FAST-10 Reading List: AI Governance, Risk, Ethics, and Responsible Use for Growing Businesses

Provider: 7formations | Course: FAST-10 | Instructor: Wadood Chaudhary | Links verified by fetch or search as of Oct 2026

Training, not legal advice. These are primary sources: frameworks (NIST, ISO), laws and regulator guidance (Massachusetts, federal, EU, one other state as an example), and security references (OWASP, CISA). Read the required items; use the optional ones as references when you build your inventory, register, test script, and policy. Rules change: check each page's date, and confirm obligations for your employer with counsel.

How to use. Required readings take 20 to 45 minutes per session and are done before (or right after) the session they support. Optional readings go deeper. 'Living page' means the publisher updates it continuously; read the current version. Reading times are estimates for the parts that matter to this course, not the whole document.

Session 1: What governance means in a small company

Required

- **AI Risk Management Framework (overview page)**
  Publisher: National Institute of Standards and Technology (NIST) | Date: Living page; AI RMF 1.0 released January 2023 | Reading time: about 10 min
  URL: https://www.nist.gov/itl/ai-risk-management-framework
  Why read: NIST's AI RMF landing page: the current version, playbook, profiles, and any revision news. Check it before relying on any summary, including ours.
- **Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1**
  Publisher: NIST | Date: January 2023 | Reading time: about 30 min
  URL: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
  Why read: AI RMF 1.0. Read Part 2 (the Core: Govern, Map, Measure, Manage). Skim the trustworthiness characteristics in Part 1.

Optional

- **NIST AI RMF Playbook**
  Publisher: NIST Trustworthy and Responsible AI Resource Center (AIRC) | Date: Living page | Reading time: about 20 min
  URL: https://airc.nist.gov/airmf-resources/playbook/
  Why read: The AI RMF Playbook: suggested actions for each subcategory. Use it to pick a few right-sized actions, not all of them.
- **ISO/IEC 42001:2023 Information technology: Artificial intelligence: Management system**
  Publisher: International Organization for Standardization (ISO) | Date: December 2023 | Reading time: about 10 min
  URL: https://www.iso.org/standard/42001

Why read: ISO/IEC 42001:2023, the certifiable AI management-system standard. Read the abstract to see how it differs from the NIST framework; customers or partners may ask about it.

- **Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1**
  Publisher: NIST | Date: July 2024 | Reading time: about 25 min
  URL: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
  Why read: NIST AI 600-1, the Generative AI Profile. Read the list of 12 generative-AI risks; it is a ready-made checklist for the inventory and register.

Session 2: Laws, contracts, and data

Required

- **General Laws Chapter 93H: Security Breaches**
  Publisher: The 194th General Court of the Commonwealth of Massachusetts | Date: Current statute text | Reading time: about 10 min
  URL: https://malegislature.gov/Laws/GeneralLaws/PartI/TitleXV/Chapter93H
  Why read: G.L. c. 93H: definition of personal information and breach notice duties.
- **201 CMR 17.00: Standards for the Protection of Personal Information of MA Residents**
  Publisher: Mass.gov, Office of Consumer Affairs and Business Regulation | Date: Regulation page dated 10/19/2017 (in force since 2010) | Reading time: about 15 min
  URL: https://www.mass.gov/regulations/201-CMR-1700-standards-for-the-protection-of-personal-information-of-ma-residents
  Why read: 201 CMR 17.00: the written information security program, including vendor oversight (17.03(2)(f)).
- **Attorney General Advisory on the Application of the Commonwealth's Consumer Protection, Civil Rights, and Data Privacy Laws to Artificial Intelligence**
  Publisher: Massachusetts Attorney General's Office | Date: April 16, 2024 | Reading time: about 15 min
  URL: https://www.mass.gov/doc/ago-ai-advisory-41624/download
  Why read: Massachusetts Attorney General advisory (April 2024): consumer-protection, anti-discrimination, and data-security laws apply to AI developers, suppliers, and users.
- **AI Act (regulatory framework overview and application timeline)**
  Publisher: European Commission, Shaping Europe's digital future | Date: Living page (updated after the AI Omnibus, 2026) | Reading time: about 10 min
  URL: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
  Why read: The European Commission's AI Act page: risk tiers and the application timeline. Relevant if you sell AI-enabled products or services into the EU.

Optional

- **201 CMR 17.00 Compliance Checklist**
  Publisher: Mass.gov, Office of Consumer Affairs and Business Regulation | Date: Living page | Reading time: about 10 min
  URL: https://www.mass.gov/info-details/201-cmr-1700-compliance-checklist

Why read: The Commonwealth's 201 CMR 17.00 compliance checklist. Good for mapping AI use into an existing WISP.

- **Summary of the HIPAA Privacy Rule**
  Publisher: U.S. Department of Health and Human Services, Office for Civil Rights | Date: Content last reviewed March 14, 2025 | Reading time: about 20 min
  URL: https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html
  Why read: HHS summary of the HIPAA Privacy Rule: business associates and minimum necessary.

- **General Laws Chapter 151B, Section 4: Unlawful practices**
  Publisher: Massachusetts Legislature | Date: Current text | Reading time: about 10 min
  URL: https://malegislature.gov/Laws/GeneralLaws/PartI/TitleXXI/Chapter151B/Section4
  Why read: G.L. c. 151B, § 4: unlawful employment practices. The anti-discrimination baseline for any AI used in hiring, promotion, or discipline.

