Sandhya Saisubramanian

School of Electrical Engineering and Computer Science
Oregon State University, USA

✉ sandhya.sai@oregonstate.edu   •   🕮 www.sandhyasai.com

Research Interests

Safe and Reliable Artificial Intelligence (AI); Foundations and Applications of Automated Planning; Transparency and Fairness in AI-Assisted Decision-Making.

Academic and Professional Experience

Oregon State University, USA   **Dec 2021-**
Assistant Professor (Tenure Track)
School of Electrical Engineering and Computer Science

NASA Ames Research Center (KBR Wyle), USA   **Jun 2019-Aug 2019**
Research Intern

University of Massachusetts, Amherst, USA   **Sep 2015-Nov 2021**
Research Assistant

Singapore Management University, Singapore   **May 2013-Jun 2015**
Research Engineer

National University of Singapore, Singapore   **Jan 2013-Apr 2013**
Graduate Student Researcher

Aricent, India   **Sep 2011- June 2012**
Software Engineer

Education

University of Massachusetts Amherst, USA   **Feb 2022**
Ph.D. in Computer Science   *(Thesis defended in Sep 2021)*
Dissertation Title: *Reliable Decision-Making with Imprecise Models*
Chair: Shlomo Zilberstein

National University of Singapore, Singapore   **Feb 2014**
Master of Computing

Pondicherry Engineering College, India   **May 2011**
Bachelor of Technology in Computer Science

Honors and Awards (Selected)

Outstanding Program Committee Member	ICAPS 2022
Distinguished Paper Award	IJCAI 2020
Selected as a Rising Star in EECS	UC Berkeley 2020
CRA-WP Early Career Mentoring Workshop Participant	2020
Conference scholarship	AAMAS 2021
Student travel grant	AIES 2020,AAAI 2019, ICAPS 2017

CS Women travel grant, UMass Amherst	2017, 2019
Robin Popplestone Fellowship, UMass Amherst	2015
Kamaraj Memorial Scholarship, Government of Pondicherry, India	2008

Teaching Experience

Oregon State University	**Spring 2023**
Instructor, AI 533: Intelligent Agents and Decision-Making	
Oregon State University	**Spring 2022**
Instructor, CS 499: Introduction to Intelligent Decision-Making	
Oregon State University	**Winter 2022**
Instructor, AI 539 Special Topics: Safe and Reliable Autonomy	
University of Massachusetts, Amherst	**Fall 2016**
Teaching Assistant, CMPSCI 383: Artificial Intelligence	

Journal Publications (Refereed)

J3. **Sandhya Saisubramanian**, Ece Kamar, and Shlomo Zilberstein. *Avoiding Negative Side Effects of Autonomous Systems in the Open-World*. In the Journal of Artificial Intelligence Research (JAIR), 2022.

J2. **Sandhya Saisubramanian**, Shlomo Zilberstein, and Ece Kamar. *Avoiding Negative Side Effects due to Incomplete Knowledge of AI Systems*. AI Magazine Winter Edition 2022.

J1. Sivakumar Nagalingam, K. Vivekanandan, Arthi Balasubramaiam, **Sandhya Saisubramanian**, and Venus Katta. *Incorporating Agent Technology for Enhancing the Effectiveness of E-learning System*. In the International Journal of Computer Science Issues (IJSCI), 2011.

Conference Publications (Refereed)

C15. Aishwarya Srivastava, **Sandhya Saisubramanian**, Praveen Paruchuri, Akshat Kumar, and Shlomo Zilberstein. *Planning and Learning for Non-Markovian Negative Side Effects Using Finite State Controllers*. In Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI), 2023.

C14. Justin Svegliato, Connor Basich, **Sandhya Saisubramanian**, and Shlomo Zilberstein. *Metareasoning for Safe Decision Making in Autonomous Systems*. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2022.

C13. Sainyam Galhotra, **Sandhya Saisubramanian**, and Shlomo Zilberstein. *Learning to Generate Fair Clusters from Demonstrations*. In Proceedings of the AAAI/ACM Conference on Artificial Intelligence, Ethics and Society (AIES), 2021.

C12. **Sandhya Saisubramanian**, Shannon C. Roberts, and Shlomo Zilberstein. *Understanding User Attitudes Towards Negative Side Effects of AI Systems*. Proceedings of the ACM CHI Conference on Human Factors in Computing Systems, Late-Breaking Work Track (CHI-LBW), 2021.

C11. **Sandhya Saisubramanian** and Shlomo Zilberstein. *Mitigating Negative Side Effects via Environment Shaping*. Extended abstract in Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2021.

