Responsible and Ethical Artificial Intelligence Framework for Virginia Tech

Virginia Tech AI Working Group Final Report

Co-Chairs

Dale Pike, Associate Vice Provost for Technology-Enhanced Learning

David Raymond, Associate Vice President for Security & Identity

Submitted to Sponsors

Executive Vice President and University Provost Cyril Clarke

Executive Vice President and Chief Operating Officer Amy Sebring

Virginia Tech

Blacksburg, Virginia

July 2025

1. Executive Summary

The current AI Working Group presents a framework for responsible artificial intelligence integration that positions Virginia Tech as a leader in higher education AI governance while advancing Virginia Tech's mission and objectives. This framework emerges from analysis of institutional readiness (see Appendix A, particularly the DEC (Digital Education Council) Ten Dimension Assessment results), extensive stakeholder engagement (see Section 2.2 and the ChatGPT Edu pilot feedback in Appendix C.4), and rigorous policy evaluation (see Appendix B.2: Detailed Findings by Domain) that collectively inform actionable recommendations for immediate implementation.

Virginia Tech approaches AI integration from a position of strength. The university's designation of AI as one of four institutional research frontiers, combined with specialized facilities including the Sanghani Center for AI and Data Analytics establishes a foundation for AI advancement in research, which in turn has the potential to inform innovation across all university functions.

Note: *Throughout this document, the working group that authored this report will be referred to as the current AI Working Group. This group recommends establishing a permanent AI Working Committee as part of IT governance. The current AI Working Group will be disbanded upon publication of this final report.*

1.1 Key Findings and Priorities

Analysis reveals implementation challenges requiring prompt action. Virginia Tech's governance infrastructure lacks formal AI policies, creating accountability gaps, while faculty demonstrate high AI adoption intent requiring structured support. Policy review identified five areas needing immediate attention, and pilot program results demonstrate community readiness for AI integration (detailed findings in Appendix A and B).

1.2 Strategic Recommendations

Virginia Tech's approach rests on six building blocks–an ethical framework, a policy gap analysis, practical guidance resources, tool recommendations, a governance model, and a phased roadmap–each feeding the five recommendations that follow. These recommendations address both immediate governance needs and long-term capacity building, with potential for returns on investment as indicated by initial analysis of pilot results.

Five strategic recommendations systematically translate Virginia Tech's research excellence into AI integration:

Immediate Governance and Policy Actions

- Establish AI Working Committee: Form a cross-functional body within the IT governance framework with defined authority and clear accountability mechanisms to provide centralized coordination and oversight for all AI initiatives. This AI Working Committee recommends developing university-wide AI policies and standards and ensuring compliance and risk management.
- Revise Five Critical Policies: Undertake immediate formal review of five foundational institutional policies, including the Undergraduate Honor Code (Policy 6000), Policy on Intellectual Property (Policy 13000), Policy on Misconduct in Research (Policy 13020), Visual

Media Policy (Policy 8205), and Safety and Security Policies (e.g. Policy 5617), as detailed in Appendix B.3: to address potential policy gaps in academic integrity, intellectual property, and operational governance brought about by AI technologies.

- **Pilot Investment**: Invest in a second phase of AI pilot programs, piloting emerging pricing models and applications. These pilots should assess needs for all constituents across the university and rigorously explore different pricing and deployment models for AI tools, building on the success and lessons learned from initial pilots like ChatGPT Edu (see Appendix C for complete results) and Microsoft 365 Copilot (see Section 7.1 for pilot outcomes).

Capacity Building and Resource Development

- **Develop Implementation practical implementation guides**: Create three domain-specific practical implementation guides (for teaching and learning, research, and administrative functions) that translate high-level governance principles into practical, operational guidance for daily AI application across the university.
- **Launch AI Literacy Initiatives**: Deploy a multi-faceted approach to cultivate AI literacy across all constituencies - students, faculty, and staff. These programs will build capacity from basic awareness to advanced implementation, ensuring the community can effectively and responsibly engage with AI.

These investments create the infrastructure necessary for sustainable AI integration while maintaining Virginia Tech's commitment to responsible innovation and its land-grant mission.

1.3 Governance Framework

The recommended governance approach establishes the AI Working Committee as a formal standing working committee within the university's IT governance framework, with delegated responsibility for AI-related issues. This structure must ensure specialized AI expertise and focus while benefiting from integration with university-wide technology governance. The AI Working Committee has responsibility to advise and inform on AI-specific policies, AI risk assessment, and AI implementation guidance while coordinating with IT governance on technology infrastructure, security, and enterprise-wide decisions. This structure addresses the complex requirements of AI governance through:

- **Strategic Oversight**: AI Working Committee within IT governance provides university-wide policy development and coordination
- **Domain Expertise**: Integrated subgroups that leverage Virginia Tech's research expertise across Teaching & Learning, Research, and Operations domains
- **Risk Management**: Three-tier assessment framework matching oversight intensity to potential impact (see supplementary document for detailed risk-tier assessment criteria)
- **Integration**: Seamless alignment with existing IT governance structures

1.4 Implementation Timeline

The approach unfolds through a three-phase approach, prioritizing governance establishment while building sustainable systems for long-term institutional AI integration. This methodology ensures that governance structures are operational before expanding to implementation across all university functions.

Phase 1: Foundation (initial implementation phase)

- Form AI Working Committee and establish governance structure (see supplementary document for AI Working Committee membership specifications)
- Publish AI Charter and risk assessment framework
- Initiate policy revision processes
- Launch initial training programs

Phase 2: Expansion (following resource allocation)

- Complete domain-specific practical implementation guides
- Launch expanded training programs and assessment processes
- Scale pilot programs based on validated metrics
- Begin policy development

Phase 3: Maturation (long-term institutionalization)

- Implement policy updates
- Establish continuous improvement processes and annual review cycles
- Evaluate and expand successful initiatives, transitioning to permanent operational structures

This phased approach ensures progress (see Section 12: Implementation Timeline and Roadmap for detailed schedule) while maintaining flexibility to adapt to emerging opportunities and challenges. Virginia Tech's distinctive combination of research excellence and commitment positions Virginia Tech not merely to adopt AI technologies but to lead in demonstrating how research universities can achieve comprehensive, responsible AI integration that advances their missions while maintaining their values.

2. Introduction and Context

2.1 Virginia Tech AI Working Group Charge and Mission

Executive Vice President and University Provost Cyril Clarke and Executive Vice President and Chief Operating Officer Amy Sebring established the current AI Working Group to address the transformational challenges and opportunities presented by artificial intelligence in higher education. Co-chaired by Dale Pike, Associate Vice Provost for Technology-Enhanced Learning, and David Raymond, Associate Vice President for Security & Identity, the group was charged with developing a framework for responsible AI use across the university's administrative, teaching, outreach, and research operations.

The current AI Working Group's mandate encompassed six interconnected deliverables designed to position Virginia Tech as a leader in responsible AI adoption. These deliverables include 1) the development of an ethical AI framework with integrated risk mitigation and oversight

mechanisms (Section 4.2: Core AI Principles and supplementary document for Risk-Tier Assessment Framework), 2) an analysis of existing university policies to identify gaps and recommend updates (Section 5), 3) the identification of critical areas requiring additional guidance and resources (Section 6), 4) the recommendation of AI tools and platforms for institutional use, based on evaluations and validation from pilot programs including a 425-participant ChatGPT Edu pilot (Section 7.2 and detailed results in Appendix C), 5) designing a sustainable governance structure that aligns with existing IT governance frameworks (Section 4.3: Governance Structure and supplementary document for Integration with Existing IT Governance), and 6) establishing a prioritized implementation timeline and roadmap for ongoing AI initiatives across the university (Section 12).

2.2 Stakeholder Engagement and Community Response

Throughout the current AI Working Group's tenure, extensive stakeholder engagement provided crucial insights into the university community's readiness for and concerns about AI integration. Feedback was systematically collected through presentations across multiple forums, open sessions designed to encourage broad participation, and pilot programs that offered practical experience with AI tools in real-world applications.

Stakeholder engagement provided insights through multiple channels, including a ChatGPT Edu pilot program that received positive reception (detailed results in Appendix C). Feedback focused on practical implementation needs, as participants requested more integration capabilities, training resources, and clear guidance on appropriate usage across different risk contexts.

2.3 Institutional Foundation and Strategic Positioning

Virginia Tech enters the AI governance conversation from a position of strength, with research infrastructure and demonstrated commitment to innovation. The Sanghani Center for AI and Data Analytics serves as the intellectual hub, bringing together 36 faculty members and over 200 graduate students working on 127 active research projects (see supplementary document for detailed capabilities). This interdisciplinary foundation is enhanced by AI's designation as one of four institutional research frontiers, bringing priority funding, strategic hiring, and cross-college collaboration incentives.

