Trajectory Predictions in Flight Management Systems'' [14]
AIAA/IEEE Digital Avionics Systems Conference (DASC), 2020

Best-of-Track Award in Unmanned Aircraft Systems (UAS)
For the paper ''ACAS-Xu: Integrated Collision Avoidance and Detect and Avoid Capability for UAS'' [10]
AIAA/IEEE Digital Avionics Systems Conference (DASC), 2019

Best-of-Track in Safety and Resilience
For the paper ''A Bayesian Network Model of Pilot Response to TCAS Resolution Advisories'' [6]
FAA/EUROCONTROL Air Traffic Management Research and Development Seminar (ATM R&D Seminar), 2017

R&D 100 Award for the Next-Gen. Unmanned Airborne Collision Avoidance System (ACAS Xu)
Research and Development World, 2016

MIT Lincoln Laboratory Team Award for the Next-Gen. Airborne Collision Avoidance System (ACAS X)
MIT Lincoln Laboratory, 2015–2016

Magna Cum Laude
Wentworth Institute of Technology, 2014

Teaching and Leadership

CS238V/AA228V: Validation of Safety-Critical Systems [2025 – Present]
Stanford University, Stanford, CA
Co-developed the course alongside the instructor, Sydney Katz, and was the head teaching assistant in 2025. The course follows our textbook *Algorithms for Validation* and covers topics including: system modeling, temporal logic specifications, optimization-based falsification, Markov chain Monte Carlo, importance sampling, reachability analysis, model checking, satisfiability, and explainability.

CS238/AA228: Decision Making Under Uncertainty [2019 – 2024]
Stanford University, Stanford, CA
Head teaching assistant in 2019 (350+ students) and TA in 2020 and 2021. Led programming projects for Bayesian network structure learning and reinforcement learning algorithms, advised student research projects, answered detailed forum questions, and led section on using beliefs for POMDP state estimation.

Tutorial on Decision Making Under Uncertainty [2022]
Center for Automotive Research at Stanford (CARS)
Stanford University, Stanford, CA
Led mini-class on how to build and solve sequential decision making problems in uncertain environments.

Stanford Center for AI Safety Project Lead [2020 – 2022]
Stanford University, Stanford, CA
Led autonomous vehicle risk assessment research project across three Stanford labs: the Stanford Intelligent Systems Laboratory (SISL), the Navigation and Autonomous Vehicles (NAV) Laboratory, and the Autonomous Systems Laboratory (ASL). Wrote research proposal and led the bi-weekly meetings interfacing with research sponsors.

Summer Undergraduate Research Fellowship (SURF) Mentor [2021]
Stanford University, Stanford, CA
Developed research agenda and mentored student from Harvard University on a deep reinforcement learning project studying the sim-to-real gap.

Julia Academy [2021]
JuliaHub, Inc. (formerly Julia Computing, Inc.)
Course creator and lecturer for ''Decision Making Under Uncertainty using POMDPs.jl''.

Volunteering

Simulation Decomposition Organization [2023 – Present]
Developing open-source Julia package for sensitivity analysis using simulation decomposition (SimDec).

Reviewer for Journals and Conferences [2018 – Present]

Journal of Artificial Intelligent Research (JAIR)
AI Access Foundation

Machine Learning Journal
Springer

Journal of Aerospace Information Systems (JAIS)
American Institute of Aeronautics and Astronautics (AIAA)

International Conference on Intelligent Robots and Systems (IROS)
Institute of Electrical and Electronics Engineers (IEEE)

Reinforcement Learning Conference (RLC)
Reinforcement Learning Journal (RLJ)

Textbook Proofreader [2013 – Present]

Algorithms for Decision Making, M. J. Kochenderfer, T. A. Wheeler, and Kyle H. Wray, *MIT Press*, 2022.

Algorithms for Optimization, M. J. Kochenderfer and T. A. Wheeler, *MIT Press*, 2019.

Decision Making Under Uncertainty: Theory and Application, M. J. Kochenderfer, *MIT Press*, 2015.

