Georgia Tech AI Tool Approval Standard (Draft)

Standard Version: 260415-006

1. Purpose and Philosophy

The Georgia Institute of Technology (**Georgia Tech, GT, Institute**) recognizes Artificial Intelligence (**AI**) as a powerful catalyst for innovation, productivity, and efficiency. To govern this, the Institute maintains the *Artificial Intelligence (AI) in Administrative and Employment Contexts Policy* and the paired *AI in Academic and Research Contexts Policy*, which together define the overarching scope and rules for the procurement, development, and use of AI across all institutional contexts.

Consistent with University System of Georgia (**USG**) mandates, these foundational policies require that the procurement, implementation, and use of all Institute technology—including free, open-source, locally developed, and publicly accessible AI Tools—must be formally assessed for risk prior to procurement and use.

This *AI Tool Approval Standard* (**Standard**) directly operationalizes that policy mandate. The specific type of assessment and the level of scrutiny required will scale based on the AI Tool's intended use and risk profile. While low-risk AI Tools may undergo an expedited review, higher-risk implementations may necessitate cybersecurity reviews, privacy assessments, and, where applicable, contract negotiations.

The purpose of this Standard is to balance the Institute's need for the rapid adoption, use, and innovation of AI Tools to enhance teaching, learning, research, and administrative operations, while balancing the obligation to protect Institutional values, secure data, and comply with all Institute and USG policies as well as applicable laws and regulations.

To achieve this, the Standard establishes a risk-based AI Tool Review Process (**Review Process**).

The Review Process is the comprehensive triage and evaluation workflow detailed throughout this document. It is a formalized system comprised of required data inputs, standardized decision trees, and the designated human subject matter experts (including Cybersecurity Governance, Risk and Compliance, Procurement, Privacy and others, as defined below) who must operate or participate in the process.

The Review Process actively conducts risk and security profiling, identifies necessary mitigations, and determines the appropriate approval paths based on the AI Tool's intended use. This ensures that AI Tools are assessed efficiently and accurately. The decision trees allow routine requests for AI Tool approval to be processed consistently and with greater speed, while the routing logic automatically escalates complex AI Tools—such as those requiring broad Institute deployment or sensitive data processing—directly to the appropriate stakeholders for further security and privacy reviews.

Delineation for Intelligent Automation: Intelligent Automation is a broad category that encompasses standard Artificial Intelligence, Agentic AI, and Robotic Process Automation (**RPA**). While all these

technologies carry inherent enterprise risk, their governance pathways differ based on the level of human oversight.

This Standard exclusively governs discrete AI Tools operated by human users. Implementations where AI models or software bots are given agency to autonomously execute multi-step workflows without or with limited human oversight (such as Agentic AI or RPA) carry unique operational risks. Therefore, these autonomous implementations must be assessed, deployed, and used in accordance with established Institute automation governance standards.

2. Roles and Responsibilities

AI accountability is a shared responsibility across the Institute. To ensure a compliant review process, this Standard defines the following key roles:

The Front Door and Local Administration:

- **GT Community Members:** Employees (including faculty and student employees) and affiliates (**GT Community Members or Users**) are responsible for submitting proposed AI Tools through a central intake form before use for any Georgia Tech work, business, or Institutional purposes. Such Users must not independently purchase, expense, or accept the Terms of Service (**ToS**) for unapproved AI Tools.
- **Local AI Points of Contact (LAIPOCs):** Designated unit-level facilitators who serve as the primary liaison between their unit and the AI Governance Program. LAIPOCs do not conduct independent security reviews; rather, they help shepherd their unit's use-cases through the Review Process. When the Review Process' rules determine an AI Tool is safe for use, the LAIPOC administratively executes the approval, expedites access, and records the **Limited Use AI Tool** on their unit's **Local AI Register**.

