AI Governance at Emory University

A Framework for Responsible Innovation

Nabile Safdar, MD MPH | Chief AI Officer, Emory University
Presented to: Office of Research Compliance and Regulatory Affairs
March 2026

Today's Agenda

01 AI Basics & The Changing Landscape

02 Why AI Governance Matters

03 Emory's Tiered Governance Framework

04 Data & AI Governance Advisory Council

05 Responsible AI Workgroup & Subcommittees

06 Enterprise AI Security Policy (Policy 5.3)

07 Resources, responsibleai.emory.edu & Next Steps

Join by Web **PollEv.com /nabiles212**

Join by Text Send **nabiles212** and your message to **22333**

What word describes how you feel when you think about AI

Loading...

Nobody has responded yet.

Hang tight! Responses are coming in.

What Is Artificial Intelligence?

Artificial intelligence (AI) is a set of technologies that enable systems to learn, reason, and perform tasks that traditionally require human intelligence — such as understanding language, recognizing patterns, and making decisions.

1

Machine Learning

Systems that learn from data to improve performance without being explicitly programmed

2

Deep Learning

Neural networks with many layers that power image recognition, speech processing, and complex predictions

3

Generative AI

Models like LLMs (ChatGPT, Claude) that generate text, images, code, and other content from prompts

4

AI Agents

Autonomous AI systems that perceive their environment, reason, and take multi-step actions to achieve goals

How AI Is Transforming Our Landscape

Research

- Accelerates literature review and hypothesis generation
- Automates data analysis and pattern recognition
- Enables multi-modal research across imaging, text, and omics
- AI tools already embedded in research workflows

Healthcare

- Clinical decision support and diagnostics
- Ambient documentation and EHR summarization
- Predictive analytics for patient outcomes
- Operational efficiency in care delivery

Academics

- Personalized learning and adaptive platforms
- AI tutors and writing assistance for students
- Challenges to academic integrity frameworks
- Faculty role evolution and pedagogy shifts

Administration

- Process automation across HR, finance, and operations
- AI-powered contract review and compliance monitoring
- Enterprise knowledge management
- Workforce augmentation strategies

AI tools are already in use across Emory — whether formally sanctioned or not. Governance is not optional.

Why AI Governance Matters

Data Privacy & Security

Sensitive patient, student, and research data may be exposed when using unapproved AI tools. Data may be used to train external models.

Research Integrity

AI-generated content in research requires transparency, reproducibility, and appropriate disclosure per IRB and journal standards.

Regulatory Compliance

HIPAA, FERPA, FDA Software as Medical Device, NIH/NSF data policies — non-compliance carries institutional and individual risk.

Intellectual Property

Automatic rights grants from AI tools may violate Emory's IP Policy. Research findings and institutional data must be protected.

Bias & Fairness

AI models can perpetuate or amplify biases present in training data — particularly consequential in clinical and research contexts.

Institutional Reputation

Uncontrolled AI use can undermine trust with patients, research partners, federal agencies, and accrediting bodies.

Emory's Tiered AI Governance Framework

Three interlocking layers — from enterprise strategy to operational execution

Corporate Governance

Board · Senior Leadership

Sets institutional AI posture, risk tolerance, and strategic direction — the 'who we are' as an institution regarding AI

Enterprise-wide risk appetite and policy approval
Final authority on high-risk AI initiatives

AI Governance (Middle-Out)

Advisory Council · Program Office

Manages policies, procedures, and specific AI case reviews — the connective tissue between strategy and operations

Data & AI Governance Advisory Council
Responsible AI Workgroup
Policy development and standards
Approves data use/sharing and AI use cases
Escalates high-risk decisions upward

Operational Governance

Schools · Units · Departments

Each area's AI roadmap, priorities, and local implementation — where governance meets day-to-day practice

AI Stewards Forum & Data Stewards Forum
Domain-specific roadmaps and priorities
Local risk assessment and triage
Escalation pathways to Tier 2

