Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

RECOMMENDED READINGS

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL) & CSAIL ALLIANCES GLOSSARY OF TERMS

CONTRBUTED BY:

Karen Savoie, Founder & President, Savoie Faire Consulting

The IRRG is aware that subject matter experts, practitioners, students, and others are busy with current and developing projects and has provided the following list of readings to share knowledge of potential interest in support of efforts relating to the resilience, security, and risk management of critical infrastructure assets and networks.

AI (Artificial Intelligence)

A broad umbrella term for systems that perform tasks associated with human intelligence, such as perception, reasoning, language, prediction, and decision-making.

AI Agent / Agentic AI

An AI system that does not just generate output but can take actions toward a goal, often by planning, using tools, and reacting to changing conditions.
CSAIL PIs to follow: Associate Professor Tim Kraska, Professor Sam Madden, Principal Research Scientist Michael Cafarella, Professor Antonio Torralba.
Dive Deeper *Agentic AI: What you need to know about AI Agents*

AI Lab (Artificial Intelligence Laboratory)

Originally a specialized group within Project MAC, the AI project was founded in 1959 under the leadership of computing pioneers like Marvin Minsky. It served as a premier global center for foundational research into robotics, computer vision, and thinking machines, before eventually merging with LCS to form CSAIL in 2003.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

AI Model

A trained computational system that maps inputs to outputs. This is the broad category that includes classifiers, vision models, language models, recommender systems, and more.
Dive Deeper *Ask CSAIL: Who's Using What Model*

Alignment

Making AI systems behave in ways that match intended goals, human values, policy constraints, or task requirements.

Algorithmic Design

A method where rules, constraints, and mathematical logic are used to generate complex forms or solutions that would be difficult to create manually. Rather than drawing a final result, the designer creates a set of instructions that the computer follows to explore a vast range of optimized possibilities.
CSAIL PIs to follow: Assistant Professor Mina Konaković Luković as leader of the Algorithmic Design Group.
Dive Deeper *Algorithmic Design with Assistant Professor Mina Konaković Luković*

Architecture (referring to Computer Architecture)

The fundamental design and organization of a computer system, including the CPU, memory, and how different hardware components interact to process instructions efficiently.
Dive Deeper *Specialized Hardware Design with MIT CSAIL Assistant Professor Rachit Nigam*

Artificial General Intelligence (AGI)

A theoretical milestone where an AI system can understand, learn, and apply knowledge across a wide range of tasks at or above human level.
Dive Deeper *The Revolutionary Potential of AI with CSAIL Professor Manolis Kellis*

Autonomous systems / Autonomy

Systems that can sense, decide, and act in the world with reduced human intervention. Examples include self-driving cars.
CSAIL PIs to follow: Many professors at MIT CSAIL explore this in the context of vehicles, robotics, or other systems, including Professor Nicholas Roy, Professor Russ Tedrake, Professor Daniela Rus, Professor John Leonard, and Associate Professor Pulkit Agrawal.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Bayesian

A method of problem-solving where an initial guess is constantly updated as new information comes in. Bayesian methods help AI systems make calculated decisions even when they are dealing with uncertain or incomplete data.

Benchmarking

Evaluating a model or system against standard datasets, tasks, or performance criteria.

Black Box

An AI system or algorithm whose internal decision-making processes are hidden or too complex for humans to understand. The inputs and outputs are visible, but exactly how the system arrived at its conclusion remains opaque.

Blue Teaming (Cybersecurity)

The defensive side of cybersecurity testing. Blue teams evaluate an organization's network security, identify vulnerabilities, and actively defend against simulated cyber attacks (carried out by Red Teams).
CSAIL PIs to follow: Senior Research Scientist Una-May O'Reilly and the Anyscale Learning For All (ALFA) Group

Chatbot

A conversational interface designed to interact with users in dialogue. A chatbot may use an LLM, but a chatbot is the application layer, not the underlying model itself.