Review Process Orchestration:

Cybersecurity Governance, Risk, and Compliance (**Cybersecurity GRC**) acts as the orchestrator of the Review Process. Cybersecurity GRC evaluates a proposed AI Tool's technical characteristics and packages these into a profile in order to collaborate with domain experts on routing and approval decisions. Key domain experts are noted below.

- **Cybersecurity GRC:** The central team responsible for shepherding requests, conducting technical security and network assessments, building risk profiles, and translating complex technical AI risks into actionable decisions for privacy and data stakeholders and for legal review.
- **Procurement and Business Services (Procurement):** Acts as the contract gateway. Procurement evaluates the AI Tool vendor's written agreement and ToS to ensure the contract language legally enforces the Institute's requirements. Where appropriate, contract language may limit the vendor from training foundational models on Georgia Tech data, guarantee data deletion, protect Georgia Tech's intellectual property (**IP**), and include other Institutional imperatives.
- **Data Stewardship & Data Governance Program:** Data Domain experts who collaborate with Cybersecurity GRC to evaluate how specific combinations of AI Tools and technical safeguards impact Institutional data. These domain experts define the boundaries for what specific categories of Protected Data (including regulated and sensitive) are permissible within certain AI pathways.
- **Data Privacy Program (Privacy):** Engages strategically when the Review Process identifies that an AI Tool will process large volumes of Protected Data, includes sensitive Personally

Identifiable Information (**PII**), or presents significant privacy risks. In these high-risk scenarios, Privacy helps conduct a formal Data Protection Impact Assessment (**DPIA**).
- **Office of General Counsel (Legal):** Advises on legal, regulatory, and IP risks, monitors shifting AI laws, and provides input on exceptions to vendor agreements.

The Institute Oversight Layer:

- **The AI Governance Officer:** Leads the overarching AI Governance Program and provides programmatic oversight of the Review Process. The AI Governance Officer establishes the overarching standards, approach, routing logic, and the necessary stakeholders required to evaluate AI Tools. While Cybersecurity GRC performs orchestration and risk profiling, the AI Governance Officer monitors the Review Process to ensure its routing logic and procedures remain aligned with Institute policies. They facilitate necessary escalations and ensure the Review Process operates correctly within its approved boundaries.
- **The AI Governance Committee (Committee):** The Institute-level oversight body responsible for formally approving the standards, approach, and logic of the Review Process recommended by the AI Governance Officer. The Review Process is authorized to operate without the Committee's direct involvement, provided it functions within its defined scope and parameters. The Committee's direct review is reserved for novel use-cases, high-risk Institute deployments, and formal escalations flagged by the Review Process or the AI Governance Officer.

3. Standardized Intake Requirements

To initiate the Review Process, GT Community Members must submit their request through the central intake form. The Review Process relies on specific facts to calculate risk; therefore, requestors are required to submit the following foundational input before a review can begin:

- **AI Tool Name and Vendor:** The basic identification of the software or platform.
- **Specific Edition and Licensing Tier:** AI Tools cannot be evaluated solely by their brand name. The requestor must specify the exact tier they intend to use (e.g., Public Free Edition, Individual Pro/Plus Subscription, or Georgia Tech-Licensed Enterprise Tier).
- **Integration Type:** Is this request for a net-new AI product, or for enabling an embedded AI feature within an existing software platform?
- **Hosting Environment:** Indicate whether the AI Tool is hosted on local Georgia Tech infrastructure or the vendor’s infrastructure.
- **Means of Access:** How the User will interact with the AI Tool, specifically noting if access is via a vendor's web browser interface, a mobile app, or a direct API integration.
- **Data and Business Purpose:** A brief but specific description of the Institutional data and/or third-party data that will interact with the AI Tool and the primary organizational benefit.
- **Dataset Aggregation:** Will the User be uploading, cross-referencing, or combining multiple distinct datasets or data sources into this AI Tool?
- **Ethical Safeguards:** A brief explanation or acknowledgment of how the User will apply human oversight and mitigate risks such as hallucination or bias.