Corporate Governance

Board & Senior Leadership — Setting Institutional Posture

TIER 1

Purpose & Authority

- Sets Emory's overall AI strategy and risk tolerance as an institution
- Defines the institutional 'posture' — how aggressive or conservative we are in AI adoption
- Approves enterprise-wide AI policies and escalated high-risk decisions
- Ensures AI investments align with Emory's mission in research, education, and healthcare
- Provides accountability to the Board and external stakeholders

Who Sits at This Level

Board of Trustees

Ultimate fiduciary oversight; sets institutional values and risk boundaries for AI at the enterprise level

President, Provost, EVPs

Executive authority over institutional strategy and policy; signals AI posture and commitment enterprise-wide

C Suite

Academic and research mission stewardship; ensures AI governance aligns with faculty, research, and student affairs

AI Governance — Middle-Out

Policies, Procedures & Specific Case Management — The Connective Tissue

TIER 2

"Middle-Out" governance sits between strategic vision and day-to-day operations. It translates enterprise policy into workable procedures, manages specific AI use cases and data use requests, and connects operational teams with executive leadership.

Policy & Standards

- Develops enterprise AI policies and data governance standards
- Maintains the Data and AI Governance Charter
- Defines responsible AI frameworks aligned with institutional values
- Sets cross-functional guidelines for AI use, data classification, and security

Case Review & Approvals

- Reviews and approves data use/sharing proposals
- Evaluates AI use cases across the four governance pillars: Application, Model, Data, Infrastructure
- Manages AI model re-classification requests (e.g., de-identified PHI)
- Escalates high-risk decisions to Tier 1

Coordination & Alignment

- Bridges University and Healthcare domains ('One Emory')
- Coordinates between Legal, Compliance, Security, Research, and Digital
- Supports escalations from Tier 3 Stewards Forums
- Reports progress and recommendations to the Steering Committee

Data & AI Governance Advisory Council

Strategic advisory body to the Steering Committee — bridging EU and EHC

The Council provides expert guidance and recommendations on data and AI governance across Emory University and Emory Healthcare — ensuring decisions reflect compliance, legal, security, research, and operational perspectives.

Co-Chairs

Abdoul Aziz Sosseh
Chief Data Analytics Officer | EU/EHC

Terrie Estes
SVP Chief Compliance Officer | EU/EHC

Dr. Nabile Safdar
Chief AI Officer | EU/EHC

Voting Members

- Melissa Hall — VP & Chief Compliance, Emory University
- Derek Spransy — Chief Information Security Officer
- Todd Sherer — Associate VP, Research (ORA)
- Scott Shacter — Sr. Dir. Sourcing/Contracting

Key Functions

- Approves data use/sharing proposals and AI use cases
- Reviews AI model re-classification requests (including de-identified PHI)
- Advises Steering Committee on enterprise policies and standards
- Coordinates compliance, legal, security, data, and AI perspectives in a single body
- Non-voting advisory members include legal counsel, data & AI strategy, cybersecurity, licensing

Responsible AI @ Emory: Enterprise Workgroup

Revised charter — cross-functional enterprise initiative with four specialized subcommittees

The Workgroup guides the ethical, compliant, and strategic use of AI across Emory — fostering collaboration, mitigating risks, and shaping institutional policies and practices. The full workgroup meets quarterly; subcommittees meet monthly.

Healthcare Subcommittee

- AI in clinical care and patient safety
- Health data governance (PHI, EHR)
- Clinical decision support oversight
- Led by: Terrie Estes (EHC Compliance)

Research Subcommittee

- Research administration & deans
- Investigator guidance on AI in research
- IRB coordination and data compliance
- Includes ORA, RCRA representatives

Academics Subcommittee

- Schools and Campus Life representatives
- AI in teaching, learning & student experience
- Academic integrity frameworks
- Provost's office coordination

Administration Subcommittee

- HR, Risk Management, and IT units
- AI impact on workforce and enterprise systems
- Operational AI deployment governance
- Enterprise risk assessment

Workgroup Scope, Deliverables & RCRA Representation

Scope Areas

Risk Mitigation

Identifies and assesses risks across healthcare, research, academics, and administration

Regulatory Compliance

Monitors AI-related regulations at local, national, and international levels; ensures alignment with laws and accreditation standards