Science and Engineering Fair Judge [2015 – 2019]

Massachusetts State Science and Engineering Fair, *Massachusetts Institute of Technology (MIT), Cambridge, MA*

Massachusetts State Middle School Science and Engineering Fair, *Worcester Technical High School, Worcester, MA*

Worcester Regional Science and Engineering Fair, *Worcester Middle School, Worcester, MA*

Invited Talks

BetaZero: Belief-State Planning for Long-Horizon POMDPs using Learned Approximations [2024]
Oral Track: Multi-Agent RL and Planning Algorithms
Reinforcement Learning Conference (RLC), Amherst, MA

Chance-Constrained POMDP Planning with Learned Neural Network Surrogates [2024]
IJCAI Workshop on Trustworthy Interactive Decision-Making with Foundation Models
International Joint Conference on Artificial Intelligence (IJCAI), Jeju, South Korea

ConstrainedZero: Chance-Constrained POMDP Planning Using Learned Probabilistic Failure Surrogates [2024]
Main Track
International Joint Conference on Artificial Intelligence (IJCAI), Jeju, South Korea

My Research Journey: Making Decisions about Decision Making [2023]
Lappeenranta University of Technology (LUT), Lappeenranta, Finland

Bayesian Safety Validation for Black-Box Systems [2023]
Modeling and Simulation Technologies
AIAA AVIATION Forum, San Diego, CA

Learn Offline Approximations for Large-Scale Sequential Decision Making Problems [2023]
SCERF Affiliates Meeting
Stanford Center for Earth Resources Forecasting (SCERF), Stanford, CA

Publishing Software Frameworks in the Journal of Open Source Software (JOSS) [2023]
Stanford SystemX Alliance
Stanford University, Stanford, CA

Combining Mixed-Fidelity Model Approximations for Efficient Planning [2022]
SCERF Affiliates Meeting
Stanford Center for Earth Resources Forecasting (SCERF), Stanford, CA

Mixed-Fidelity Model Approximations [2022]
Introductory Research Review
KoBold Metals, Stanford, CA

Stanford Intelligent Systems Laboratory (SISL): An Overview [2022]
Stanford Center for Earth Resources Forecasting (SCERF), Stanford, CA

Quotes and Lessons From "Letters to a Young Scientist" by Edward O. Wilson [2022]
CS239/AA229: Advanced Topics in Sequential Decision Making
Stanford University, Stanford, CA

An Efficient Framework for Modular Autonomous Vehicle Risk Assessment (MAVRA) [2022]
Workshop on Safety Validation of Connected and Automated Vehicles
IEEE International Conference on Intelligent Transportation Systems (ITSC)

Transferring Aviation Safety Lessons to the Road [2021]
Automated Road Transportation Symposium
The National Academies of Sciences, Engineering, and Medicine

Adaptive Stress Testing of Trajectory Predictions in Flight Management Systems [2020]
Validation and Verification
AIAA/IEEE Digital Avionics Systems Conference (*DASC*)

Simplified VAPS Algorithm for Online Stigmergic Policies [2020]
CS239/AA229: Advanced Topics in Sequential Decision Making
Stanford University, Stanford, CA

A Bayesian Network Model of Pilot Response to TCAS RAs [2017]
Safety and Resilience
Air Traffic Management Research and Development Seminar, Seattle, WA

A Decision Theoretic Approach to Future Aircraft Collision Avoidance [2017]
Data Science Tea
UMass Amherst Data Science Seminar, Amherst, MA

Using Julia as a Specification Language for the Safety-Critical Systems [2016]
Safety and Validation
nuTonomy, Cambridge, MA

Using Julia as a Specification Language for the Next-Gen. Airborne Collision Avoidance System [2015]
Scientific Applications
JuliaCon, Cambridge, MA

Models of the Rotation Curve for the Milky Way Galaxy [2014]
Undergraduate Research and Outreach
American Physical Society (*APS*), Savannah, GA

Publications

[31] **Model-Fidelity Analysis for Sequential Decision-Making Systems using Simulation Decomposition: Case Study of Critical Mineral Exploration**
R. J. Moss, M. Kozlova, A. Corso, and J. Caers
Routledge, 2024

[30] **SimDec Algorithm and Guidelines for its Usage and Interpretation**
M. Kozlova, **R. J. Moss**, P. Roy, A. Alam and J. S. Yeomans
Routledge, 2024