Cloud Computing

The delivery of computing services—such as data storage, processing power, servers, and databases—over the internet ("the cloud") rather than relying entirely on local servers or personal devices.
Dive Deeper *The Landscape of Cloud Computing with Professor Christina Delimitrou*

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Compiler

A specialized software program that translates human-readable source code written in a high-level programming language into low-level machine code that a computer's hardware can execute.
Dive Deeper *Specialized Programming Languages and Compilers with Professor Saman Amarasinghe*

Computational Biology

The application of computer science, algorithms, and mathematical modeling to understand and solve complex biological problems, such as genome sequencing, disease mechanisms, and protein folding.
CSAIL PIs to follow: Professor Manolis Kellis, head of the Computational Biology Group.
Dive Deeper *Programmable Therapeutics: A modular approach to precision medicine with MIT CSAIL Professor Manolis Kellis*

Computer Assisted Design (CAD)

The use of specialized software to aid in the creation, modification, analysis, and optimization of digital models. It is heavily used in engineering, architecture, and digital fabrication.
CSAIL PIs to follow: Professor Wojciech Matusik as leader of the Computational Design and Fabrication Group.

Computer Vision

AI methods for interpreting images and video—recognizing objects, scenes, motion, depth, and visual relationships.
CSAIL PIs to follow: The Visual Computing Group at CSAIL is led by Associate Professor Jonathan Ragan-Kelley and includes over a dozen CSAIL professors including Professor Fredo Durand, Professor Antonio Torralba, Professor Polina Golland, and Professor William Freeman.

Dive Deeper *How AI is Reshaping Medical Imagery with MIT CSAIL Professor Polina Golland*

Cryptography

The science of secure communication. It involves encrypting and decrypting data to ensure that information remains private, secure, and unaltered by unauthorized parties.
CSAIL PIs to follow: Professor Yael Kalai, Professor Vinod Vaikuntanathan.
Dive Deeper The Expanding Frontier of Modern Cryptography with MIT CSAIL Professor Yael Kalai , The Future of Cryptography with Professor Vinod Vaikuntanathan

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

CSAIL

The Computer Science and Artificial Intelligence Laboratory at MIT. Formed in 2003 by the merger of the AI Lab and the Laboratory for Computer Science (LCS), it is the largest interdepartmental research lab at MIT.

Deep Learning

A set of machine-learning approaches that use neural networks with many layers to learn hierarchical representations directly from data. These methods are especially effective for high-dimensional, unstructured data such as language, images/video, audio, and sensor streams, and they power modern systems for recognition, generation, and decision-making/control.

Deterministic

A deterministic system or algorithm is one that always produces the same output from a given input, following a predictable, fixed sequence of states (as opposed to stochastic).

Diffusion Model

A generative model that learns to create images, molecules, audio, or other data by starting from random noise and gradually "denoising" it into a coherent sample. During training, the model learns how data gets progressively corrupted by noise; during generation, it reverses that process step by step, using its learned denoising predictions to produce realistic new outputs.

Digital Fabrication

A manufacturing process where computer-controlled machines (such as 3D printers, laser cutters, and CNC routers) are used to create physical objects from digital designs.
CSAIL PIs to follow: Assistant Professor Mariana Popescu.
Dive Deeper *Knitting the Future of Construction with MIT CSAIL Assistant Professor Mariana Popescu*

Distributed Systems

A foundational computing area for CSAIL involving multiple computers that communicate and coordinate to achieve a common goal. Distributed systems underlie cloud services, large databases, content delivery networks, and many modern AI/data platforms.
CSAIL PIs to follow: Professor Barbara Liskov won the 2008 Turing Award in part for her groundbreaking work on distributed systems.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Dive Deeper *Women's History Month Profiles: The Pioneering Women of CSAIL*

EGI (Engineering General Intelligence)

Applying AI to support and automate engineering work end-to-end—capturing engineering knowledge and using it to help generate, evaluate, and refine designs.
CSAIL PIs to follow: Coined by researchers in Professor Wojciech Matusik's lab and a central idea behind the CSAIL Spinout Foundation EGI.

Dive Deeper *CSAIL Spinout Foundation EGI Pioneers New Area of AI: Engineering General Intelligence*

Embodied AI / Embodied Intelligence

AI for agents that have to operate in a physical (or realistically simulated) environment, where intelligence is expressed through sensing, movement, and interaction, not just producing text or predictions. Typical examples include mobile robots, drones, robot arms, and other autonomous machines that manipulate objects, navigate spaces, and collaborate with people.
CSAIL PIs to follow: The Embodied Intelligence Community of Research, led by Professor Nicholas Roy, includes more than 20 CSAIL PIs.

Explainability / Interpretability

The set of techniques and strategies for understanding how and why an AI model arrives at its decisions or what patterns it has learned from data. Interpretability focuses on making a model's inner workings transparent; explainability goes further by providing human-understandable reasons or justifications for specific predictions. Both are critical for trust, debugging errors, detecting bias, meeting regulatory requirements, and effective human collaboration.