The university's research excellence spans multiple domains, from the Virginia Tech Transportation Institute's autonomous vehicle safety breakthroughs to the College of Veterinary Medicine's 95% accuracy in AI-powered cancer detection (see supplementary document for details). Strategic investments, including the 2021 $10 million Sanghani gift and the Commonwealth Cyber Initiative's $25 million annual investment, strengthen these capabilities while the Innovation Campus in Alexandria positions Virginia Tech to shape national AI policy (see supplementary document for strategic positioning details).

While this research excellence provides unique advantages, an institutional readiness assessment identified critical implementation challenges requiring systematic attention (detailed analysis in Appendix A.3: Assessment Summary and Strategic Implications). The assessment findings inform this report's recommendations for translating research strengths into operational capabilities.

2.4 Report Structure and Recommendations

Building on the findings of the AI Readiness Assessment (see Section 3 and Appendix A), this report proposes six interconnected components that form the foundational framework for responsible AI integration. These components translate readiness insights into concrete governance, policy, resource, and implementation actions. From them, this report derives five key actionable recommendations to guide immediate and long-term efforts:

#	Strategic Component	Purpose & Highlights
1	**Responsible & Ethical AI Framework**	Establishes core principles, a risk-tier assessment process, and supporting resources that translate ethics into day-to-day practice.
2	**Policy Gap Analysis**	Reviews 168 policies, pinpoints gaps, and prescribes formal revisions or supplemental guidance to ensure AI alignment across academic, research, and operational domains.
3	**Guidance & Capability Development Resources**	Delivers three domain-specific practical implementation guides and a tiered AI-literacy program that builds practical skills for faculty, staff, and students.
4	**Evidence-based Tool & Platform Recommendations**	Recommends AI services (e.g., Microsoft Copilot, ChatGPT Edu) validated through pilot data showing user-reported productivity gains.
5	**Adaptive Governance Model**	Creates a AI Working Committee embedded in existing IT governance to coordinate policy, oversee risk, and steward resources.
6	**Phased Implementation Roadmap**	Lays out short-, mid-, and long-range milestones that coordinate approvals, resource rollouts, policy updates, and ongoing refinement.

Together, these six framework components provide a coherent, actionable pathway for responsible AI adoption, anchored in principled governance, strengthened by empirical pilot evidence, and supported by detailed templates in the appendices. The five actionable recommendations (as outlined in the Executive Summary) are the direct initiatives for implementing this pathway.

3. Virginia Tech AI Readiness Assessment

3.1 Readiness Assessment Overview

Virginia Tech's AI Working Group applied the Digital Education Council (DEC) Ten Dimension AI Readiness Assessment between February and April 2025 (see Appendix A, Assessment Methodology for complete framework details). The instrument, published in 2025 by a DEC working group of 34 universities, offers descriptive rubrics rather than statistically normed scores. Formal reliability or validity studies have not yet been released. The analysis was based on several facilitated workshops, document analysis and follow-up interviews, as well as drawing on preliminary infrastructure inventories and a spring 2025 faculty survey (n = 118).

3.2 Key Findings of Readiness Assessment

Virginia Tech's AI Readiness Assessment paints a dual portrait: the university already operates from strong research platform yet translating that strength into institution-wide impact requires targeted investment in policy, operations, and community capability. The assessment confirms that Virginia Tech is not starting from scratch; instead, it should build bridges between established research excellence and the day-to-day practices of teaching, learning, and administration.

Building on this context, five findings stand out as the most consequential for planning and resource allocation (detailed analysis available in Appendix A, Detailed Assessment Results by Dimension):

- Research Leadership - Established:
  Virginia Tech's designation of AI as one of four research frontiers and extensive research portfolio demonstrates strong research leadership (detailed analysis in Appendix A).
- Faculty Engagement - High and Broad-Based:
  Faculty demonstrate strong AI adoption intent, with usage patterns showing exploratory rather than ubiquitous adoption (see supplementary document for detailed usage data).
- Governance Infrastructure - Emerging and Urgent:
  No formal AI policies currently exist, making governance establishment the top priority (detailed policy analysis in Appendix B.1).
- Operational Readiness - Strong Hardware, Limited Support Systems:
  Powerful infrastructure exists but lacks support systems for institution-wide impact (detailed infrastructure assessment in Appendix A.2).
- AI Literacy & Capability Development - Promising Start, Needs Scale:
  Multiple training modules exist but comprehensive, role-specific programs are needed (detailed capability analysis in Appendix A.2).

Intent outpaces current practice: 75% of all VT faculty responding to the survey–including many present non-users–expect to incorporate AI in future teaching (vs. 86% globally) (see supplementary document for barriers analysis). Coupled with VT's strengths in engagement and integrity applications, this forward-leaning intent argues for targeted capability-building rather than wholesale persuasion.

These findings confirm Virginia Tech's competitive advantage in research while clarifying where governance, operational infrastructure, and community capability should evolve to achieve responsible, campus-wide AI integration.

1. Scale adoption - move from the current 52% already using AI to the 75% who intend to adopt.
2. Leverage differentiators - material creation & administration dominate, but VT's edge lies in student-engagement and integrity use cases.
3. Deepen usage - most adopters sit at limited or moderate intensity; discipline-specific workflows can unlock higher impact.
4. Remove practical & perceptual barriers - time, guidance and risk reassurance are bigger hurdles than technology.

3.3 Implications

Virginia Tech's research excellence (detailed in Appendix A, Virginia Tech's AI Ecosystem) provides a unique advantage for AI adoption that distinguishes the university from institutions pursuing reactive implementation approaches. The university's AI expertise and research credibility enable it to lead responsible AI integration across all operations.

This distinctive foundation enables three critical success factors that should determine the effectiveness of institutional AI adoption.

- First: The university must systematically translate its research leadership into institutional transformation support, leveraging scholarly expertise to inform practical implementation decisions.
- Second: The university must establish permanent governance structures to address the implementation gaps identified in the readiness assessment, ensuring that research excellence translates into operational effectiveness.
- Third: Sustainable systems must be developed to connect research capabilities with operational needs, creating pathways for knowledge transfer and practical application.

To ensure that advancements in artificial intelligence effectively support both academic research and educational goals, it is recommended to establish formal collaboration mechanisms between research and both administrative and educational units.

3.4 Implementation Strategy

The readiness assessment findings (summarized in Appendix A.3: Assessment Summary and Strategic Implications) directly inform an implementation strategy designed to leverage Virginia Tech's strengths while addressing identified development needs. This strategy emphasizes continuous improvement through regular readiness assessments that track progress across all ten dimensions of the DEC instrument, ensuring that implementation efforts maintain alignment with institutional capabilities and emerging requirements.

The implementation approach establishes baseline metrics and cross-functional teams specifically designed to address identified gaps through collaborative action. These teams will leverage the university's collaborative strengths to build implementation capacity, transforming

theoretical expertise into practical operational capabilities. Assessment insights will be systematically analyzed to prioritize implementation actions, creating a responsive management framework that adapts to evolving institutional needs while maintaining coherence and measurable progress toward AI integration.

See Appendix A for complete readiness assessment methodology (Section A.1), detailed findings by dimension (Section A.2), and strategic implications (Section A.3).

4. Responsible and Ethical AI Framework

4.1 Framework Overview

Virginia Tech's Responsible and Ethical AI Framework establishes an integrated approach to responsible AI use across the university's teaching, research, and administrative functions. The framework builds upon Virginia Tech's Principles of Community (https://www.vt.edu/principles-of-community.html) and institutional values while providing practical guidance for implementation. This initial, living framework is designed to guide current efforts while evolving through ongoing experience, feedback, and technological and regulatory change.

The Responsible and Ethical AI framework consists of four interconnected components: **Principles** (articulating fundamental values and commitments), **Governance** (establishing clear accountability and decision-making structures), **Processes** (addressing operational implementation through concrete methods), and **Resources** (providing practical tools and materials for responsible AI use).

4.2 Core AI Principles

As a land-grant research university, Virginia Tech is committed to *Ut Prosim*–That I May Serve. Artificial intelligence amplifies our ability to teach, discover, and serve the Commonwealth and the world, but only when its adoption reflects our Principles of Community and the public trust those principles embody.

These principles emerged from months of collaborative development and extensive stakeholder engagement across the university community. They are designed to serve as enduring guidance–a "constitution" for AI governance–while remaining flexible enough to adapt as technology and our understanding evolve.

