

Sajit Rao

☎ (+1) 617-510-2469 | ✉ sajit.rao@gmail.com

Professional Profile

I have spent nearly four decades building cognitive systems across academic research, industrial R&D, entrepreneurship, commercial product development, and university teaching. I design end-to-end intelligent systems that perceive and act in the physical world in real time, combining scientific inquiry with practical engineering, system reliability, and operational responsiveness.

As Founder and Chief Science Officer of Co57 Systems, I won and led a four-year, multimillion-dollar DARPA program in automated human-activity understanding. My work bridges computer vision, machine learning, embodied intelligence, statistical modeling, and the translation of research into systems and products that work reliably in the real world.

Selected Leadership and Technical Highlights

- Founder and Chief Science Officer of Co57 Systems; set the scientific direction and system architecture for a competitive four-year, multimillion-dollar DARPA AI program.
- Built operationally responsive, real-time perception systems, including a hybrid C++ human-detection pipeline that achieved an order-of-magnitude improvement in effective processing speed during DARPA field evaluation.
- Originated two successful DARPA research seedlings at MIT CSAIL, including an \$800,000 effort whose results helped form the basis for DARPA's MindsEye program.
- Originated or co-authored proposals for two independent European Union research programs that subsequently received €5 million in combined funding.
- Designed and developed an end-to-end statistical product for advertising-impact measurement using large-scale commercial, demographic, and market data.
- Six issued U.S. patents from computer-vision research at Xerox PARC; four MIT degrees, including a PhD in Artificial Intelligence.
- Invited to seven European Commission expert review cycles in robotics and artificial intelligence; recognized repeatedly among Koç University's top 10% of instructors.
- Established engineering practices for nightly regression tests and long-duration stress testing of research systems intended for continuous real-world operation.

Education

Massachusetts Institute of Technology (MIT)	Cambridge, MA, USA
PhD, Artificial Intelligence , Department of Computer Science <i>Dissertation: Visual Routines and Attention</i>	Feb 1998
SM, Electrical Engineering and Computer Science <i>Thesis: Repairing Learned Knowledge Using Experience</i>	Feb 1992
SB, Brain and Cognitive Science	Jun 1992
SB, Mathematics with Computer Science (MIT Course 18-C)	Feb 1990

Technical Expertise

Artificial Intelligence and Perception: Computer vision, machine learning, active vision, visual attention, human activity understanding, embodied intelligence, and multisensory perception.

Machine Learning and Data Science: PyTorch, NumPy, and pandas; statistical modeling, anomaly detection, advertising impact analysis, and large-scale commercial data analysis.

Real-Time and Performance-Oriented Research Engineering: Python for research and machine learning; C++ for performance-critical vision and video processing, efficient algorithm implementation, and integration with systems libraries such as FFmpeg; design of hybrid algorithms and processing pipelines that balance scientific accuracy, computational speed, and operational responsiveness.

Research Software Reliability: Established team practices for Jenkins-based nightly regression testing against known results and for long-duration stress and soak testing, often running systems continuously for hours or days to identify crashes, performance degradation, and anomalous behavior.

AI Leadership, Research and Product Development

Koç University *Instructor, Department of Computer Engineering*

Istanbul, Turkey [Sep 2023–Present]

Teach undergraduate computer science and graduate-level Embodied Artificial Intelligence, connecting foundational concepts with real-world systems through examples from research, industry, and entrepreneurship. Developed the graduate course from the ground up, supervise AI capstone projects, and have been recognized repeatedly among the university's top 10% of instructors based on student evaluations.

Leela AI *Consultant*

Boston, MA, USA [Dec 2022–Mar 2023]

Provided technical consulting on computer vision and AI for industrial productivity monitoring in unconstrained outdoor environments.

- Determined that the target event could not be robustly inferred from vision alone and designed a multisensory perception approach combining computer vision and inertial sensing.
- Developed a practical recommendation for integrating AI into real-world construction workflows operating under unconstrained environmental conditions.