Policy Evaluation

Reviews and recommends updates to institutional policies and AI workflows

Stakeholder Engagement

Hosts workshops, training sessions, and speaker series to promote responsible AI literacy

Deliverables & RCRA Members

Deliverables (reported to Executive Committee)

- Periodic reports on subcommittee activities, risk assessments, and recommendations
- Centralized repository of best practices, policies, and educational resources
- Formal policy update recommendations submitted to the Executive AI Governance Committee

RCRA Representatives in Research Subcommittee

Tier 3: Operational Governance

Domain-Level Roadmaps, Priorities & Day-to-Day AI Management

TIER 3

Domain Execution Bodies

AI Stewards Forums

Monitors AI model performance and bias, maintains the AI inventory, ensures compliance with Responsible AI standards, and escalates AI issues to the Advisory Council

Data Stewards Forum

Triages local data quality issues, maintains data definitions, resolves cross-domain data conflicts, trains data custodians, and reports issues to the Advisory Council

What Operational Governance Looks Like

- Each school, unit, or department develops its own AI roadmap and priorities
- Local teams identify AI use cases and initiate governance review if needed
- Stewards maintain domain-specific AI and data inventories
- First-line risk assessment — triaging what needs escalation versus local resolution
- Responsible AI Workgroup subcommittees (Research, Healthcare, Academics, Administration) operate at this level
- Accountability flows upward through clear escalation pathways

Enterprise AI Security Policy — Policy 5.3

Effective March 5, 2026 | Applies to all Emory faculty, staff, and students

EASAT — Emory Approved Secure AI Technology

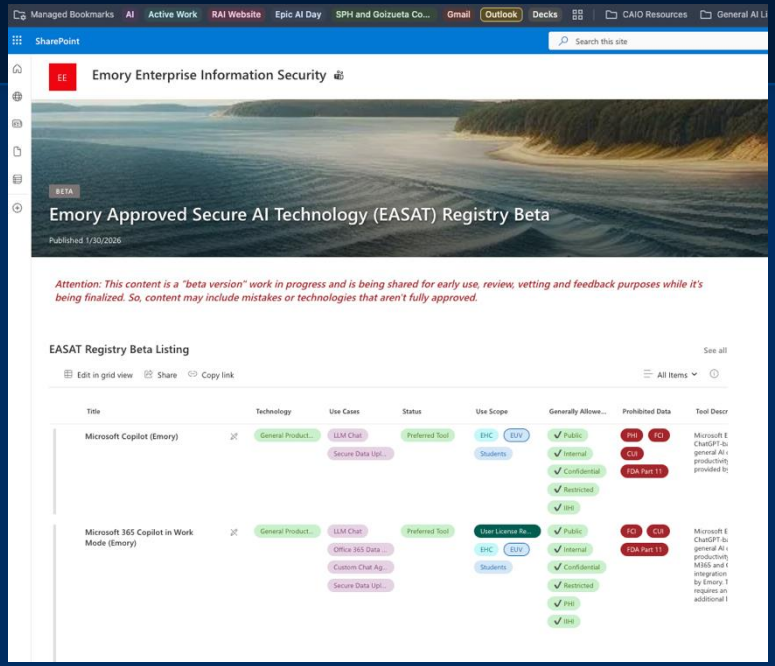

Research-Specific AI Compliance Considerations

What Policy 5.3 and Emory governance mean for research compliance officers and investigators

Data Security Requirements for Research

Emory's data classification (Internal → Confidential → Restricted) directly governs which AI tools researchers may use. Federal requirements (NIST 800-171, CMMC, NIH genomics policy) layer on top for federally funded work. See Appendix for full federal map.

IRB-Specific EASAT Pathway

AI tools can qualify as EASAT for a specific IRB-approved protocol with written OIT Security review evidence. This enables research-appropriate tool use even if not on the general registry.

Confidential & Restricted Data in Research

PHI, genomic data, student records: EASAT-approved tools only. Federal award data involving CUI may also require NIST 800-171-compliant environments — not just EASAT approval.