The Seven Core Principles

Each principle below includes both a plain-language commitment and the fuller context necessary for meaningful implementation:

1. Mission Alignment

We use AI only in ways that further Virginia Tech's teaching, research, and outreach mission and honor our Principles of Community.

The use of AI should support Virginia Tech's core missions and align with our Principles of Community, enhancing learning, fostering discovery, promoting engagement, improving the

human condition, advancing knowledge, and mitigating potential harms. Every AI implementation should be evaluated against these fundamental purposes.

2. Innovation for Good

We explore new AI tools boldly while weighing benefits against risks.

We will foster a culture of continuous learning and improvement and embrace working with innovative technologies to create new opportunities for our students and employees. We acknowledge that the inherent complexity and rapidly evolving nature of AI technologies will necessitate an approach in which we continually question, analyze, and evaluate AI solutions for robustness, suitability, and benefits that outweigh risks.

3. Human-Centered Benefit

AI should extend–not replace–human insight, creativity, and well-being.

AI serves us best when its purpose is clear and its tools are well-matched to complement or enhance outcomes–whether by informing decisions, strengthening results, or improving efficiency–while working in concert with human expertise. We commit to leveraging AI technologies in ways that amplify and enhance human intelligence, creativity, and decision-making capabilities. Further, we commit to considering the psychological and social impacts of AI integration across university functions, recognizing that "beneficence" means actively working for the common good while avoiding harm.

4. Responsible & Ethical Use

We consider environmental impact, intellectual property rights, and social consequences before deploying AI.

We will promote understanding and discourse regarding the responsible and ethical use of AI. We will closely monitor and consider the implications and impact of the intersection of AI and intellectual property. When choosing to use AI systems, we will consider quality, sustainability, and environmental impact. Further, we will consider the findability, accessibility, interoperability, and optimization of reuse of AI digital assets.

5. Fairness & Transparency

We design, procure, and use AI systems that are explainable and strive to reduce bias. We disclose AI use to those affected by it.

Transparency regarding AI usage is essential to maintain public trust. We should promote attribution and transparency regarding data collection practices and AI use in decision-making. We will strive for continuous improvement in inclusivity and equality of opportunity, treatment, and impact, considering fairness in all aspects of AI use. We acknowledge that there are several potential sources of bias in AI and that our responsibility to educate ourselves about and respond to bias extends to our interactions with AI. We will prioritize explainable AI solutions that clearly outline how specific results are produced and why, allowing users to understand the reasoning behind the outputs, anticipate expected impacts, and identify potential biases.

6. Human Judgment & Accountability

People remain accountable for decisions influenced by AI. A human should always be in the loop for important decisions.

The university should preserve human judgment and accountability, with AI informing rather than replacing decision-making. Individuals using AI systems are responsible for adhering to existing university policies, standards, and security review processes. Virginia Tech intends to support an AI-informed workforce and student population and, where feasible, will provide AI-related upskilling, reskilling, and educational opportunities. This principle extends beyond decision-making to preserving human relationships and connections that are fundamental to our educational mission.

7. Data Security & Privacy

We safeguard personal and institutional data used in AI systems and interactions.

We will prioritize the safety, security, privacy, and protection of our community, valuing the unique benefits that human interactions bring to university processes. We recognize that people are ultimately responsible for decisions, especially in situations that directly impact others. Further, we commit to safeguarding individual privacy rights as well as proprietary university data by understanding what data is being used by AI solutions and who can access it. We should protect data from disclosure and, where appropriate, obtain consent from individuals who interact with AI.

Implementation and Oversight

These principles apply to all students, faculty, and staff, contractors, and affiliated partners who design, procure, or use AI at Virginia Tech. While the principles provide enduring guidance, their application requires practical judgment and will be supported by:

- Domain-specific practical implementation guides with concrete examples and use cases
- Risk assessment frameworks to guide decision-making
- Approved tool catalogs identifying validated AI services
- Regular training programs tailored to different roles and needs

The AI Working Committee, working within the IT governance framework, will oversee the continuous evolution of these principles and their implementation, ensuring they remain relevant and actionable as technology advances and our understanding deepens.

Living Principles

These principles are not static rules but living guidance that will evolve through practice and experience. They are intentionally written to balance specificity with flexibility, providing clear direction while avoiding prescriptive constraints that would quickly become outdated. Regular review cycles will ensure they continue to serve Virginia Tech's mission while adapting to emerging challenges and opportunities in the AI landscape.

Risk-Tier Assessment Framework

Virginia Tech employs a three-tier risk assessment framework that matches oversight intensity to potential impact:

- Low Risk Applications: Basic productivity tools, general writing assistance, simple data analysis. These require self-certification with standard guidelines and periodic audit.
- Medium Risk Applications: Advanced analytical tools, decision support systems, educational applications. These require departmental approval and regular monitoring.
- High Risk Applications: AI systems affecting personnel decisions, student outcomes, or resource allocation. These require review by the AI Working Committee, which will provide recommendations to executive leadership through the IT governance structure. Final approval authority rests with appropriate executive leaders.

Complete assessment criteria, processes, and implementation procedures are detailed in the supplementary document.

4.3 Governance Structure

The governance structure establishes the AI Working Committee as a standing working committee within the IT governance framework, reporting to the IT Governance Advisory Committee with delegated responsibility for AI-related issues across the university. This structure ensures specialized AI expertise and focus while benefiting from integration with university-wide technology governance. The AI Working Committee maintains authority over AI-specific policies, risk assessment, and implementation guidance while coordinating with IT governance on technology infrastructure, security, and enterprise-wide decisions. Complete structural details, membership specifications, and operational procedures are provided in the supplementary document.

4.4 AI Working Committee Composition and Authority

The AI Working Committee serves within the university's IT governance structure, bringing together academic, administrative, and technical perspectives.

As a standing working committee within the IT governance framework, the AI Working Committee provides advisory stewardship over shared practices, policies, and resources that enable responsible AI adoption. The committee will make recommendations to campus leadership through established governance channels, with operational implementation decisions remaining with appropriate campus leaders and units.

Complete membership specifications, authority frameworks, and operating procedures are detailed in the supplementary document.

Escalation and broader governance: Actions with institution-wide impact–such as structural changes, policies affecting multiple units, or major budget proposals–will continue to move through established governance bodies, ensuring alignment with the university's overall decision-making framework.

4.4.1 Decision Authority Framework (To Be Determined)

The specific decision rights and authority levels for the AI Working Committee are currently being developed as part of the broader IT governance reform. The following framework will be finalized through the IT governance process:

- **Advisory Functions**: The AI Working Committee will provide guidance and recommendations on AI policies, standards, and implementations
- **Operational Decisions**: Day-to-day AI implementation decisions will remain with campus operational leaders and units
- **High-Risk Decisions**: The AI Working Committee will advise executive leadership through the IT governance structure on high-risk AI implementations
- **Policy Development**: Specific authority for AI policy development and approval will be determined through the IT governance reform process

Note: These decision rights will be formalized upon establishment of the AI Working Committee within the reformed IT governance structure.

4.5 Integrated Subgroups

The AI Working Committee will establish domain-focused subgroups to ensure comprehensive and complementary guidance across general and domain-specific AI applications. These integrated subgroups will incorporate perspectives from Teaching & Learning, Research, and Operations. The specific structure, membership, and operating procedures for these subgroups will be determined by the AI Working Committee upon its formation, including:

- Formal charter development for each subgroup
- Membership appointment processes
- Reporting relationships and cadence
- Integration mechanisms with the main committee

Note: The term 'integrated' reflects the need for these subgroups to work collaboratively to ensure comprehensive coverage across all university functions.

The AI for Teaching and Learning subgroup develops practical faculty guidance through practical implementation guides, creates syllabus language templates and assessment adaptation strategies, and establishes student AI literacy curriculum alongside faculty and staff training programs.

The AI for Research subgroup develops research-specific guidance, creates intellectual property guidelines for AI-generated research outputs, and establishes publication standards and data handling protocols for AI-assisted research.

The AI for Operations subgroup focuses on administrative AI implementations through operations practical implementation guides, develops guidelines for automated decision systems with appropriate human oversight requirements, and establishes vendor evaluation procedures alongside staff training programs.

4.6 IT Governance Integration

The AI Working Committee operates as a formal standing working committee within the university's IT governance structure, parallel to other standing committees such as the Research Technology Working Committee and Teaching and Learning Working Committee, ensuring that AI governance benefits from enterprise-wide technology coordination while maintaining the focused expertise and authority needed for responsible AI implementation. This integration addresses the university's vision for unified technology governance while preserving the specialized functions identified as critical by the Virginia Tech AI working group.