Mandatory Credential Protection Attestation: At the time of intake, the requestor must formally attest that they are not currently using, and will not use, their Georgia Tech organizational email addresses or login credentials to access or register for unapproved, free, or personally purchased tiers of AI Tools.

4. The Risk Assessment Components

AI Tools vary widely in how they process information, where they send data, and what legal terms govern their use. For example, a locally hosted open-source model analyzing public research data carries a fundamentally different risk profile than a public web AI Tool analyzing employee accommodation records. To accurately capture these differences and match the AI Tool to the correct level of scrutiny, the Review Process evaluates proposed AI Tools across four primary risk variables during the central intake process:

A. Categorization of the AI Integration. The Review Process first evaluates the nature of the AI to determine the necessary level of review:

- **AI-First Tools:** Purpose-built models (e.g., ChatGPT, Claude) where AI is the core functionality. These AI Tools go through the full technical and data scoping assessments.
- **Embedded AI in Existing Tools:** When an already-approved vendor integrates new AI capabilities into a standard software platform the Institute already uses. Because the overarching software is already governed, the Review Process bypasses a full vendor review from scratch. Instead, the Review Process strictly evaluates the new AI module to determine its risk profile—ranging from low-risk incidental features (like enhanced search algorithms) to high-risk generative models that may attempt to share Protected Data back to the vendor for training. These embedded AI components must have their data boundaries and associated risks specifically vetted by the Review Process before the new functionality is enabled for Users.
- **Incidental AI:** Basic search algorithms, standard spell-checkers, or ubiquitous operating system features. The Review Process identifies these as low-risk, incidental features and generates a standard "Approved Scope."

B. State-Banned Vendors and Products, Intellectual Property, and Contractual Risk. The Review Process evaluates AI Tools for legal and contractual risks before any use is authorized.

- **State-Banned Vendors and Products:** State law strictly prohibits the use of software and services from countries the U.S. considers adversarial or geopolitically sensitive. Any software, product, or service originating from vendors designated under GTA Standard SS-22-002—such as DeepSeek and Baidu—will be automatically rejected by the Review Process, even if downloaded for local use.
- **Credential Protection and Domain Conflicts:** GT Community Members must never use their Georgia Tech organizational email addresses or login credentials to access or register for unapproved, free, or personally purchased AI Tool tiers. Beyond the immediate data exfiltration risk, this creates administrative conflicts if the Institute later executes an Institute contract and needs to claim that domain to provision licensed seats.
- **The Risks of Public Terms of Service:** Even if an AI Tool is free and not on the banned list, GT Community Members must never independently click "accept" on a public ToS or click-wrap agreement. Doing so introduces two enterprise risks:
    1. **Intellectual Property Loss:** Public AI Tools often include ToS language that claim ownership over uploaded materials, inadvertently surrendering Georgia Tech's or a third-party's proprietary information and/or IP to an external vendor.

2. **Compromising Future Institute Negotiations:** When individuals bypass the central intake and accept public terms, it severely weakens the Institute's bargaining power. All AI Tools must go through the Review Process so the Institute can negotiate on our terms, requiring vendors to guarantee data deletion, implement privacy safeguards and opt-outs for model training, and other Institutional imperatives.

C. The Deployment Pathway. An AI Tool's risk profile and its eventual approved data scope depend heavily on what it is and how it is accessed. The Review Process evaluates the Tool's deployment pathway to determine data exfiltration risks:

- **Direct Vendor Access (Editions and Tiers):** AI Tools cannot be evaluated solely by their brand name; the specific licensing tier fundamentally alters the risk. If accessed via a Public Free Tier, the data-exfiltration risk is high, and the AI Tool is strictly restricted to Public Data. However, if accessed under a formally negotiated Georgia Tech-Licensed Enterprise Tier, the contract requires verified data protections and model-training opt-outs, allowing the AI Tool's approved scope to potentially include Protected Data.
- **Means of Access:** The Review Process also factors in how the User will interact with the AI Tool. Direct Application Programming Interface (API) integrations often carry fundamentally different ToS and data handling implications compared to the same vendor's public web browser interface or mobile app.
- **Enterprise Gateways:** There may be multiple pathways to access the same AI model, and each pathway carries different data handling implications. Accessing third-party models (e.g., Claude) through secure, Institute-contracted cloud environments – whether via an end-user interface like Microsoft Copilot or a developer platform like Azure Foundry – ensures data remains within the Institute's contracted tenant. This secure pathway allows for the greater potential of an approved scope that permits Protected Data, whereas accessing that same model directly through the public vendor's platform is generally restricted to public or non-sensitive data.
- **Locally Hosted & Homegrown Models:** Custom models built from scratch by Georgia Tech personnel or downloaded open-source models run entirely on local Georgia Tech infrastructure will bypass the Procurement contract review, as data does not exfiltrate to an external vendor. However, the Review Process must still route these to Cybersecurity GRC for a mandatory network safety and vulnerability assessment before they can be fast-tracked to the Local AI Register.

D. Data Domains and Derived Information. The Review Process recognizes that not all Protected Data carries the same operational or privacy risk. Rather than getting bogged down in specific, isolated data elements, the Review Process evaluates the broader domains, sub-domains, or collections of data types a User intends to process. Differentiating between the domain of general academic operational data versus a collection of sensitive employee health records allows the Review Process to accurately match the right AI Tool with the appropriate security safeguards.

Furthermore, the assessment process evaluates the risk of derived information. Because AI Tools can aggregate and analyze multiple sources of data to derive sensitive insights, the Review Process is configured to flag use-cases where a User intends to cross-reference disparate datasets. If a User proposes combining seemingly harmless collections of data (such as anonymized donor lists and public tax records), the Review Process escalates the request to Privacy and Data Stewards to evaluate whether the combination could accidentally unmask or expose confidential identities. The Institute explicitly prohibits submitting combinations of data elements into AI Tools that could reasonably identify an individual without proper approval.

5. Establishing Repeatable Approval Pathways

The Review Process does not rely on a single person making ad-hoc decisions for every new AI Tool request. Instead, it utilizes a collaborative, rule-based methodology to build repeatable approval pathways. This methodology ensures consistent security while allowing routine requests to be processed rapidly.

Step 1: Packaging the Variables When a new type of AI Tool or deployment method is introduced, Cybersecurity GRC and Procurement first establish the technical and contractual facts. They determine the AI Tool's deployment pathway and verify the contractual safeguards for the specific edition (e.g., Enterprise Tier vs. Public Free Tier), and validate the means of access (e.g., API vs. web interface).

Step 2: Collaborative Rule-Building Cybersecurity GRC packages these technical facts and presents them to the appropriate Data Stewards, Privacy experts, and Legal. Rather than asking these stakeholders to approve a single brand-name AI Tool or individually assess every possible specific data element, Cybersecurity GRC asks them to evaluate the combination of variables. Together, they determine exactly what data domains or collections of Institutional data are permissible under those specific, verified technical conditions.

Step 3: Creating the Repeatable Pathway Once the stakeholders agree on the data boundaries for that specific combination of technical safeguards, it is integrated into the Review Process as an approved rule. This process establishes a repeatable pathway.

Step 4: Assessment and Expedited Routing When a GT Community Member submits a request through the central intake form, the Review Process plots their specific use-case against these established rules.