AI-Generated Content Disclosure

Any AI-generated text, figures, tables, or code in research outputs must carry a clear disclosure notice until appropriate human review and validation has occurred.

De-identification & Re-classification

AI models trained on de-identified PHI remain Confidential Data unless the Advisory Council approves re-classification. Full documentation of the de-identification methodology is required.

AI Code in Research Pipelines

AI-generated code used in research data pipelines requires documented SME review. Degree of rigor scales with data sensitivity — higher standards for Restricted/CUI Data.

Intellectual Property

Automatic rights grants from AI tools may violate Emory's IP Policy. Research findings and novel discoveries generated with AI require careful review of tool terms of service.

responsibleai.emory.edu

Emory's Central Hub for AI Governance, Policy, and Resources

https://responsibleai.emory.edu

AI Policies & Guidelines

Enterprise AI Security Policy (5.3), Responsible AI principles, and institutional frameworks — all in one place

EASAT Registry

Browse and search the list of AI tools approved for use at Emory, including approved data classifications and scope limitations

Governance Structure

Full documentation of Emory's tiered AI governance bodies, charters, and escalation pathways

Education & Training

Workshops, guidance documents, and training resources on responsible AI use for researchers, faculty, staff, and students

Submit a Request

Pathways to submit AI use cases, data use proposals, EASAT applications, and requests for the Advisory Council

Contact & Support

Geoffrey Parsons (OIT Security) and the Data Governance team are the primary contacts for policy and EASAT questions