The AI Working Committee structure integrates systematically with existing IT governance frameworks through coordinated decision-making processes, shared membership, and aligned technology standards. Detailed integration mechanisms and procedures are specified in the supplementary document.

4.7 Resources and Support

Implementation support ensures that ethical AI framework adoption receives appropriate institutional resources and expertise across all operational levels. This support infrastructure addresses the practical challenges of translating ethical principles into effective operational practices while helping community members at all levels of AI expertise.

Support programs span the complete spectrum from basic AI literacy for general university community members through advanced implementation guidance (see supplementary document for curriculum concepts) for technical specialists and administrative leaders. Technical support services include help desk services for routine questions, implementation consulting for complex deployments, and security review processes that ensure compliance with institutional standards. The policy library maintains current resources including general institutional policies, domain-specific implementation guidelines, and practical examples that demonstrate successful ethical AI implementation across diverse university functions. Assessment tools provide resources including risk evaluation forms, compliance checklists, and ongoing monitoring frameworks that enable continuous improvement and accountability.

4.8 Domain-Specific Applications

The ethical framework addresses three primary operational domains that encompass the full spectrum of university activities, ensuring that AI implementation maintains consistency with institutional values while accommodating the distinct requirements and challenges of different functional areas.

Teaching and learning applications require specialized attention to faculty and staff training programs (see Section 6.2 for Teaching and Learning practical implementation guide specifications) that develop pedagogical expertise in AI integration, student AI literacy initiatives that prepare learners for AI-enhanced educational environments, assessment adaptation that maintains academic rigor while accommodating AI tools, and academic integrity frameworks that address the evolving challenges of AI in educational contexts.

Research applications demand attention to data management protocols that ensure research integrity while enabling AI-enhanced inquiry, intellectual property frameworks that protect institutional and researcher interests, publication standards that maintain scholarly credibility in AI-assisted research, and ethical review processes that address the unique considerations of AI research methodologies.

Administrative operations require transparent automated decision systems that maintain accountability and fairness, staff training programs that enable effective AI utilization, robust privacy protection mechanisms that safeguard community and institutional data, and efficiency measurement protocols that demonstrate AI value while maintaining service quality.

4.9 Implementation Strategy

This implementation strategy directly responds to the readiness assessment findings (Section 3.3), addressing the need to translate Virginia Tech's research excellence into operational AI capabilities through permanent governance structures and sustainable systems.

The strategy prioritizes establishing robust governance foundations before expanding operational implementation across the institution. This sequenced approach ensures that AI initiatives build upon stable structures while maintaining flexibility as technology and institutional needs evolve.

See Section 12: Implementation Timeline and Roadmap for detailed schedule, deliverables, and responsible parties.

5. Policy Gap Analysis and Recommendations

5.1 Assessment Overview

An evaluation of 168 policies across nine organizational divisions provided critical insights into institutional readiness for AI integration while identifying specific gaps requiring attention (see Appendix B: Policy Gap Analysis - Methodology Details, Section B.1). Subject matter experts applied rigorous consistent criteria that evaluated AI integration potential, data governance requirements, ethical use standards, and operational oversight mechanisms to ensure thorough assessment of institutional policy infrastructure.

The evaluation findings demonstrate Virginia Tech's strong foundational policy framework, with 145 policies (86%) requiring no immediate action, indicating substantial alignment between existing institutional governance and AI implementation requirements. However, the analysis identified 18 policies (11%) requiring supplemental guidance to address AI-specific considerations and 5 policies (3%) demanding formal revision to ensure coverage of critical AI governance areas. This distribution suggests that while Virginia Tech's existing policy infrastructure provides a solid foundation, focused attention on specific high-impact areas will enable AI governance integration.

These findings directly informed the governance recommendations (Section 4.3) and implementation roadmap (Section 12).

5.2 Critical Findings

The policy analysis revealed four priority areas requiring immediate attention to ensure AI governance across all institutional functions. These areas represent critical intersections between AI capabilities and existing university operations where policy gaps could significantly impact institutional effectiveness and compliance.

Academic integrity emerges as the most immediate priority, with the Honor Code requiring review and guidance to address AI-assisted academic work standards (see Appendix B: Academic Affairs findings, Section B.2.2), including clear definition of appropriate usage, attribution requirements, and adjudication procedures. Research integrity policies demand focused attention to intellectual property frameworks and research misconduct definitions that address AI-specific considerations including ownership of AI-generated works, collaboration agreements, and disclosure requirements for AI-assisted research methodologies. The need for clear AI usage standards aligns with faculty concerns identified in the readiness assessment (Section 3.2, Faculty Engagement findings).

Human resources policies present substantial opportunities for AI integration enhancement, particularly in recruitment processes, performance evaluation systems, and staff development programs that can leverage AI capabilities while maintaining fairness and transparency. IT governance frameworks require integration of AI Responsible Use guidelines into existing acceptable use policies, ensuring that AI applications align with institutional security, privacy, and operational standards while enabling innovation and efficiency improvements.

5.3 Policies Requiring Formal Revision

Five institutional policies, representing 3% of those reviewed, require immediate formal revision to address critical AI governance gaps that could significantly impact university operations and compliance. These policies span multiple domains and require coordinated revision efforts that balance innovation enablement with appropriate oversight and accountability.

The Undergraduate Honor Code (Policy 6000) would benefit from review to establish clear standards for AI-assisted academic work (see Appendix B: Academic Affairs - Detailed Policy Assessment, Section B.2.2), including precise definitions of appropriate usage, attribution requirements that maintain academic integrity, and detection and adjudication procedures for AI-related misconduct that ensure fair and consistent enforcement.

The Policy on Intellectual Property (Policy 13000) would benefit from clarification of ownership frameworks for AI-generated works (see Appendix B: Research Domain findings, Section B.2.3) and collaboration agreements that address the complex intersection of human creativity and artificial intelligence capabilities. This revision must systematically address patent and copyright implications for AI-assisted research while protecting both institutional and individual researcher interests.

The Policy on Misconduct in Research (Policy 13020) would benefit from establishment of supporting citation standards for large language model usage (see Appendix B: Research - Detailed Policy Assessment, Section B.2.3) and disclosure requirements that maintain research integrity and transparency. The standards and procedures must define clear parameters for AI

data fabrication and falsification that address emerging challenges in AI-assisted research methodologies.

The Visual Media Policy (Policy 8205) must be reviewed to consider the growing challenges of AI-generated imagery and deepfake technology (see Appendix B: Student Affairs findings, Section B.2.4), and associated consent implications that affect student safety and institutional reputation. The policy would benefit from clear attribution requirements for AI-created visual content that maintain transparency and accountability.

Finally, Safety and Security Policies (e.g. Policiy 5617) would benefit from guidelines for AI-enhanced surveillance systems (see Appendix B: Safety and Security considerations, Section B.2.8) that balance security effectiveness with privacy protection and community trust. These revisions must address oversight considerations for AI-powered security applications while maintaining institutional commitment to community safety and individual privacy rights.

5.4 Implementation Strategy

The policy revision implementation prioritizes critical governance needs while building sustainable systems for ongoing AI policy management. This approach addresses immediate risks through formal revision of five identified policies while establishing frameworks for long-term governance.

Policy revision follows the institutional implementation timeline established in Section 12, with critical policies addressed in the initial phase and comprehensive review extending through maturation phase. The strategy coordinates with AI Working Committee establishment to ensure policy development aligns with governance structure creation and operational capacity building.

See Section 12: Implementation Timeline and Roadmap for specific timing and milestones.

6. Areas Needing Additional Guidance and Resources

6.1 Resource Gap Analysis

Analysis combining working group research and user feedback from the ChatGPT Edu pilot (see Appendix C.4: Qualitative Findings) and faculty survey responses (see supplementary document for current AI usage data) identified critical areas requiring enhanced resources to facilitate responsible AI adoption across university functions. These gaps will be systematically addressed through specialized, domain-specific “practical implementation guides” providing practical implementation guidance.

6.2 Three Priority Implementation Guides

Three domain-specific practical implementation guides will provide practical implementation guidance that translates institutional AI principles (see Section 4.2: Core AI Principles) into actionable operational frameworks. These practical implementation guides address the distinct challenges and opportunities within teaching, research, and administrative functions while maintaining consistency with overall institutional AI governance.

The Teaching and Learning practical implementation guide provides guidance for academic AI integration through syllabus language templates that establish clear expectations for AI use disclosure while maintaining academic integrity standards. The practical implementation guide includes assessment adaptation strategies that account for AI tool availability while preserving educational rigor, academic integrity frameworks that address the evolving challenges of AI-assisted academic work (responding to concerns identified in Appendix B.2.2: Academic Affairs policy gaps, particularly regarding the Undergraduate Honor Code), and specialized faculty training modules that develop pedagogical expertise in AI integration across diverse curriculum design contexts.