[29] **Kov: Transferable and Naturalistic Black-Box LLM Attacks Using Markov Decision Processes and Tree Search**
R. J. Moss
arxiv 2408.08899, 2024
Code: https://github.com/sisl/Kov.jl

[28] **Bayesian Safety Validation for Failure Probability Estimation of Black-Box Systems**
R. J. Moss, M. J. Kochenderfer, M. Gariel, and A. Dubois
AIAA Journal of Aerospace Information Systems (*JAIS*), 2024
Code: https://github.com/sisl/BayesianSafetyValidation.jl

[27] **Chance-Constrained POMDP Planning with Learned Neural Network Surrogates**
R. J. Moss, A. Jamgochian, J. Fischer, A. Corso, and M. J. Kochenderfer
IJCAI Workshop on Trustworthy Interactive Decision-Making with Foundation Models, 2024
Code: https://github.com/sisl/ConstrainedZero.jl

[26] **ConstrainedZero: Chance-Constrained POMDP Planning Using Learned Probabilistic Failure Surrogates and Adaptive Safety Constraints**
R. J. Moss, A. Jamgochian, J. Fischer, A. Corso, and M. J. Kochenderfer
International Joint Conference on Artificial Intelligence (*IJCAI*), 2024
Code: https://github.com/sisl/ConstrainedZero.jl

[25] **BetaZero: Belief-State Planning for Long-Horizon POMDPs using Learned Approximations**
R. J. Moss, A. Corso, J. Caers, and M. J. Kochenderfer
Reinforcement Learning Journal (*RLJ*), 2024
Code: https://github.com/sisl/BetaZero.jl

[24] **Bayesian Safety Validation for Black-Box Systems**
R. J. Moss, M. J. Kochenderfer, M. Gariel, and A. Dubois
AIAA AVIATION Forum, 2023
Code: https://github.com/sisl/BayesianSafetyValidation.jl

[23] **Formal and Practical Elements for the Certification of Machine Learning Systems**
J.-G. Durand, A. Dubois, and **R. J. Moss**
AIAA/IEEE Digital Avionics Systems Conference (*DASC*), 2023

[22] **Uncovering Heterogeneous Effects in Computational Models for Sustainable Decision-Making**
M. Kozlova, **R. J. Moss**, J. S. Yeomans, and J. Caers
Environmental Modelling & Software (*under review*), 2023

[21] **Prioritizing Emergency Evacuations under Compounding Levels of Uncertainty**
L. J. Einstein, **R. J. Moss**, M. J. Kochenderfer
IEEE Global Humanitarian Technology Conference (*GHTC*), 2022
Code: https://github.com/sisl/EvacuationPOMDP.jl

[20] **Algorithms for Efficient Validation of Black-Box Systems**
R. J. Moss
Stanford University, M.S. Thesis, 2021

[19] **A Survey of Algorithms for Black-Box Safety Validation of Cyber-Physical Systems**
A. Corso, **R. J. Moss**, M. Koren, R. Lee, and M. J. Kochenderfer
Journal of Artificial Intelligence Research (*JAIR*), 2021

[18] **Autonomous Vehicle Risk Assessment**
R. J. Moss, S. Gupta, R. Dyro, K. Leung, M. J. Kochenderfer, G. X. Gao, M. Pavone, E. Schmerling, et al.
Stanford Center for AI Safety, 2021
Code: https://github.com/sisl/AutonomousRiskFramework.jl

[17] **Predictive Risk for Efficient Black-Box Validation of Autonomous Vehicles**
R. J. Moss
Stanford University, CS229: Machine Learning, 2021

[16] **POMDPStressTesting.jl: Adaptive Stress Testing for Black-Box Systems**
R. J. Moss
Journal of Open Source Software (*JOSS*), 2021
Code: https://github.com/sisl/POMDPStressTesting.jl

[15] **Certification Considerations for Adaptive Stress Testing of Airborne Software**
M. Durling, H. H. Zapana, B. Meng, M. Meiners, J. Hochwarth, N. Visser, R. Lee, **R. J. Moss**, and V. T. Valapil
IEEE/AIAA Digital Avionics Systems Conference (*DASC*), 2021