Fine-Tuning

The process of taking a pre-trained foundation model and continuing its training on a smaller, targeted dataset to adapt it for a specific task.

Foundation Model

A large, general-purpose model trained on broad data that can be adapted to many downstream tasks.

Generative AI

AI systems trained to generate new, original content—such as text, images, code, audio, video, molecules, or designs—by learning patterns from large amounts of training data.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Hallucinate (in the context of AI)

When a generative AI model confidently produces false, fabricated, or nonsensical information that is not grounded in its training data or the user's input prompt.

Human-Centered AI

An approach to designing AI that prioritizes human needs, values, and contexts, ensuring systems are usable, fair, transparent, and aligned with people's goals. It emphasizes collaboration (AI as a partner or assistant, not a replacement), clear communication of capabilities and limits, accountability, and design choices that support trust, safety, and meaningful human control.

Human-Computer Interaction (HCI)

The interdisciplinary study and practice of designing, evaluating, and improving the ways people interact with computers and digital systems. HCI combines psychology, design, engineering, and computer science to create interfaces and interactions that are usable, accessible, efficient, and satisfying—ensuring technology supports human goals, minimizes errors, and fits real-world workflows and contexts.
CSAIL PIs to follow: Associate Professor Stefanie Mueller as leader of the HCI Engineering Group at CSAIL.
Dive Deeper *Understanding Generative AI with Stefanie Mueller*

IAP (MIT specific)

Independent Activities Period, a special four-week term at MIT every January where students and faculty participate in short, non-traditional classes, hackathons, and experimental projects.
Dive Deeper *Sony Interactive Entertainment Hosts MIT IAP Course on Video Games & AI*

Inference

The "live" stage where a trained AI model is put to work to make predictions or generate content using new, real-world data. It covers the entire journey from receiving a user's request to delivering a final answer, balancing how fast the system responds (latency) with how many users it can handle at once (throughput).

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Large Language Model (LLMs)

A model trained to predict or generate language, trained on vast text corpora to generate, summarize, reason over, or transform language. An LLM is the model; a chatbot is one possible interface built on top of it.

Latency

The time delay between a user or system initiating a request and receiving a response. In AI and robotics, low latency is critical for real-time systems to function safely and effectively.

LCS (Laboratory for Computer Science)

Originally founded in 1963 as Project MAC, LCS was created to develop a computer system accessible to a large number of people and sponsored by the Department of Defense. In 2003, LCS and the AI Lab merged to create CSAIL.

Liquid Neural Network

A type of AI model inspired by the brains of tiny organisms (like roundworms) that can adapt behavior "on the fly" as new data comes in. Unlike traditional AI, which is fixed after training, a Liquid Neural Network uses flexible, continuous-time equations that allow it to remain stable and accurate even when environmental conditions change.
CSAIL PIs to follow: Originally developed in Professor Daniela Rus's lab; serves as the foundational technology of the CSAIL Spinout Liquid AI.

Machine Learning (ML)

A subset of AI that focuses on building systems that learn and improve from experience without being explicitly programmed for every task. While AI is the broad vision of creating "smart" machines, Machine Learning is the specific engine that uses data and statistics to find patterns and make decisions.

Moore's Law

An observation and historical trend stating that the number of transistors on a microchip roughly doubles every two years, which has historically driven exponential increases in computing power and decreases in cost. Now that physical limitations have effectively ended this trend, the industry has shifted toward specialized AI hardware and architectural innovations to continue driving performance gains.
Dive Deeper *The Death of Moore's Law: What it means and what might fill the gap going forward*, *Software Performance Engineering in a Post-Moore Era with Professor Charles Leiserson*, *Adapting to the Slowing Rate of Technological Improvement with CSAIL Research Scientist Neil Thompson*

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Multimodal AI

AI that works across multiple kinds of input or output, such as text, images, video, audio, sensor data, or robotics state.

Natural language processing (NLP)

The field focused on making computers understand, generate, and reason with human language.
CSAIL PIs to follow: Associate Professor Jacob Andreas, Professor Tommi Jaakkola, and Associate Professor Yoon Kim as leaders of the Natural Language Processing Group.
Dive Deeper *The Changing Landscape of Natural Language Processing with Assistant Professor Yoon Kim*, *AI's Language Leap: MIT CSAIL Associate Professor Jacob Andreas Explores NLP and LLMs*

Neural Network

A computational model made of interconnected units that learns patterns from data; the main building block of most modern deep learning systems.