The Research practical implementation guide addresses the complex intersection of AI capabilities and scholarly inquiry through detailed IRB alignment procedures that ensure AI-assisted research protocols meet institutional ethical standards. he practical implementation guide provides data handling decision trees that guide appropriate AI tool usage in research contexts, establishes publication standards and attribution requirements for AI-generated content that maintain scholarly integrity (addressing gaps identified in Appendix B.2.3: Research policies, specifically Policies 13000 and 13020), and develops intellectual property guidelines that protect both institutional and researcher interests in AI-enhanced research outputs.

The Administrative Operations practical implementation guide The practical implementation guide focuses on responsible AI integration in university operations through guidelines for automated decision systems that maintain appropriate human oversight requirements and institutional accountability (building on the risk-tier assessment framework detailed in Section 4.2 and supplementary document). The practical implementation guide includes staff training protocols that enable effective AI tool integration in administrative workflows, detailed vendor evaluation frameworks that ensure AI service procurement aligns with institutional standards and values, and transparency standards that maintain community trust in AI-assisted administrative processes.

6.3 AI Literacy and Capability Development

A multi-faceted approach to cultivating AI literacy across the Virginia Tech community ensures that all constituencies develop appropriate competencies for effective and responsible AI engagement. This responds directly to readiness assessment findings (see Appendix A.2.4: AI Literacy and Ethical Use - Developing level) indicating that while educational content exists, implementation gaps remain. This capability development framework addresses diverse learning needs while building institutional capacity for sustained AI integration and innovation.

Virginia Tech's University Libraries have established a strong foundation for AI literacy through existing initiatives and resources. The proposed expanded literacy programs will build upon these established efforts, leveraging the Libraries' expertise and existing infrastructure to achieve university-wide impact.

The literacy development strategy includes foundational AI literacy modules designed for universal deployment across all students, faculty, and staff, addressing the specific barrier identified in the faculty survey where 38% of non-users cited 'uncertainty about how to use AI' (see supplementary document for barriers analysis), ensuring baseline understanding of AI capabilities, limitations, and ethical considerations. Building upon this foundation, role-specific

training programs will address the unique needs and applications associated with different functional areas, providing targeted expertise that enables effective AI integration within specific operational contexts. Advanced implementation guidance supports complex AI integration scenarios that require a sophisticated understanding of technical capabilities, governance requirements, and ethical implications. Continuous professional development programs ensure ongoing capability enhancement that keeps pace with rapidly evolving AI technologies and institutional needs, creating sustainable systems for long-term AI literacy and capacity advancement.

6.4 Implementation Strategy

The implementation strategy prioritizes immediate educational needs while building capability development infrastructure for long-term institutional AI integration. Resource development aligns with the timeline in Section 12, with the Teaching and Learning practical implementation guide addressing urgent faculty needs first, followed by Research and Administrative Operations practical implementation guides.

The development sequence reflects critical dependencies: practical implementation guide creation follows risk-tier assessment framework establishment, and supplemental policy guidance depends on working group formation. Training program deployment coincides with practical implementation guide completion to maximize effectiveness.

Continuous improvement processes embedded from the outset ensure guidance materials remain relevant through regular user feedback and iterative updates.

See Section 12: Implementation Timeline and Roadmap for specific development milestones and scheduling.

7. AI Tools and Platform Recommendations

7.1 Evaluation Framework

AI tool recommendations emerge from an evaluation framework that balances multiple critical factors to ensure institutional AI deployments align with university values, operational requirements, and objectives. This assessment methodology considers security requirements that protect institutional data and community privacy, educational value that advances teaching and learning objectives, administrative efficiency that enhances operational effectiveness, cost-effectiveness that optimizes resource utilization, and integration capabilities that leverage existing university technology infrastructure while minimizing operational disruption.

This framework has been validated through pilot program evaluation, including the ChatGPT Edu pilot involving 425 participants over a four-month period (see Appendix C.1: Pilot Overview and Methodology, which details the January-May 2025 implementation). The pilot provided evidence of tool effectiveness, user adoption patterns, and integration requirements that directly inform institutional AI tool selection criteria.

The university also partnered with an external vendor, Cloudforce, to conduct a January-to-March 2025 Microsoft 365 Copilot pilot. A cohort of 177 faculty and staff worked with Copilot for 38 workdays, supported by six instructor-led training sessions and a dedicated Teams channel for peer feedback. Post-pilot analytics showed that 94% of participants reported a daily time saving, averaging 38 minutes per workday, and that overall satisfaction reached 80% with 92% willing to recommend the tool (see news release https://news.vt.edu/notices/2025/05/it-copilot-study-report-next-steps.html) . These results confirm that a balance of security, educational value, administrative efficiency, cost, and integration requirements should inform implementation across multiple tool types.

Finally, Rolai, an AI-tutoring add-on now under evaluation in selected Canvas courses through Fall 2025, will be folded into the same assessment pipeline once outcome data are available.

7.2 Initial Approved Tools and Risk Categories

The university's initial AI tool deployment strategy focuses on three carefully selected pilot programs that demonstrate successful integration with existing infrastructure while providing measurable value across diverse operational contexts.

Microsoft Copilot Suite

Selected because it takes advantage of existing M365 infrastructure investments to provide productivity enhancement capabilities that integrate seamlessly with established workflows and security protocols. Microsoft Copilot Suite was given approval for campus use after demonstrating secure integration with the university's existing M365 tenant and documented productivity gains–especially in Outlook and Word, where satisfaction exceeded 80% (see Section 7.1 for pilot study results). While the suite meets all baseline security controls, limitations with Excel analytics and feature consistency require continued monitoring; if future releases add automated decision functions, an escalation to Tier 3 may be necessary.

ChatGPT Edu

Validated through extensive pilot evaluation (detailed results in Appendix C). Key applications validated include research support, writing assistance, and administrative workflows. This platform has not been officially selected as an enterprise platform for Virginia Tech, but pilot experiences suggest it may be a top contender.

Advanced Research Computing

AI Resources provide specialized GPU cluster access that supports advanced AI research and development initiatives, enabling advanced scholarly inquiry while maintaining appropriate security and oversight protocols.

These pilot programs operate within a risk categorization framework that guides evaluation and oversight intensity.

- Low-risk applications encompass basic productivity tools that operate with standard security requirements and minimal operational impact, enabling streamlined deployment and user adoption.

- Medium-risk applications include educational and research implementations that require enhanced oversight due to their potential impact on academic integrity, intellectual property, and research methodology.
- High-risk applications involve administrative decision systems that require assessment and ongoing monitoring due to their potential impact on institutional operations, community welfare, privacy, and regulatory compliance.

7.3 Implementation Strategy

The implementation strategy builds upon successful pilot program experiences while establishing sustainable systems for ongoing AI tool evaluation and deployment.

Initial effort will focus on expanding current pilot programs through integration of user feedback and security assessment results, ensuring that initial deployments provide validated models for broader institutional adoption.

A second stage enables deployment of approved tools through university-managed licensing and support structures that provide consistent user experience while maintaining institutional control over costs, security, and compliance requirements.

Finally, establishing ongoing evaluation processes for emerging AI technologies and tools, creating mechanisms for continuous assessment of new capabilities that align with evolving institutional needs and technological advancement will finalize the implementation efforts.

7.4 Cost and Support Framework

The bottom line is that high impact use requires secure access to highly capable tools, and secure access to highly capable tools is relatively expensive. A fundamental paradox is that the level of investment required to provide broad access to highly capable AI interfaces would logically require evidence of high impact use cases, but the development of high impact use cases cannot occur without safe and reliable access to highly capable tools.

To continue exploration of piloted tools, we recommend the following.

Recommendation and Anticipated Costs

ChatGPT Edu

Investment recommendation: $120,000 annually for 1,000 licenses plus variable credit costs (detailed cost analysis in Appendix C for cost analysis). Success requirements include improved integration, enhanced onboarding, and accessibility improvements.

M365 Copilot

Copilot follows Microsoft's enterprise licensing. Units that opt in pay the per-seat fee published on the Software Service Center site and can expect prorated charges for the remainder of FY 25 (https://software.vt.edu). Because Copilot runs entirely within the university tenant, there is no additional storage or data-egress cost, but effective use still hinges on robust support. Pilot participants highlighted the value of "Art of the Possible" workshops and weekly clinics; CCS will

therefore keep those offerings and add short "feature-focus" videos whenever Microsoft ships major updates. Accessibility testing is ongoing; until Excel analytics ships with full keyboard and screen-reader support, units relying heavily on data analysis should weigh Copilot's benefits against that gap.