C10. **Sandhya Saisubramanian**, Ece Kamar, and Shlomo Zilberstein. *A Multi-Objective Approach to Mitigate Negative Side Effects*. In Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), 2020. **Distinguished Paper Award.**

C9. **Sandhya Saisubramanian**, Ece Kamar, and Shlomo Zilberstein. *Mitigating the Negative Side Effects of Reasoning with Imperfect Models: A Multi-Objective Approach*. Extended abstract in Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2020.

C8. **Sandhya Saisubramanian***, Sainyam Galhotra*, and Shlomo Zilberstein. *Balancing the Tradeoff Between Clustering Value and Interpretability*. In Proceedings of the AAAI/ACM Conference on Artificial Intelligence, Ethics and Society (AIES), 2020.

C7. **Sandhya Saisubramanian**, Connor Basich, Shlomo Zilberstein, and Claudia Goldman. *Satisfying Social Preferences in Ridesharing Services* . In Proceedings of the 22nd IEEE Intelligent Transportation Systems Conference (ITSC), 2019.

C6. **Sandhya Saisubramanian**, Kyle Hollins Wray, Luis Pineda, and Shlomo Zilberstein. *Planning in Stochastic Environments with Goal Uncertainty*. In Proceedings of the 32nd IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2019.

C5. **Sandhya Saisubramanian** and Shlomo Zilberstein. *Adaptive Outcome Selection for Planning With Reduced Models*. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2019.

C4. **Sandhya Saisubramanian**, Shlomo Zilberstein, and Prashant Shenoy. *Planning Using a Portfolio of Reduced Models with Cost Adjustments*. Extended abstract in Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2018.

C3. **Sandhya Saisubramanian**, Pradeep Varakantham, and Hoong Chuin Lau. *Risk Based Optimization for Improving Emergency Medical Systems*. In Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), 2015.
Media Coverage: [Phys.org (Jan 2016)](), [Asian Scientist (Jan 2016)]().

C2. Matthew Brown, **Sandhya Saisubramanian**, Pradeep Varakantham, and Milind Tambe. *STREETS: game-theoretic traffic patrolling with exploration and exploitation*. In Proceedings of Innovative Applications in Artificial Intelligence (IAAI) at 28th AAAI Conference on Artificial Intelligence, 2014.

C1. Sivakumar Nagalingam, K. Vivekanandan, and **Sandhya Saisubramanian**. *Testing Agent-Oriented Software by Measuring Agent's Property Attributes*. In Proceedings of International Conference on Advances in Computing and Communications (ICACC), 2011.

Workshop Papers (Refereed)

W7. Justin Svegliato, Connor Basich, **Sandhya Saisubramanian**, and Shlomo Zilberstein. *Using Metareasoning to Maintain and Restore Safety for Reliably Autonomy*. In workshop on Robust and Reliable Autonomy in the Wild (R2AW), Internationla Joint Conference on Artificial Intelligence (IJCAI) 2021.

W6. Minori Narita, **Sandhya Saisubramanian**, and Shlomo Zilberstein. *Identifying Missing Features in State Representation for Safe Decision-Making*. In workshop on Human-AI Collaboration in Sequential Decision-Making, International Conference on Machine Learning (ICML), 2021.

W5. **Sandhya Saisubramanian**, Minh Do, Jeremy Frank, and Shlomo Zilberstein. *Efficient Integration of Complementary Solvers for Quantum Circuit Compilation*. In Scheduling and Planning Applications workshop (SPARK), International Conference on Automated Planning and Scheduling (ICAPS), 2020.

W4. **Sandhya Saisubramanian**, Connor Basich, Shlomo Zilberstein, and Claudia Goldman. *The Value of Incorporating Social Preferences in Dynamic Ridesharing*. In Scheduling and Planning Applications workshop (SPARK), International Conference on Automated Planning and Scheduling (ICAPS), 2019.

W3. **Sandhya Saisubramanian** and Shlomo Zilberstein. *Minimizing the Negative Side Effects of Planning with Reduced Models*. In AI Safety Workshop, AAAI, 2019.

W2. **Sandhya Saisubramanian** and Shlomo Zilberstein. *Safe Reduced Models for Probabilistic Planning*. In ICML/IJCAI/AAMAS Workshop on Planning and Learning (PAL), 2018.

W1. **Sandhya Saisubramanian**, Shlomo Zilberstein, and Prashant Shenoy. *Optimizing Electric Vehicle Charging Through Determinization*. In Scheduling and Planning Applications workshop (SPARK), 27th International Conference on Automated Planning and Scheduling (ICAPS), 2017.

Dissertation

Sandhya Saisubramanian. *Reliable Decision-Making with Imprecise Models*. Ph.D. Dissertation. Manning College of Information and Computer Sciences, University of Massachusetts Amherst, Oct 2021.