[14] **Adaptive Stress Testing of Trajectory Predictions in Flight Management Systems**
R. J. Moss, R. Lee, N. Visser, J. Hochwarth, J. G. Lopez, and M. J. Kochenderfer
AIAA/IEEE Digital Avionics Systems Conference (*DASC*), 2020

[13] **Adversarial Weakness Recognition for Efficient Black-Box Validation**
R. J. Moss
Stanford University, CS230: Deep Learning, 2020
Code: https://github.com/sisl/WeaknessRecognition.jl

[12] **Cross-Entropy Method Variants for Optimization**
R. J. Moss
arxiv 2009.09043, 2020
Code: https://github.com/mossr/CrossEntropyVariants.jl

[11] **Automated Trash Collection using Markov Decision Processes**
R. J. Moss
Stanford University, CS221: Artificial Intelligence: Principles and Techniques, 2020
Code: https://github.com/mossr/TrashMDP.jl

[10] **ACAS Xu: Integrated Collision Avoidance and Detect and Avoid Capability for UAS**
M. P. Owen, A. Panken, **R. J. Moss**, L. Alvarez, and C. Leeper
AIAA/IEEE Digital Avionics Systems Conference (*DASC*), 2019

[9] **A Parallel Simulation Approach to ACAS X Development**
A. Gjersvik and **R. J. Moss**
IEEE High Performance Extreme Computing Conference (*HPEC*), 2019

[8] **Alternative Gravity Rotation Curves for the LITTLE THINGS Survey**
J. G. O'Brien, T. L. Chiarelli, J. Dentico, M. Stulge, B. Stefanski, **R. J. Moss**, and S. Chaykov
The Astrophysical Journal, 2018

[7] **Bayesian Network Model of Pilot Response to Collision Avoidance System Resolution Advisories**
E. H. Londner and **R. J. Moss**
Journal of Air Transportation (*JAT*), 2018

[6] **A Bayesian Network Model of Pilot Response to TCAS Resolution Advisories**
E. H. Londner and **R. J. Moss**
Air Traffic Management Research and Development Seminar (*ATM R&D Seminar*), 2017

[5] **Automated Dynamic Resource Allocation for Wildfire Suppression**
J. D. Griffith, M. J. Kochenderfer, **R. J. Moss**, V. V. Mišić, V. Gupta, and D. Bertsimas
Lincoln Laboratory Journal, 2017

[4] **Recent Advancements in Conformal Gravity**
J. G. O'Brien, S. S. Chaykov, J. Dentico, M. Stulge, B. Stefanski, and **R. J. Moss**
Journal of Physics: Conference Series, 2017

[3] **Rotation Curve for the Milky Way Galaxy in Conformal Gravity**
J. G. O'Brien and **R. J. Moss**
Journal of Physics: Conference Series, 2015

[2] **A Comparison of Monte Carlo Tree Search and Mathematical Optimization for Large Scale Dynamic Resource Allocation**
D. Bertsimas, J. D. Griffith, V. Gupta, M. J. Kochenderfer, V. V. Mišić, and **R. J. Moss**
arXiv 1405.5498, 2014

[1] **Rotation Curve Modeler: A Modeling and Simulation Tool for Arbitrary Galaxies**
R. J. Moss and J. G. O'Brien
Wentworth Institute of Technology, 2014
Code: https://github.com/RoCMSOCM/RoCM

Textbooks

Algorithms for Validation [2025]
Mykel J. Kochenderfer, Sydney M. Katz, Anthony L. Corso, and **Robert J. Moss**
MIT Press

Course Notes

CS109 Notes: Probability for Computer Scientists [2020]
Stanford University, Stanford, CA

CS229 Notes: Machine Learning [2021]
Stanford University, Stanford, CA

CS234 Notes: Reinforcement Learning [2021]
Stanford University, Stanford, CA

CS221 Notes: Algorithms for Artificial Intelligence [2020]
Stanford University, Stanford, CA

Markov Decision Processes (MDPs) [2021]
Decision Making Under Uncertainty using POMDPs.jl
Julia Academy

Partially Observable Markov Decision Processes (POMDPs) [2021]
Decision Making Under Uncertainty using POMDPs.jl
Julia Academy