7.5 Pilot Program Lessons Learned

The ChatGPT Edu pilot provided critical insights for institutional AI deployment strategy, highlighting needs for improved integration, training, and accessibility (detailed findings in Appendix C).

Copilot results reinforced several lessons from the ChatGPT Edu pilot (see news release at https://news.vt.edu/notices/2025/05/it-copilot-study-report-next-steps.html for complete pilot results). Participants praised its ability to surface emails, recap meetings, and draft content while keeping data inside the secure M365 tenant. At the same time, two limitations stood out: the absence of advanced Excel analytics and variations in feature depth across M365 apps, both of which tempered initial expectations. These findings confirm that usability and feature parity, not just security, drive adoption. They also underscore the need for clear communication about product roadmaps and for tiered rollout plans that let high-dependency units wait until critical functions mature.

Implementation Strategy insights include the value of conservative security approaches during initial deployment, the importance of onboarding processes that account for delayed participation, and the need for feedback collection mechanisms that inform ongoing policy and resource development.

See Appendix C for detailed tool evaluation matrices, security assessment criteria, and implementation timelines.

8. Universal Access Considerations

8.1 Addressing the Digital Divide

Virginia Tech's AI integration strategy prioritizes ensuring that all advancements bridge rather than exacerbate existing digital divides within the university community. The institution is committed to proactively addressing potential barriers including cost constraints, varying technology requirements, geographic limitations, and diverse technical expertise levels among students, faculty, and staff.

These considerations have been validated through pilot program experience (see Appendix C.5: Accessibility and Equity Assessment for detailed findings). The pilot demonstrated the critical importance of proactive accessibility assessment, leading to enhanced evaluation criteria that prioritize universal access in AI tool selection and deployment processes.

8.2 Accessibility Compliance Requirements

Virginia Tech maintains steadfast commitment to full compliance with accessibility mandates, particularly ADA Title II Subpart H requiring adherence to WCAG 2.1 AA standards by Spring 2026.

This requirement includes all digital content including web applications, mobile platforms, digital communications, videos, documents, LMS systems, and social media.

AI tool evaluation criteria explicitly incorporate rigorous accessibility standards with testing protocols ensuring compatibility with assistive technologies (as demonstrated in the ChatGPT Edu pilot accessibility findings detailed in Appendix C.5). The university will implement alternative access methods for core functions where necessary, supported by centralized policy governance, resource provision, and compliance auditing processes.

8.3 Algorithmic Fairness and Bias Mitigation

Virginia Tech implements bias mitigation strategies that ensure AI systems maintain accountability for fairness and unbiased outcomes across all applications, regardless of whether underlying algorithms are proprietary or institutionally developed. This approach addresses the challenges of bias detection and correction in AI systems through multiple interconnected mechanisms.

The university conducts regular algorithmic auditing processes that systematically evaluate AI systems for bias detection across diverse demographic groups (building on the risk-tier assessment framework outlined in Section 4.2 and detailed assessment procedures in supplementary document), ensuring that technological solutions do not inadvertently perpetuate or amplify existing structural challenges. These auditing processes should ensure that multiple perspectives inform critical decisions about AI system design, implementation, and ongoing assessment. Outcome monitoring with corrective action protocols provides ongoing oversight that enables rapid identification and resolution of bias-related issues as they emerge. Transparency requirements for AI decision-making processes, particularly in high-impact applications affecting student success, employment decisions, or resource allocation, will ensure that community members understand how AI systems influence decisions that affect their university experience.

8.4 Data Sovereignty and Cultural Considerations

Respect for data sovereignty and cultural considerations is integral to the ethical AI framework. Data sovereignty includes both legal compliance with regional regulations and community rights to control data. All AI applications must adhere to ethical data use protocols and applicable data transfer compliance requirements.

8.5 Implementation Strategy

The universal access implementation strategy ensures that design principles that seek to reduce divides are embedded throughout all AI initiatives rather than addressed as secondary considerations. This approach integrates access requirements into evaluation, deployment, and monitoring processes that maintain institutional commitment to excellence.

Immediate actions focus on integrating accessibility criteria as primary factors in all AI tool evaluations, addressing the specific gaps identified in the ChatGPT Edu pilot (see Appendix C.5 for detailed findings) and aligning with the three-tier risk assessment process (Section 4.2 and supplementary document), ensuring that accessibility compliance is assessed before rather than

after deployment decisions. Accessibility review of draft implementation practical implementation guides (described in Section 6.2: Three Priority practical implementation guides) ensures that practical guidance materials meet broader user needs and capabilities. Support programs for users with varying technical capabilities will provide assistance that enables effective AI tool utilization across the entire university community regardless of technical background or experience.

Ongoing monitoring includes regular access audits overseen by the AI Working Group (governance structure detailed in Section 4.4 and supplementary document) that systematically monitor demographic impact and ensure that AI implementations advance rather than hinder institutional objectives. Accessibility testing for AI tools provides proactive compliance assessment that identifies and addresses accessibility challenges before they affect community members. Continuous improvement processes ensure ongoing compliance with evolving accessibility standards while maintaining institutional leadership in technology deployment practices.

9. Stakeholder Engagement and Communication

9.1 Engagement Strategy

Stakeholder engagement ensures AI initiatives align with university community needs and values, as validated through the ChatGPT Edu pilot's 425 participants (see Appendix C.2: Participant Demographics and Representation) and extensive Working Group consultations. The strategy includes feedback collection, transparent communication, and ongoing collaboration across all university constituencies.

9.2 Key Stakeholder Groups

Effective AI governance requires engagement with diverse stakeholder communities whose perspectives and expertise inform institutional decision-making while ensuring that AI initiatives serve the full spectrum of university constituencies. This stakeholder-informed approach includes both internal university community members and external partners whose collaboration enhances institutional AI capabilities, building on the engagement patterns identified during pilot programs (see Appendix C.2: diverse representation across colleges) and policy review processes (see Appendix B.4: Stakeholder Engagement Framework).

Internal stakeholders represent the core university community whose daily experience with AI systems directly influences institutional success. This constituency includes faculty, staff, and students across all colleges and departments who bring diverse disciplinary perspectives and operational requirements to AI governance discussions. University leadership, governance bodies, and administrative units provide oversight and policy expertise that ensures AI initiatives align with institutional mission and regulatory requirements. The research community, teaching professionals, and support services contribute specialized expertise in AI applications, educational integration, and operational implementation that informs practical governance decisions.

External stakeholders provide essential perspectives that connect institutional AI initiatives with broader technological, regulatory, and societal contexts. Industry partners, peer institutions, and regulatory bodies offer expertise in AI implementation best practices, emerging technological capabilities, and compliance requirements that inform strategic planning and risk management. Community organizations, alumni networks, and government agencies contribute insights into public expectations, workforce development needs, and regional economic impacts that ensure AI initiatives serve broader institutional mission objectives.

9.3 Stakeholder Communication

A stakeholder communication strategy establishes channels for information sharing, feedback collection, and collaborative decision-making that maintain transparency while enabling responsive governance. This approach ensures that all stakeholders receive appropriate information while contributing meaningfully to ongoing AI governance development, addressing the communication gaps identified in the readiness assessment (see Appendix A.2.2: Institutional Governance - 'Limited formal collaboration between units').

Regular communication channels provide information sharing through university-wide updates on AI initiatives (responding to pilot participant feedback requesting 'more practical onboarding resources' and 'advanced implementation guidance' – see Appendix C for training and support requirements), policy changes, and implementation progress that keep the entire community informed of institutional AI development. Domain-specific guidance delivers targeted communications to relevant constituencies that address operational needs and disciplinary requirements. Interactive forums facilitate ongoing feedback collection and community dialogue that enables continuous engagement rather than episodic consultation.

The feedback integration process ensures that stakeholder input systematically influences AI governance through collection and analysis of stakeholder input on AI policies and implementations. Regular adjustment of initiatives based on community feedback and emerging needs demonstrates institutional responsiveness to stakeholder concerns and evolving requirements. Transparent reporting on how feedback influences AI governance decisions maintains accountability and demonstrates institutional commitment to collaborative governance processes.

9.4 Ongoing Engagement Strategy

Continuous stakeholder engagement through quarterly forums, annual assessment surveys, and meaningful integrated task team participation, building on the high engagement rates demonstrated in pilots (78% weekly active users, see Appendix C for engagement metrics) and faculty survey responses (118 participants representing diverse disciplines, see supplementary document for details) ensures that AI governance remains responsive to evolving community needs and rapidly changing technological developments. This sustained engagement approach creates opportunities for stakeholder input while building institutional capacity for adaptive governance that maintains effectiveness in dynamic technological and regulatory environments.