- **Established Pathways:** If the User's requested AI Tool, specific edition/licensing tier, deployment method, and data types match a pre-approved combination already established (and routinely verified) by the Institute stakeholders, the Review Process rapidly generates the "Approved Scope" and clears the AI Tool for use. The AI Tool is then automatically routed to the appropriate operational inventory—either directly to the Institute AI Register for broad Institute tools, or to a Local AI Register where a Local AI Point of Contact is available to administratively shepherd the localized access if needed.
- **Novel or High-Risk Pathways:** If a User requests a novel combination of variables (e.g., a new data type mixed with a new deployment method) or proposes a high-risk Institute-wide deployment, the Review Process halts. Cybersecurity GRC must then convene the stakeholders to formally evaluate the new scenario, potentially triggering a DPIA, before a new pathway can be established.

6. The Final Output: Approved Scopes & Dual-Context Boundaries

The Standard does not issue blanket approvals. Successfully reviewed AI Tools are assigned a highly defined "Approved Scope". For every approved AI Tool, the Register will dictate:

- **The Approved Edition and Licensing Tier:** Approvals are strictly bound to the specific evaluated tier. If an AI Tool is approved at the Georgia Tech-Licensed Enterprise Tier, that approval does ***not*** cascade down to the free, public version of the same AI Tool name. Users must strictly utilize the edition authorized on the Register.

- **The Approved Data Scope:** Precisely what categorization of data is permitted to be input into the AI Tool (e.g., Public Data only, or specific Protected Data).
- **The Approved Use Scope:** Who is authorized to use the AI Tool and for what specific business, academic, or research scope.

The Dual-Context Boundary: Approvals are bound by the context in which they were granted. It is highly common for individuals to operate in both Academic & Research (A&R) and Administrative & Employment (A&E) contexts simultaneously. Regardless of context, all AI Tools must be approved through this Review Process. However, an AI Tool approved strictly for a localized research experiment **cannot be repurposed** to analyze administrative human resources data or official student financial records without triggering a formal reassessment.

Appealing a Denied Request: The Review Process' routing rules are designed to protect the Institute, but decisions to deny the use of an AI Tool are not always final. If an AI Tool request is rejected, the requestor has a formal pathway to appeal the decision:

- **Tier 1: Resubmission with New Information:** The fastest method of appeal is resubmission. A requestor may submit an updated intake form that augments the original request with new mitigating safeguards, updated vendor contractual terms, or additional business context that specifically addresses the security or privacy feedback that led to the original rejection.
- **Tier 2: AI Governance Committee Appeal:** If a request is denied during the central review process, the requestor may formally appeal the denial to the full AI Governance Committee for a broader evaluation.
- **Tier 3: Executive Sponsor Appeal:** If the AI Governance Committee issues a denial, the final avenue of appeal rests with the Institute's Executive Sponsors, who serve as the ultimate decision-makers.

7. Relief Valves

Interim Conditional Access (Pending Reviews) Georgia Tech recognizes that Institute contract negotiations and formal DPIAs can take time, and faculty and staff require access to modern tools to remain competitive. If a proposed AI Tool is currently listed as "Pending Review" by the Review Process, GT Community Members ***may*** be granted Interim Conditional Access to use the AI Tool responsibly while the formal review is underway, strictly provided that **all** of the following conditions are met:

- **Not a Prohibited Vendor:** The AI Tool does not originate from a State-Banned Vendor (e.g., DeepSeek, Baidu,) as designated under GTA Standard SS-22-002.
- **Public Data Only:** Only **Public Data** or otherwise Institutionally permitted data is used (no Georgia Tech Protected Data, including Regulated Data such as student records, medical data, CUI, export-controlled data, or sensitive PII).
- **No Model Training:** The vendor allows Users to disable the use of prompts/outputs for model training, and the User formally opts out.
- **Data Ownership:** The Institute, through the User, retains ownership of inputs and generated outputs.
- **Data Deletion:** The vendor allows the deletion of User content and documents its retention practices.
- **Security Standards:** Baseline security protections are in place, including encryption in transit and at rest.
- **Human Oversight:** Users apply human judgment and do not treat AI output as authoritative.