Open Source

Software, hardware, or research artifacts released with licenses that permit inspection, reuse, modification, and redistribution. CSAIL Alliances members have access to a searchable database of CSAIL open source tools here.
Dive Deeper *Open Source Projects Search Tool* (available to CSAIL Alliances Members)

Perception

How a machine extracts useful information from raw sensory input like images, lidar, audio, or force data.

Physical AI

Systems that combine perception, reasoning, and action in embodied systems like robots or manufacturing platforms which must interact with the tangible world.
CSAIL PIs to follow: Originally coined by Professor Daniela Rus, this term has been used in recent CSAIL materials to frame physically intelligent robots.

Planning (in Robotics)

Choosing a sequence of actions to achieve a goal, often under constraints and uncertainty.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Platform (for AI)

A comprehensive software architecture and infrastructure that provides the necessary tools, computing environments, and services to build, deploy, train, and manage AI applications.

Program Synthesis

The automated generation of software from a high-level description of its desired behavior. Instead of a human writing the code line-by-line, the system writes the code to meet the user's constraints.
CSAIL PIs to follow: Professor Armando Solar-Lezama as leader of the Computer-Aided Programming Group.
Dive Deeper *The Reality & Potential of Program Synthesis with Professor Armando Solar-Lezama*

Project MAC (Multiple Access Computing or Machine-Aided Cognition)

Founded at MIT in 1963 as a project sponsored by the Department of Defense to develop a computer system accessible to a large number of people, this was the early research hub that laid the foundation for modern computing and serves as the direct ancestor of CSAIL. A precursor of LCS.

Prompt Engineering

The skill or process of structuring text input to a generative AI model to yield the most accurate or useful output.

Quantum Computing

An emerging paradigm of computing that harnesses the laws of quantum mechanics (like superposition and entanglement) to solve certain highly complex problems exponentially faster than classical computers ever could.
Dive Deeper *The State of Quantum Computing with Professor Peter Shor & Professor William Oliver*

Red Teaming (in cybersecurity)

The practice of rigorously playing the role of an adversary to test the security or safety of a system. In AI, this involves intentionally trying to break an AI model's safety guardrails to expose vulnerabilities before deployment.
CSAIL PIs to follow: Senior Research Scientist Una-May O'Reilly and the Anyscale Learning For All (ALFA) Group.
Dive Deeper *Evolutionary Algorithms, Adversarial Intelligence, and Cybersecurity with MIT CSAIL Senior Research Scientist Una-May O'Reilly*

Reinforcement Learning (RL)

A learning framework in which an agent learns by trial and error to maximize reward through interaction with an environment.

Retrieval-Augmented Generation (RAG)

A setup in which a language model is connected to external documents or databases so it can retrieve facts before generating an answer.

Robotics

The design of systems that sense, compute, and act in the physical world.
CSAIL PIs to follow: Many professors at CSAIL work on robotics, including Professor Daniela Rus, Professor Russ Tedrake, Associate Professor Pulkit Agrawal, Professor Nicholas Roy, Professor Leslie Kaelbling, Professor Tomás Lozano-Pérez, Professor John J. Leonard, Assistant Professor Andreea Bobu.
Dive Deeper *Human Centric Robot Learning with MIT CSAIL Assistant Professor Andreea Bobu*, *The Potential of Robotic Manipulation with Professor Russ Tedrake*, *Force-Centric Dexterous Robotic Manipulation with MIT CSAIL Associate Professor Pulkit Agrawal*, *The Long Horizon of Robotics with MIT CSAIL Professor Leslie Kaelbling*, *The Future of Robotics Hype, Breakthroughs, and Challenges with Professor Russ Tedrake*

Schwarzman College of Computing (Building 45)

A cross-cutting academic hub announced in 2018 to advance computer science, AI, and data science education and research while integrating these technologies across all disciplines at the Institute. Headquartered in Building 45, the college represents the first fundamental change of MIT's academic structure in nearly 70 years.

Self-Supervised Learning

A technique where a model teaches itself by hiding parts of its own unlabeled data and trying to predict them, such as guessing the next word in a sentence. This allows AI to recognize patterns from massive datasets without the need for manual human labeling.

Sim-to-Real

Training or validating in simulation, then transferring the learned behavior to real-world robots or systems.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Stata Center (Building 32)

The iconic Frank Gehry-designed building that serves as the physical home of CSAIL.