10. Innovation and Continuous Improvement

Virginia Tech's approach to Artificial Intelligence extends beyond mere adoption to actively fostering a culture of innovation and continuous improvement. This section outlines how the university will cultivate an environment where AI's potential is consistently explored, measured, and refined, ensuring sustained relevance and leadership in the evolving AI landscape.

10.1 Innovation Support

To catalyze AI-driven innovation, Virginia Tech will establish a robust support approach designed to encourage experimentation and the scalable development of new AI applications. This approach will encompass various mechanisms, including the provision of secure sandbox environments referenced within forthcoming practical implementation guides (see Section 6.2 for domain-specific practical implementation guide specifications), allowing faculty, researchers, and staff to safely explore AI tools and concepts. Furthermore, the university will leverage existing funding mechanisms and explore new avenues to support promising AI initiatives, facilitating their progression from conceptualization to impactful implementation. Emphasis will also be placed on identifying and promoting pathways for scaling successful pilot projects and fostering cross-disciplinary collaboration, recognizing that many of AI' s most profound applications emerge at the intersection of diverse fields.

10.2 Assessment and Metrics

Effective integration of AI necessitates a disciplined approach to assessment and measurement. Virginia Tech will develop metrics to strategically monitor the impact and progress of AI adoption across the institution. These metrics will provide high-level insights into various facets of AI integration, including the proportion of new projects undergoing risk assessments, the growth in the catalog of approved AI tools, the rates of participation in AI literacy training programs, and overall user satisfaction. Furthermore, the university should track the outcomes of innovation projects, evaluate cost savings and efficiency gains attributable to AI, and monitor research productivity enhancements. This data-driven approach will enable informed decision-making and demonstrate the tangible value of AI investments.

These metrics have been validated through pilot program evaluation (detailed in Appendix C). The pilot provided a measured approach for measuring AI integration success that can inform ongoing assessment and continuous improvement processes.

10.3 Continuous Improvement Processes

Recognizing the rapid pace of AI evolution, Virginia Tech is committed to embedding continuous improvement into its AI governance and implementation frameworks. This involves establishing regular review cycles for the foundational guidance and associated standards, ensuring their ongoing relevance. practical implementation guides providing domain-specific guidance will undergo quarterly updates to reflect new insights, technological advancements, and community feedback. Similarly, the ai.vt.edu web site will be refreshed regularly, serving as a dynamic repository of knowledge. Ongoing evaluation of AI tools and consistent benchmarking against peer institutions will further inform adaptive strategies, ensuring that Virginia Tech remains at the forefront of responsible AI integration in higher education.

10.3.1 Addressing Critical Concerns and Challenges

Virginia Tech's commitment to continuous improvement requires honest acknowledgment of substantive concerns raised by faculty, particularly those in writing-intensive and other creative disciplines. The Working Group recognizes that AI adoption presents challenges extending beyond technical implementation to fundamental questions about academic integrity, scholarly values, and educational mission (see stakeholder concerns documented in Section 9 and policy gaps identified in Appendix B.2.2).

Faculty have identified critical gaps in current guidance regarding AI use in dissertations, theses, and examinations (see Appendix B.2.2: Academic Affairs policy analysis) – contexts where original thought and independent scholarship are paramount. The humanities and social sciences face particular challenges as these disciplines cultivate critical thinking through the act of writing itself. Additionally, the unreliability of AI detection tools creates enforcement dilemmas, while the opacity of large language models makes meaningful validation of AI-assisted research extremely difficult, especially in qualitative methodologies where interpretive nuance is essential.

The framework's emphasis on individual responsibility risks repeating past mistakes with technology adoption, where institutions failed to actively shape usage patterns and address systemic impacts. Faculty rightly note the environmental costs of AI systems, and the concerning opacity around how submitted data may be used beyond immediate academic purposes. These concerns echo across disciplines and deserve sustained institutional attention rather than wholesale delegation to individual users.

To address these challenges systematically, the AI Working Group should ensure strong representation from across the campus and explicitly incorporate these critical perspectives into its regular review cycles. The AI Working Group's annual assessment should evaluate not just technical effectiveness but also impacts on academic integrity, pedagogical values, and environmental sustainability. The integrated working groups, particularly those addressing Teaching & Learning and Research domains, should prioritize developing nuanced guidance for high-stakes academic assessments and discipline-specific research standards. By acknowledging these tensions and trade-offs, Virginia Tech demonstrates the thoughtful, critical engagement with technology that we seek to cultivate in our students, positioning Virginia Tech to lead through reasoned analysis rather than uncritical adoption.

10.4 External Partnerships and Leadership

Virginia Tech is positioned to be a thought leader in the responsible integration of AI within higher education, a role that necessitates active engagement beyond its immediate campus. The university should foster external partnerships with industry leaders, government agencies, and peer institutions to share best practices, collaborate on research, and collectively navigate emerging challenges. This includes developing a formal commitment to industry collaboration and strengthening networks with peer institutions to exchange insights on AI policy and implementation. The university should actively pursue grant and funding opportunities to advance its AI initiatives and seek opportunities to influence policy discussions at regional, national, and international levels, reinforcing its commitment to the land-grant mission.

10.5 Next Steps and Campus Recommendations

To drive this commitment to innovation and continuous improvement, several actions are recommended. The university should initiate the first review cycle of its AI efforts, leveraging the established metrics (Section 10.2) and governance framework (supplementary document for Annual Review Process) to assess progress and identify areas for refinement. Discussions and planning will commence to consider an annual AI symposium to showcase internal innovations and engage with external experts. Furthermore, guidelines for external partnerships will be developed to streamline collaborations. These steps will ensure that Virginia Tech' s AI journey is characterized by ongoing learning, adaptation, and a proactive pursuit of excellence.

11. Conclusion and Call to Action

11.1 Summary of Integrated Approach

This report provides an integrated approach for Virginia Tech's responsible engagement with Artificial Intelligence. It moves beyond theoretical discussions to offer both analysis and concrete implementation tools. The framework is built upon the seven principles (detailed in Section 4.2, Core AI Principles), supported by evolving standards, and translated into practical application through domain-specific practical implementation guides (specifications in Section 6.2, Three Priority practical implementation guides). A robust governance structure, centered around the AI Working Group, ensures oversight and accountability. This holistic approach directly addresses stakeholder feedback, delivering actionable guidance and a clear roadmap for Virginia Tech to navigate the complexities and capitalize on the opportunities presented by AI.

11.2 Critical Leadership Decisions Required

Successful implementation of Virginia Tech's AI framework requires five critical leadership decisions that establish institutional commitment and provide necessary authority for advancement. University leadership must approve the charge and documentation that articulates institutional principles (see Section 4.2, Core AI Principles and Section 4.3, Governance Structure) and announces the approach to governance. The appointment of AI Working Committee members through the IT governance process (proposed structure in Section 4.4, AI Working Committee Composition and Authority; detailed specifications in supplementary document) represents a fundamental governance decision that establishes expertise and accountability for institutional AI advancement. The AI Working Committee will serve in an advisory capacity, providing guidance and recommendations while operational decisions remain with campus leaders.

Resource allocation decisions for implementation (see Section 7.4, Cost and Support Framework for pilot-validated investments) determine the scope and pace of AI integration across university functions while demonstrating institutional commitment to responsible AI adoption. Leadership commitment to an annual review cycle ensures that AI governance remains responsive to technological advancement and institutional needs. Finally, leadership support for the cultural shift toward responsible AI adoption provides essential organizational backing for community engagement and change management throughout the implementation process.

11.3 Implementation Approach

Virginia Tech's AI implementation recognizes that sustainable transformation emerges through engagement and empowerment rather than mandate. The approach builds on existing strengths identified in the readiness assessment, particularly high faculty interest and established research excellence (detailed findings in Appendix A). This philosophy accommodates varying adoption paces across constituencies through flexible support mechanisms and multiple engagement pathways.

The implementation strategy emphasizes adaptive change management that acknowledges both the potential and inherent uncertainties of artificial intelligence. Different units may accelerate adoption based on readiness and resources, with the AI Working Group supporting both baseline and accelerated paths. Critical success factors include sustained executive commitment, adequate funding and staffing, active community participation, and adaptability to rapid technological change.

Central to success is the recognition that resistance often stems from legitimate concerns requiring thoughtful response. Faculty concerns about academic integrity, staff apprehensions about role changes, and student questions about fairness all deserve respectful engagement. The implementation approach addresses these through education, demonstration of benefits, and inclusive governance structures that ensure all voices are heard.