responsibleai.emory.edu — Live Site

https://responsibleai.emory.edu | Emory's central hub for AI governance

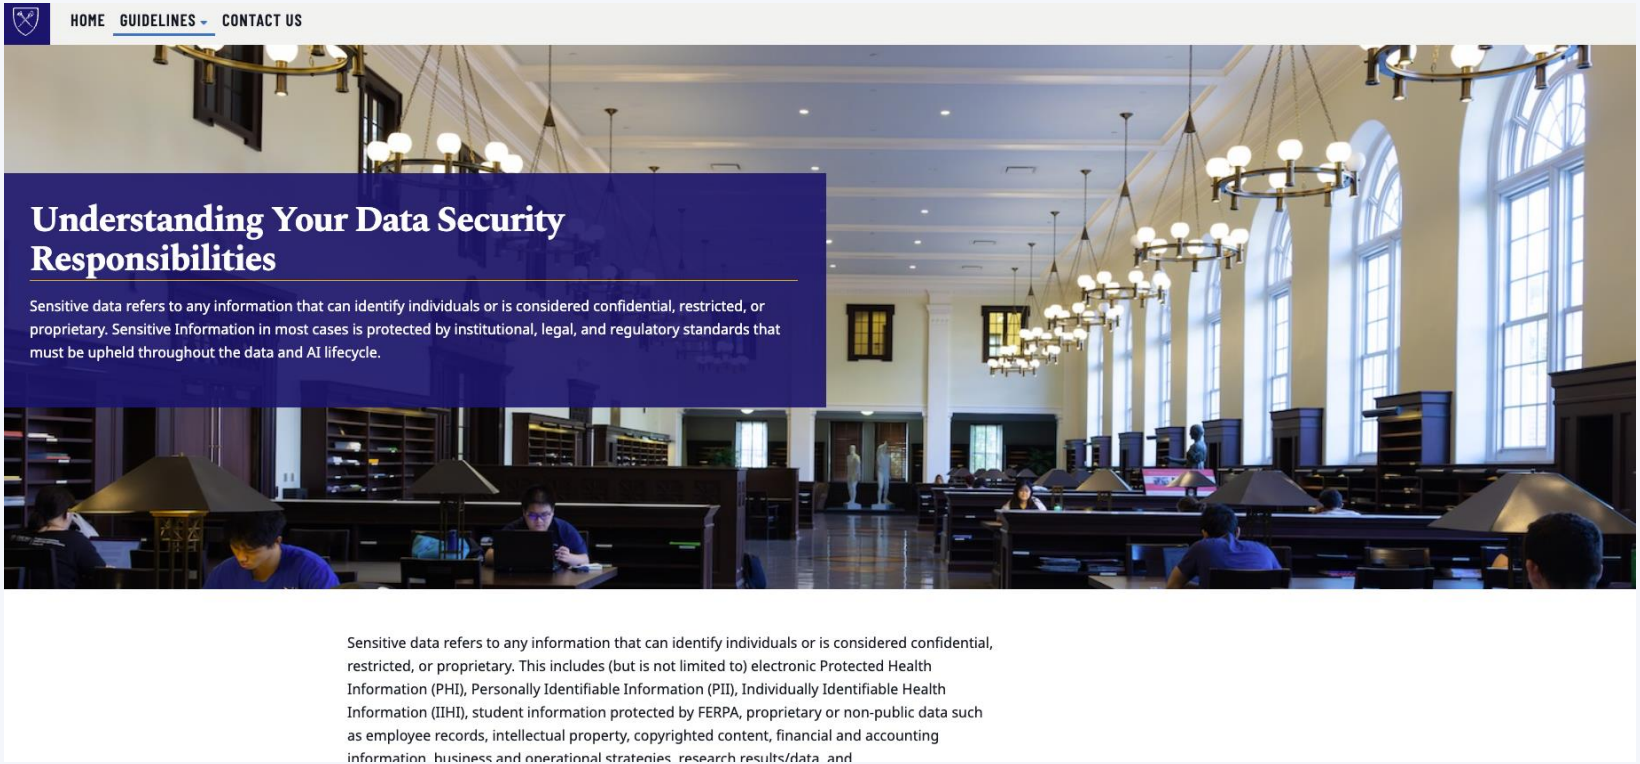

Key Takeaways

1

AI governance at Emory is tiered and comprehensive

Three layers — corporate, middle-out, and operational — ensure accountability flows from board-level strategy down to day-to-day practice in each unit.

2

Research compliance is embedded in the governance structure

RCRA representatives sit on the Research Subcommittee. IRB-specific EASAT approval pathways exist specifically for research. The Advisory Council handles data use cases.

3

Policy 5.3 is in effect — know the requirements

All AI tool use must be EASAT-approved or in process. Research involving Confidential/Restricted Data requires approved tools only. AI code and content carry review obligations.

4

The Responsible AI Workgroup connects strategy to practice

Four domain-specific subcommittees (Healthcare, Research, Academics, Administration) translate policy into actionable guidance for each part of the institution.

5

Start at responsibleai.emory.edu

Policies, the EASAT Registry, governance charters, training resources, and request pathways are all accessible at Emory's central responsible AI hub.

Questions?

Dr. Nabile Safdar, MD
Chief AI Officer, Emory University
nmsafdar@emory.edu

responsibleai.emory.edu

APPENDIX

Federal Research Data Security Requirements

CMMC · NIST 800-171 · NIH Genomics · NSF · FISMA · NSPM-33

Federal Research Data Security: Executive Overview

There is no single universal rule — requirements are triggered by funder, data type, and award structure

Three Triggers for Research Data Security Obligations

1. Who funds the research (DoD vs. NIH vs. NSF vs. other)
2. What kind of data is involved (CUI, genomic, export-controlled, PHI, PII)
3. How the award is structured (contract vs. grant, flow-down clauses, subcontract terms)

CMMC

DoD + CUI only

- Applies only when: DoD funding + CUI/FCI present + DFARS clauses flowed into the award
- Does NOT apply to fundamental research unless CUI is involved
- Level 2 requires NIST 800-171 (110 controls) + formal assessment
- CMMC 2.0 finalized late 2024; phasing into contracts starting 2025

NIST 800-171

De facto cross-agency baseline

- Required by NIH, DoD, NSF, DOE, DHS, DOJ when CUI is present
- Governs CUI on non-federal systems (i.e., university infrastructure)
- Requires: SSPs, access controls, logging, encryption, incident response
- Increasingly required by attestation — not just 'best effort'

NIH Genomics

Strictest for research data

- NOT-OD-24-157 (effective Jan 25, 2025): NIST 800-171 required for controlled-access genomic data
- Institutions must attest to compliance — not just aspire to it
- Applies to cloud platforms and third-party tools
- Compliance costs are allowable and expected in proposals