- **Consumer Financial Protection Circular 2022-03: Adverse action notification requirements in connection with credit decisions based on complex algorithms**
  Publisher: Consumer Financial Protection Bureau | Date: May 26, 2022 | Reading time: about 8 min
  URL: https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/
  Why read: CFPB Circular 2022-03: adverse-action reasons are required even with complex models.

- **Massachusetts law about artificial intelligence**
  Publisher: Mass.gov, Trial Court Law Libraries | Date: Last updated August 28, 2026 | Reading time: about 10 min
  URL: https://www.mass.gov/info-details/massachusetts-law-about-artificial-intelligence
  Why read: Trial Court Law Libraries guide to Massachusetts law about AI (updated 2026). Check it for new bills or guidance.

- **SB26-189 Automated Decision-Making Technology**
  Publisher: Colorado General Assembly | Date: Signed May 14, 2026; applies to decisions on or after January 1, 2027 | Reading time: about 10 min
  URL: https://www.leg.colorado.gov/bills/sb26-189
  Why read: Colorado SB26-189 (May 2026), which replaced the 2024 Colorado AI Act. An example of a state law that can reach out-of-state businesses making consequential decisions about residents.

- **AI Omnibus enters into force**
  Publisher: European Commission, Shaping Europe's digital future | Date: July 27, 2026 | Reading time: about 5 min
  URL: https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force
  Why read: The AI Omnibus (in force July 2026) that moved high-risk deadlines to December 2027 (Annex III) and August 2028 (Annex I).

- **When does enforcement start? (AI Act Service Desk FAQ)**
  Publisher: European Commission, AI Act Service Desk | Date: Living page | Reading time: about 5 min
  URL: https://ai-act-service-desk.ec.europa.eu/en/ai-act/faq/when-does-enforcement-start
  Why read: AI Act Service Desk FAQ: when enforcement starts for each part of the Act.

Session 3: Measuring and testing risk

Required

- **OWASP GenAI LLM Top 10 2026**
  Publisher: OWASP Gen AI Security Project | Date: Published August 4, 2026 | Reading time: about 20 min
  URL: https://genai.owasp.org/resource/owasp-genai-llm-top-10-2026/
  Why read: OWASP Top 10 for LLM Applications 2026. Read the one-line descriptions of all ten; use LLM01, LLM02, LLM03, LLM07, and LLM08 in your test script.

Optional

- **OWASP Top 10 for LLM Applications**
  Publisher: OWASP Gen AI Security Project | Date: 2025 edition | Reading time: about 5 min
  URL: https://genai.owasp.org/llm-top-10/
  Why read: OWASP GenAI project home for the Top 10, with earlier versions and related guidance.
- **AI Data Security: Best Practices for Securing Data Used to Train and Operate AI Systems**
  Publisher: Cybersecurity and Infrastructure Security Agency (CISA) with partners | Date: May 2025 | Reading time: about 20 min
  URL: https://www.cisa.gov/resources-tools/resources/ai-data-security-best-practices-securing-data-used-train-operate-ai-systems
  Why read: CISA and partners on securing data used to train and operate AI systems. Useful for vendor questions about data provenance and protection.
- **Navigating the Jagged Technological Frontier (Dell'Acqua et al.), HBS Working Paper 24-013**
  Publisher: Harvard Business School | Date: September 2023 | Reading time: about 15 min
  URL: https://www.hbs.edu/faculty/Pages/item.aspx?num=64700
  Why read: Why AI quality varies sharply by task: the reason testing must be task-specific.

Session 4: Policy and incidents

Required

- **Requirements for Data Breach Notifications**
  Publisher: Mass.gov, Office of Consumer Affairs and Business Regulation | Date: Living page | Reading time: about 10 min
  URL: https://www.mass.gov/info-details/requirements-for-data-breach-notifications
  Why read: Mass.gov on data breach notification: who to notify and the online breach report form.
- **Breach Notification Rule**
  Publisher: U.S. Department of Health and Human Services, Office for Civil Rights | Date: Living page | Reading time: about 10 min
  URL: https://www.hhs.gov/hipaa/for-professionals/breach-notification/index.html
  Why read: HIPAA Breach Notification Rule: individual, HHS, and media notice, and the 60-day outer limit.

Optional

- **Training and Employment Notice 07-25: The U.S. Department of Labor's Artificial Intelligence Literacy Framework**
  Publisher: U.S. Department of Labor, Employment and Training Administration | Date: February 13, 2026 | Reading time: about 10 min
  URL: https://www.dol.gov/agencies/eta/advisories/ten-07-25
  Why read: U.S. Department of Labor AI Literacy Framework (TEN 07-25). Use its five content areas when you plan staff training.
- **Configure a secure and governed foundation for Microsoft Copilot**
  Publisher: Microsoft Learn | Date: Living page | Reading time: about 15 min
  URL: https://learn.microsoft.com/en-us/microsoft-365/copilot/configure-secure-governed-data-foundation-microsoft-365-copilot
  Why read: Microsoft 365 shops: how to set data governance foundations (sensitivity labels, oversharing controls) before broad AI rollout.
- **Generative AI in Google Workspace Privacy Hub**
  Publisher: Google Workspace Admin Help | Date: Living page | Reading time: about 10 min
  URL: https://knowledge.workspace.google.com/admin/generative-ai/generative-ai-in-google-workspace-privacy-hub
  Why read: Google Workspace shops: how Gemini in Workspace handles your organization's data.

Link verification

Links verified by fetch or search as of Oct 2026. Links that could not be confirmed were dropped and are not listed. The instructor re-checks links before each cohort.