Beliefs: State Uncertainty [2020]
CS238/AA228: Decision Making Under Uncertainty
Stanford University, Stanford, CA

Reinforcement Learning Algorithms and Equations [2020]
Stanford University, Stanford, CA

Markov Decision Process: Chain Rule [2020]
Stanford University, Stanford, CA

Loss Functions in Machine Learning [2020]
Generated using TeX.jl
Stanford University, Stanford, CA

Unconstrained Optimization [2020]
CS361/AA222: Engineering Design Optimization
Stanford University, Stanford, CA

Constrained Optimization [2020]
CS361/AA222: Engineering Design Optimization
Stanford University, Stanford, CA

Constrained Optimization and Expression Optimization [2020]
CS361/AA222: Engineering Design Optimization
Stanford University, Stanford, CA

Implementation: Learning Policies with External Memory (VAPS) [2020]
CS239/AA229: Advanced Topics in Sequential Decision Making
Stanford University, Stanford, CA

Deriving the Quadratic Formula [2020]
Generated using TeX.jl
Stanford University, Stanford, CA

Open-Source Code

Open-source repositories for research and teaching, primarily written in Julia: https://github.com/mossr

Research Code

BetaZero.jl [2023 – Present]
Belief-state planning for POMDPs using learned approximations; integrated into POMDPs.jl.
https://github.com/sisl/BetaZero.jl

BayesianSafetyValidation.jl [2022 – Present]
Estimate probability of failure using reframed Bayesian optimization.
https://github.com/sisl/BayesianSafetyValidation.jl

POMDPStressTesting.jl [2020 – Present]
Adaptive stress testing of black-box systems within POMDPs.jl.
https://github.com/sisl/POMDPStressTesting.jl

POMDPModelFidelityFramework.jl [2021 – Present]
Framework for sensitivity analysis of POMDP model fidelities.
https://github.com/sisl/POMDPModelFidelityFramework.jl

SimulationDecomposition.jl [2023 – Present]
Sensitivity analysis using simulation decomposition (SimDec) in Julia.
https://github.com/Simulation-Decomposition/SimulationDecomposition.jl

SignalTemporalLogic.jl [2022 – Present]
Signal temporal logic (STL) formulas and robustness in Julia.
https://github.com/sisl/SignalTemporalLogic.jl

PropositionalLogic.jl [2020]
Simple propositional logic in Julia.
https://github.com/mossr/PropositionalLogic.jl

AutonomousRiskFramework.jl [2021 – 2023]
Framework for autonomous vehicle risk assessment using CARLA.
https://github.com/sisl/AutonomousRiskFramework.jl

EvacuationPOMDP.jl [2022]
Prioritizing emergency evacuations using POMDPs: A case study of the Afghanistan evacuation.
https://github.com/sisl/EvacuationPOMDP.jl

WeaknessRecognition.jl [2020]
Black-box validation through encoded failure representation.
https://github.com/sisl/WeaknessRecognition.jl

GaussianDiscriminantAnalysis.jl [2021]
Gaussian discriminant analysis in Julia.
https://github.com/mossr/GaussianDiscriminantAnalysis.jl

CrossEntropyVariants.jl [2020]
Cross-entropy method variants for optimization in Julia.
https://github.com/mossr/CrossEntropyVariants.jl

TrashMDP.jl [2020]
Automated trash collection using Markov decision processes.
https://github.com/mossr/TrashMDP.jl

Teaching Code

BeautifulAlgorithms.jl [2020 – Present]
Concise and beautiful algorithms written in Julia (1.3k stars).
https://github.com/mossr/BeautifulAlgorithms.jl

machine_learning_book [2021]
Stanford's CS229 Machine Learning lecture notes compiled into a Tufte-style textbook.
https://github.com/mossr/machine_learning_book

StateEstimation.jl [2020]
Belief state estimation for Stanford's CS238/AA228 Decision Making Under Uncertainty.
https://github.com/mossr/StateEstimation.jl

Gradescope.jl [2020]
Julia interface for Gradescope autograding.
https://github.com/sisl/Gradescope.jl

RedPen.jl [2019]
Grading project assignments in Julia.
https://github.com/sisl/RedPen.jl