CMMC: DoD-Only, But Very Real

Cybersecurity Maturity Model Certification — applies when ALL three conditions are met

1. Funding is from the Department of Defense
2. The award involves Controlled Unclassified Information (CUI) or Federal Contract Information (FCI)
3. Requirement is explicitly flowed into the award via DFARS clauses

CMMC Levels

Level 1 — Basic Cyber Hygiene

FAR 52.204-21 (17 controls); FCI only; annual self-assessment

Level 2 — Advanced (CUI)

NIST SP 800-171 (110 controls); C3PAO third-party assessment or government-confirmed self-assessment depending on contract criticality

Level 3 — Expert

NIST SP 800-172+ controls; government-led assessment; very specialized contracts

Practical Impact on Universities

- CMMC applies to specific labs, enclaves, or environments — NOT the whole university
- Common triggers: defense-sponsored AI/cyber/materials/aerospace research; subcontracts from defense primes; UARC-like units
- Fundamental research exclusion: CMMC does NOT apply if no CUI is generated or received
- Emory researchers on DoD awards should check DFARS 252.204-7012/7019/7021 clauses
- Non-compliance with flowed-down CMMC terms is a contract breach — not just a policy violation

NIH, NSF, FISMA & NSPM-33: Research-Relevant Requirements

Non-DoD federal requirements that directly affect Emory research programs

NIH — Now Among the Strictest

NOT-OD-24-157 (eff. Jan 25, 2025)

- NIST SP 800-171 required for controlled-access human genomic data
- Institutions must formally attest to compliance — aspirational compliance is insufficient
- Applies to cloud platforms and third-party AI tools used with NIH genomic data
- Costs for secure environments are allowable and expected in proposals
- Emory RCRA has published institution-specific guidance on secure enclave requirements

NSF — Quieter, But Real

Participates in federal CUI program

- NSF-funded research involving sensitive infrastructure data, international datasets, or PII at scale may involve CUI
- When CUI is present: proper marking, controlled sharing, secure storage and destruction required
- NSF does not generally mandate CMMC, but expects NIST-aligned protections for CUI
- NSF data management plans should address CUI identification and handling

FISMA — Indirect but Foundational

Applies to agencies; flows down to grantees

- FISMA applies to federal agencies directly, not universities
- However, agencies use NIST 800-53 / 800-171 to meet FISMA — and flow those down
- Universities are not 'FISMA-compliant' per se, but must implement FISMA-derived controls
- System Security Plans (SSPs) required when FISMA-derived controls apply

NSPM-33 & CHIPS Act — Research Security

Institutional research security overlay

- Requires institutions receiving large federal funding volumes to have formal research security programs
- Covers: cybersecurity, export controls, research security training, foreign talent programs
- Does not mandate CMMC directly, but reinforces expectation of data protection capability
- Institutions may need to certify their security capability to federal agencies

Federal Requirements Comparison & AI Implications

Quick-reference map and why AI changes everything for research data governance

Requirement	Applies When	Core Standard	Emory Trigger Examples
CMMC Level 2	DoD research + CUI/FCI present	NIST SP 800-171 (110 controls)	Defense-sponsored AI, materials, cyber labs
NIH Genomics (NOT-OD-24-157)	Controlled-access NIH genomic data	NIST SP 800-171	Any study using dbGaP or NIH controlled repositories
NSF CUI	NSF awards with CUI data	CUI program + NIST-aligned controls	Sensitive infrastructure, international datasets, PII at scale
FISMA (indirect)	Federal flow-down to grantees	NIST 800-53 / 800-171	Most federally sponsored research with sensitive data
NSPM-33	Large federal funding recipients	Programmatic security controls	Institutional-level compliance program required

Why This Matters for Responsible AI & Research Governance

- AI training data can become CUI — and inherit its security obligations
- Model outputs may inherit the data classification of what was used to train them
- Cloud AI platforms (including EASAT tools) must be vetted against NIST controls when CUI is present
- 'Research vs. operations' is no longer a safe distinction — AI blurs those lines