Stochastic

Processes, models, or systems governed by random probability rather than fixed rules. Unlike deterministic systems, which always produce the same output for a given input, stochastic models involve an element of chance, meaning their results can vary each time they run.

Supervised Learning

A method of training an AI model using data that has already been labeled by humans.

Symbolic Programming

A computer programming paradigm where programs manipulate human-readable symbols, logic, and abstract expressions as data rather than relying solely on numerical calculation. The primary advantage of symbolic programming is its ability to reason about abstract concepts and logic, allowing a computer to solve complex, unstructured problems like mathematical proofs, natural language processing, and advanced AI reasoning.

Throughput

The volume of requests a system can handle within a specific timeframe (e.g., 1,000 queries per second).

Tool Use

When an AI system calls external tools—search, code execution, APIs, databases, simulation engines—to complete a task. This is a core ingredient in many agentic systems.

Traceability (of Data)

The ability to track the origin, lineage, and lifecycle of data used to train AI models. This is crucial for verifying data integrity, managing copyright attribution, and ensuring legal compliance.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Transformers

A breakthrough deep learning architecture introduced in 2017 that relies on "attention" mechanisms to weigh the importance of different parts of a sequence. This allows the model to understand long-range context and relationships in data, making it the foundational technology behind most modern Large Language Models (LLMs) and generative AI.

Trustworthy AI

A broad term covering reliability, safety, privacy, robustness, explainability, and governance. Dive Deeper *Risks and considerations of AI in healthcare with Associate Professor Marzyeh Ghassemi*

Turing Award

Often referred to as the "Nobel Prize of Computing," 10 CSAIL faculty members and alumni have won this top honor.

UI (User Interface)

The visual and interactive points of contact—such as screens, menus, buttons, and voice commands—through which a human communicates with a computer system or software application.

Uncertainty (in Computer Science)

A state in which a system must operate with incomplete, noisy, or unpredictable information. Managing uncertainty is a major focus in robotics, probabilistic reasoning, and AI decision-making.

Unsupervised Learning

A method of training an AI model using raw, unlabeled data, leaving the system to find hidden patterns and structures entirely on its own.

Vector Database

A specialized database designed to store and search high-dimensional mathematical representations of data rather than traditional tables of text. They are the foundational architecture for Retrieval-Augmented Generation (RAG), allowing AI to instantly find the most relevant context from millions of documents to ground its answers in facts.

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Vertical AI

AI systems designed and trained for a specific industry or domain, such as healthcare, law, or manufacturing, rather than as a general-purpose assistant. By focusing on deep, specialized data, these models can solve complex, niche problems with a level of precision and expert knowledge that general AI often lacks.

Vibe Coding

A popular informal term for coding by prompting AI systems in a lightweight, exploratory, or improvisational way rather than building or writing everything manually. It describes a shift in productivity where a software engineer (or even programming novice) focuses on high-level intent and iterative "vibes" rather than writing code by hand.

Vision-Language Model (VLM)

AI systems designed to jointly reason over both text and visual inputs, such as images or video. This integrated understanding allows them to perform complex tasks like describing a scene, answering questions about a photo, or following text-based instructions to navigate a physical space.

Visual Computing

A broad umbrella spanning computer vision, graphics, imaging, and sometimes visualization—basically computation involving visual data and visual representation.

W3C (World Wide Web Consortium)

The main international standards organization for the internet, originally founded and hosted at MIT/CSAIL by Tim Berners-Lee.

World Model

An emerging class of AI designed to move beyond pattern matching toward a deep, internal understanding of how the physical and conceptual world works. By mastering the underlying rules of reality—such as physics, causality, and the consequences of actions—these models are intended to give AI and robotics the human-like common sense and planning abilities that current models lack.
Dive Deeper *The Revolutionary Implications of World Models with MIT CSAIL Assistant Professor Vincent Sitzmann*

Carleton University
Infrastructure Resilience Research Group (IRRG)
Office of the Dean, Faculty of Engineering and Design

Zero-Knowledge Proof (ZKP)

A cryptographic method allowing one party to prove to another that a statement is true without revealing any information beyond the validity of the statement itself. (Professors Shafi Goldwasser and Silvio Micali won the 2012 Turing Award for pioneering this at MIT).

Zero-Shot Learning

The ability of an AI model to successfully perform a task or recognize an object it has never explicitly been trained on.