Obfuscatee.jl [2019]
Obfuscate Julia source code.
https://github.com/sisl/Obfuscatee.jl

julia-tufte-beamer [2020]
Tufte-style beamer template with Julia integration.
https://github.com/mossr/julia-tufte-beamer

julia-mono-listings [2020]
LaTeX listings style for Julia and Unicode support for the JuliaMono font.
https://github.com/mossr/julia-mono-listings

PlutoNotebooks [2020 – Present]
Collection of Pluto notebooks.
https://github.com/mossr/PlutoNotebooks

TikzNeuralNetworks.jl [2020]
Visualize neural networks using Ti*k*Z in Julia.
https://github.com/mossr/TikzNeuralNetworks.jl

TeX.jl [2020]
tex macro for generating LaTeX PDFs from Julia code with descriptions.
https://github.com/mossr/TeX.jl

Miscellaneous Code

Tetris.jl [2020]
One-line Tetris in the Julia REPL.
https://github.com/mossr/Tetris.jl

Snake.jl [2020]
The game of snake in the Julia REPL.
https://github.com/mossr/Snake.jl

Donut.jl [2021]
Andy Sloane's rotating donut in Julia.
https://github.com/mossr/Donut.jl

NoisyLandscapes.jl [2020]
Generate landscapes using Simplex noise in Julia.
https://github.com/mossr/NoisyLandscapes.jl

Organizations

Stanford Intelligent Systems Laboratory (SISL) [2019 – Present]
Part of the Stanford Artificial Intelligence Laboratory (SAIL)
Stanford University, Stanford, CA

Stanford Center for AI Safety [2020 – Present]
Stanford School of Engineering
Stanford University, Stanford, CA

Stanford Institute for Human-Centered Artificial Intelligence (HAI) [2022 – Present]
Stanford School of Engineering
Stanford University, Stanford, CA

Stanford Mineral-X Group [2023 – Present]
Stanford Doerr School of Sustainability
Stanford University, Stanford, CA

Stanford Center for Earth Resources Forecasting (SCERF) [2022 – Present]
Stanford Doerr School of Sustainability
Stanford University, Stanford, CA

Institute of Electrical and Electronics Engineers (IEEE) [2023 – Present]
IEEE Graduate Student Member

American Institute of Aeronautics and Astronautics (AIAA) [2023 – Present]
AIAA Student Member

American Physical Society (APS) [2014]
APS Student Member

RTCA, Inc. [2013 – 2019]
Member and Working Group Co-Chair
Washington, DC

Technical Skills

Programming languages: Julia, MATLAB, Python, JavaScript, C++, Java, LaTeX, PGF/Ti*k*Z, POMDPs.jl

Technical skills: Sequential decision making under uncertainty, machine learning, deep learning, optimization, validation of safety-critical systems, Bayesian networks, probability and statistics, modeling and simulation, algorithm development, open-source development, code golf.

Biography

Robert Moss is a computer science Ph.D. student at Stanford University studying algorithms for safe planning and algorithms to validate safety-critical autonomous systems. He holds an M.S. in computer science from Stanford where his research received the best computer science master's thesis award and he also received the Centennial TA award for his teaching efforts. He earned his B.S. in computer science with a minor in physics from the Wentworth Institute of Technology in Boston, MA. Robert was an associate research staff member at MIT Lincoln Laboratory where he was on the team that designed, developed, and validated the next-generation aircraft collision avoidance system for commercial aircraft, unmanned vehicles, and rotorcraft. Robert was also a research engineer at the NASA Ames Research Center, developing decision support tools for the VIPER autonomous Lunar rover mission searching for water deposits on the Moon. Robert is a member of the Stanford Intelligent Systems Laboratory, the Stanford Mineral-X Group, and part of the Stanford Center for AI Safety conducting research on methods for high-dimensional planning under uncertainty using low-dimensional surrogate models, risk assessment of autonomous vehicles, and efficient safety validation.

External Links

- CV: https://bit.ly/moss-cv
- Résumé: https://bit.ly/moss-resume
- Website: https://robert-moss.com
- GitHub: https://github.com/mossr
- Google Scholar: https://bit.ly/moss-scholar
- LinkedIn: https://www.linkedin.com/in/robert-j-moss