Autometrics *Chief Science Officer*

Boston, MA, USA [Jul 2017–Dec 2019]

Personally designed and developed an end-to-end statistical system for measuring the impact of individual television advertisements on consumer vehicle-purchase intent using large-scale demographic and market data.

Co57 Systems *Founder and Chief Science Officer*

Cambridge, MA, USA [Sep 2010–Aug 2016]

Founded Co57 Systems as an AI research and engineering company focused on active perception, computer vision, and real-time intelligent systems.

- As Chief Science Officer and Principal Investigator, set the scientific direction and system architecture for a large-scale DARPA program in automated human activity understanding, leading the development of biologically inspired real-time perception systems.
- Personally designed and implemented a hybrid human-detection pipeline in C++, using motion-based detection for continuous video and invoking slower static-frame detectors only at approximately 60-frame intervals. The approach provided an order-of-magnitude improvement in effective processing speed and enabled operationally responsive performance during DARPA's final field evaluation in an unconstrained forest environment.

Massachusetts General Hospital *Consultant*

Boston, MA, USA [Jul 2005–Oct 2005]

Designed and developed an AI-based document-understanding prototype to parse and categorize patient reports, bringing difficult-to-interpret paper content into the Electronic Health Record system.

Designed a web-based educational course-creation product and led its development and delivery in less than a year. The product accelerated course production by a factor of four.

Conducted business-process studies with major airline and airport clients. Separately developed a consumer-facing travel-planning prototype based on a concept originated by Ascent's founder.

Co-founded an early image-based search venture with two fellow MIT AI Laboratory graduate students, at a time when mainstream search was almost entirely text-based. The founding concept won the Grand Prize in the year-long 1997 MIT \$50K Entrepreneurship Competition. Explored commercial applications including visual search for trademarks and logos.

SELECTED ACADEMIC AI RESEARCH

Conceived and led a DARPA-funded research seedling investigating how common-sense reasoning can emerge from perceptual mechanisms rather than explicit symbolic representations.

The success of this effort led to a white paper that became the basis for DARPA's MindsEye program on automated visual understanding of human activities.

Conducted research on embodied approaches to general intelligence within a seven-laboratory EU Cognitive Vision consortium spanning robotics, computer vision, neuroscience, and developmental psychology.

Conceived a three-stage developmental architecture for robotic cognition and implemented core components of its first stage. The work was singled out by EU project reviewers as outstanding and novel.

Conducted doctoral research under Professors Rodney Brooks and Patrick Winston on active vision and visual routines, examining how intelligent visual behavior emerges through task-directed interaction with the environment, rather than passive image interpretation.

The dissertation introduced computational mechanisms for directing visual attention through the active control of perception. This work laid much of the conceptual foundation for my later work on embodied intelligence, developmental learning, and active perception.

Research Contributions

- **Embodied Intelligence**

Developed computational approaches to intelligence grounded in perception, action, and interaction with the physical world, rather than symbolic reasoning alone.

- **Developmental Cognition**

Conceived a three-stage developmental architecture for robotic cognition and implemented foundational components, grounding incremental cognitive development in perception, action, and interaction with the physical world.

- **Active Vision and Visual Attention**

Developed computational approaches to task-directed visual perception, visual routines, and active control of attention, beginning with doctoral research at MIT and continuing through later work on real-time perception systems.

- **Human Activity Understanding**

Led research on automated interpretation of human activities from video, including biologically inspired models of perception, representation, and common-sense reasoning.

- **Real-Time Perception Systems**

Designed complete perception-to-action systems capable of operating in real-world environments, integrating computer vision, machine learning, randomized attention-selection mechanisms, and robust software engineering.

- **Technology Translation**

Translated research in computer vision, machine learning, and statistical modeling into deployable systems for defense and commercial applications.

Selected Funded Research Programs

DARPA MindsEye Program — Principal Investigator, Co57 Systems 2010–2014 | Multimillion-dollar award

Won and led a competitive DARPA research contract for automated visual understanding of human activities. Directed the scientific program, system architecture, and development of biologically inspired real-time perception and activity-understanding systems.