Invited Talks & Guest Lectures

AI in the Real World: Applications and Challenge
Guest Lecture, CS 331: Introduction to Artificial Intelligence, Oregon State University   May 2022

Introduction to Sequential Decision Making
Guest Lecture, AI530: Big Ideas in AI, Oregon State University   May 2022

Avoiding Negative Side Effects of Autonomous Systems in the Open World
Robotics Seminar, Oregon State University   Jan 2022

AI in the Open World: Leveraging Human Guidance to Mitigate
Undesirable Impacts of Incomplete Specification
Responsibility in AI Systems and Experiences Seminar, University of Washington   Jan 2022

Learning to Avoid Negative Side Effects
AI Safety Group, Vector Institute, Canada   Oct 2021

In Pursuit of Reliable Autonomy in the Open World
Tech Talk Tuesday, Oregon State University   Oct 2021

Challenges and Directions in Avoiding Negative Side Effects
Keynote at AI Safety Workshop at AAAI   Feb 2021

Reliable Decision-Making in the Open World
Rising Stars, UC Berkeley   Nov 2020

Invited Panel Discussions

Invited panelist, Reliable Data-Driven Planning and Scheduling Workshop, ICAPS   June 2022

Invited panelist, AAAI undergraduate consortium   Feb 2020

Advising and Mentoring

Research Group...

○ Yashwanthi Anand, Ph.D. in Computer Science, Oregon State University   2021-

○ Pulkit Rustagi, Ph.D. in Robotics, Oregon State University   2022-

Thesis Committee Member..

○ Devin Crowley, Dual M.S. in AI and Robotics, Oregon State University   2022
(Co-advisors: Alan Fern and Jonathan Hurst)

○ Alp A. Aydeniz, Ph.D. in Robotics, Oregon State University   2020-
(Advisor: Kagan Tumer)

○ Skand, Ph.D. in AI, Oregon State University   2021-
(Co-advisors: Fuxin Li and Stefan Lee)

○ Neha Girish Pusalkar, Ph.D. in Robotics, Oregon State University   2021-
(Advisor: Julie A. Adams)

Research Mentoring...

Aishwarya Srivastava, M.S. student (Research Project), IIIT Hyderabad	2021-2022
Moumita Choudhury, Ph.D. student (Research Project), UMass	2021-2022
Saaduddin Mahmud, Ph.D. student (Research Project), UMass	2021-2022
Minori Narita, M.S. student (M.S. Thesis), UMass	2020-2021
Roman Ganchin, M.S. student (Research Project), UMass	2017-2018
Connor Basich, Ph.D. student (Research Project), UMass	2017-2018

Workshop Organization

Co-chair and Lead Organizer, "Robust and Reliable Autonomy in the Wild", IJCAI 2021

Reviewing and Chairing

Journal..

Reviewer, Journal of Artificial Intelligence Research (JAIR)	2021
Reviewer, Frontiers in Artificial Intelligence (section: Machine Learning and Artificial Intelligence)	2021
Reviewer, AI and Society	2021

Conference...

Co-chair, Doctoral Consortium, ICAPS	2024
Session Chair, ICAPS	2022
Program Committee Member, ICAPS	2022
Program Committee Member, IJCAI	2022
Program Committee Member, AAAI	2022
Senior Program Committee Member, IJCAI	2021
Program Committee Member, ICAPS	2021
Program Committee Member, IJCAI	2020
Program Committee Member, AAAI	2021
Program Committee Member, AAAI	2020
Auxiliary Reviewer, ITSC	2019
Auxiliary Reviewer, ICAPS	2018

Workshop...

Program Committee Member, XAIP Workshop at ICAPS	2022
Program Committee Member, AI Safety Workshop at AAAI	2022
Program Committee Member, Multi-Objective Decision Making Workshop (MODeM)	2021
Program Committee Member, AI Safety Workshop at AAAI	2021

Program Committee Member, AI Safety Workshop at AAAI   2020
Program Committee Member, AI Safety Workshop at AAAI   2019
Program Committee Member, AI Safety Workshop at IJCAI   2020
Program Committee Member, AI Safety Workshop at IJCAI   2019

Outreach and University Service

AI Program Committee Member,   2021-2023
Oregon State University

Member, Women in Machine Learning   2020-current

Graduate Student Representative,   Sep 201-Sep 2018
College of Information and Computer Sciences, UMass Amherst

Mentor, Welcoming World to Amherst (WW2A),   Jan-Sep 2016
College of Information and Computer Sciences, UMass Amherst

Volunteer, Eureka! One-Week Workshop (Sponsor: Girls Inc.)   Jul 2018, 2017, 2016
College of Information and Computer Sciences, UMass Amherst

Volunteer, Women in Engineering and Computing Career Day   2016, 2015
College of Information and Computer Sciences, UMass Amherst