DARPA Perceptual Priming for Language Learning Seedling — Proposal Originator, MIT CSAIL Proposal accepted Feb 2010

Conceived and authored a successful DARPA seedling proposal investigating how perceptual priming could support language learning. Following my transition to Co57 Systems, Patrick Winston and his students led the funded research.

DARPA Visuospatial Reasoning Seedling — Principal Investigator, MIT CSAIL 2006–2009 | \$800,000

Conceived and secured funding for research investigating how common-sense reasoning could emerge from perceptual mechanisms. The successful seedling produced a white paper that became the basis for DARPA's MindsEye program.

CONTACT EU Research Proposal — Originating Proposal Author Original proposal 2004 | Later award: €3M

Originated and wrote the initial proposal. Following revision and resubmission by the Italian research group after my departure, the independent project received €3 million in funding.

COGVIS EU Collaborative Research Project — Contributing Proposal Author 2001–2004 | €2M

Contributed to the jointly authored proposal that secured €2 million for a seven-laboratory European research consortium.

Selected Publications

1. P. Fitzpatrick, G. Metta, L. Natale, **S. Rao**, and G. Sandini, “*Learning about Objects through Action: Initial Steps toward Artificial Cognition*,” Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2003.
2. L. Natale, **S. Rao**, and G. Sandini, “*Learning to Act on Objects*,” Proceedings of the International Workshop on Biologically Motivated Computer Vision (BMCV), Springer-Verlag, 2002.
3. G. Metta, L. Natale, **S. Rao**, and G. Sandini, “*Development of the Mirror System: A Computational Model*,” Conference on Brain Development and Cognition in Human Infants: Emergence of Social Communication—Hands, Eyes, Ears, Mouths, Acquafredda di Maratea, Italy, Jun 2002.
4. D. R. Magee, C. J. Needham, P. E. Santos, and **S. Rao**, “*Inducing the Focus of Attention by Observing Patterns in Space*,” IJCAI Workshop on Modeling Others from Observations (MOO), pp. 47–52, 2005.
5. P. H. Winston and **S. Rao**, “*Repairing Learned Knowledge Using Experience*,” in *Artificial Intelligence at MIT: Expanding Frontiers*, vol. 1, ch. 14, MIT Press, 1990.
6. J. Heel and **S. Rao**, “*Temporal Integration of Visual Surface Reconstruction*,” Image Understanding Workshop, 1990.

U.S. Patents

Co-inventor on six issued U.S. patents assigned to Xerox Corporation: 5,455,898; 5,513,271; 5,522,022; 5,537,491; 5,544,267; and 5,659,639. The patents cover computer-vision methods for analyzing graphical representations, node-link structures, proportioned-parts graphs, stable grouping, and editing operations.

Professional Service and Recognition

European Commission Expert Panel in Robotics and Artificial Intelligence 2007–2014 | Seven review cycles

Invited to evaluate EU research proposals, interview research teams, assess scientific and execution plans, and contribute to funding recommendations across seven review cycles.

- Recognized repeatedly among Koç University's top 10% of instructors based on student evaluations.
- Selected as a Kauffman Fellow following a competitive application and personal interview, 2008.
- Grand Prize, MIT \$50K Entrepreneurship Competition, 1997.
- Elected Full Member, Sigma Xi, The Scientific Research Honor Society.
- MIT Peter J. Eloranta Undergraduate Research Award, 1988.

Selected Invited Talks

- *The Influence of Brain and Cognitive Science on My Work in AI*, Koç University Neuroscience Festival, 2025.
- *Demystifying AI: What Innovators Need to Know*, Tufts University Gordon Institute, 2025.
- *Constructing Intelligence*, Norway Financial Summit, 2018.
- *A Model of Visual Attention for Event Perception and Recognition*, University of Coimbra, 2015.
- *Language Grounding*, DARPA Language Workshop, 2008.
- *Visuospatial Reasoning*, DARPA Spatial Reasoning Workshop, 2006.