11.3.1 Transition from Current to Future State

The transition from the current AI Working Group to the permanent AI Working Committee will proceed as follows:

1. **Report Publication**: Upon acceptance of this report, the current AI Working Group completes its charge and is formally disbanded
2. **Committee Formation**: The IT Governance Executive Committee will charter the AI Working Committee as a standing working committee within 30 days of report acceptance
3. **Member Appointment**: Initial AI Working Committee members will be appointed according to IT governance procedures within 60 days
4. **Knowledge Transfer**: Key members of the current AI Working Group may be invited to serve on or advise the new AI Working Committee to ensure continuity
5. **Operational Handoff**: All recommendations, frameworks, and guidance developed by the current AI Working Group become the foundation for the AI Working Committee's ongoing work

11.4 Vision for Virginia Tech AI Leadership

Virginia Tech stands poised to model responsible AI adoption in higher education. By combining principled governance, practical implementation tools, and commitment to continuous improvement, the university can demonstrate how land-grant institutions serve their communities in the AI era. This leadership extends beyond technological deployment to

encompass ethical considerations, equitable access, and the preservation of human agency that defines educational excellence.

The path forward requires immediate action on the governance foundations outlined in this report, followed by sustained effort to build institutional capacity and community engagement. Through this commitment, Virginia Tech will not merely adapt to the AI revolution but will help shape how higher education harnesses these powerful technologies to advance knowledge, enhance learning, and serve the greater good.

The detailed implementation timeline and roadmap with specific milestones for achieving this vision are presented in Section 12.

12. Implementation Timeline and Roadmap

This implementation follows a phased approach with three distinct periods, subject to resource availability:

- **Initial Phase**: Establishing governance foundations
- **Expansion Phase**: Building operational capacity
- **Maturation Phase**: Achieving institutional transformation

Each period builds upon previous accomplishments, with specific milestones and deliverables ensuring measurable progress toward AI integration.

Implementation timelines align with Virginia Tech's budget planning cycles. As dedicated funding has not yet been allocated, initial activities focus on efforts achievable within existing resources while comprehensive budget proposals are developed for subsequent phases.

12.1 Initial Phase: Governance Foundation

The initial implementation phase establishes critical governance foundations and initiates capacity building across the university, incorporating accessibility compliance requirements aligned with Spring 2026 WCAG 2.1 AA standards deadline (see Section 8.2 for detailed compliance requirements and Appendix C.5 for accessibility assessment findings). Guidance approval and publication through the Provost's Office provides institutional authorization and community awareness of AI governance principles and commitments (detailed in Section 4.2, Core AI Principles). The formal seating of the AI Working Group creates operational governance capacity with defined authority and accountability for institutional AI advancement (see Section 4.4 for AI Working Group composition and supplementary document for complete membership specifications).

Responsible Parties for Initial Phase Deliverables:

The Provost's Office will lead the approval and publication of AI governance principles, working closely with the EVP/Chief Operating Officer to ensure institutional alignment. The Executive Leadership team (Provost and EVP/COO offices) will oversee AI Working Group formation and member appointment. The newly formed AI Working Group, with support from Legal Counsel and the Office of Audit, Risk, and Compliance, will publish the risk-tier assessment framework. Policy

owners will lead the revision processes for the five critical policies, with coordination support provided to the AI Working Group. TLOS (Technology enhanced Learning and Online Strategies), University Libraries and CETL (Center for Excellence in Teaching and Learning) will develop and deploy foundational AI literacy modules, while the IT Division manages digital infrastructure setup with Communications support.

Publication of the risk-tier assessment matrix as interim guidance through collaboration between the office of the EVP/Chief Operating Officer and Vice President for Research provides immediate practical frameworks for AI implementation decisions while policies are developed (risk-tier framework detailed in Section 4.2 and supplementary document). Initiation of domain-specific practical implementation guide development processes for Teaching and Learning, Research, and Operations begins translation of governance principles into practical operational guidance (practical implementation guide specifications outlined in Section 6.2 and supplementary document). The launch of foundational AI literacy modules creates baseline community competencies that support effective and responsible AI engagement across all university constituencies.

12.2 Expansion Phase

The capacity building phase focuses on operational development and stakeholder engagement that builds institutional AI implementation capabilities. Completion of three draft practical implementation guides with dedicated administrative support for the AI Working Committee provides detailed practical guidance for teaching and learning, research, and administrative AI applications (see Section 6.2 for practical implementation guide content overview and supplementary document for subgroup charters responsible for development). The open pilot of online risk-assessment forms enables evaluation processes that support consistent and efficient AI implementation decisions.

Responsible Parties for Expansion Phase Deliverables:

Cross-functional subgroups that incorporate perspectives from Teaching & Learning, Research, and Operations, each co-led by a AI Working Committee member and domain expert, will develop their respective practical implementation guides: the Teaching & Learning subgroup (Teaching & Learning practical implementation guide), the Research subgroup (Research practical implementation guide), and the Operations subgroup (Administrative Operations practical implementation guide). The Division of IT will manage tool deployment decisions with AI Working Committee feedback. Policy owners will develop supplemental guidance with support from relevant working groups. Human Resources will lead the comprehensive HR policy review with AI Working Committee guidance. TLOS, the University Libraries and CETL will deploy role-specific training programs based on completed practical implementation guides. The AI Working Committee, with Communications support, will coordinate stakeholder feedback sessions.

Publication of the initial approved AI tools catalog provides preliminary access to vetted technologies while demonstrating institutional commitment to evaluation and deployment processes (initial tools detailed in Section 7.2; evaluation framework in Section 7.1). The first stakeholder feedback sessions create opportunities for community input on governance effectiveness and implementation priorities, ensuring that AI initiatives remain responsive to institutional needs and values.

12.3 Maturation Phase

The assessment and expansion phase establishes evaluation and sustainability mechanisms that ensure long-term governance effectiveness and continuous improvement. The first round of unit-level self-audits conducted under AI Working Committee oversight provides assessment of AI implementation progress and identifies areas requiring additional support or policy development (assessment metrics framework outlined in Section 10.2).

Responsible Parties for Maturation Phase Deliverables:

The AI Working Committee, supported by the Office of Audit, Risk, and Compliance, will advise on a strategy for unit-level self-assessments. The AI Working Committee, working with the Division of IT and Institutional Research, will advise on the development and maintenance of analytics dashboards. The AI Working Committee will sponsor annual community listening sessions with support from stakeholder group representatives. Working groups will manage their respective annual review cycles. Individual initiative owners will lead scaling efforts under AI Working Committee oversight. The AI Working Committee, with Research and Communications divisions, will coordinate external engagement activities.

Development and presentation of analytics dashboards to appropriate campus leadership demonstrates measurable progress toward institutional AI objectives while maintaining governance accountability and oversight. Annual community listening sessions for guidance and standards updates ensure that governance frameworks remain responsive to technological advancement, regulatory changes, and evolving institutional needs (stakeholder engagement approach detailed in Section 9.3). evaluation of pilot program outcomes and scaling of successful initiatives creates evidence-based approaches to AI expansion while maintaining institutional commitment to responsible and effective implementation.

12.4 Implementation Dependencies and Coordination

Successful AI framework implementation requires attention to task dependencies and coordination mechanisms that ensure logical sequencing and optimal resource utilization. Guidance approval serves as the foundational enabler for all subsequent implementation actions, providing institutional authorization and governance framework that supports advancement (see Section 11.2 for critical leadership decisions required). The risk assessment matrix must be completed before practical implementation guide development to ensure that practical guidance incorporates appropriate risk evaluation and mitigation strategies (risk-tier categories defined in supplementary document).

Practical implementation guide completion should complement any AI tool deployment to ensure that community members have access to guidance and support resources before engaging with institutional AI systems. Continuous feedback collection through email, focus groups, and surveys will provide ongoing community engagement and appropriate responses to emerging questions and concerns throughout the implementation process.

12.5 Organizational Implementation Strategy

Effective implementation requires organizational coordination and accountability mechanisms that ensure consistent progress toward institutional AI objectives. Task ownership assignment

per established timelines creates clear accountability and enables progress monitoring across all implementation domains. Establishment of project management coordination through the DoIT Project Management Office provides professional implementation support and coordination across multiple organizational units.

Development of progress tracking dashboards enables real-time monitoring of implementation milestones while providing leadership with assessment capabilities (dashboard specifications and metrics outlined in Section 10.2). Regular quarterly leadership reviews ensure that implementation progress receives appropriate oversight while enabling adaptive management responses to emerging challenges or opportunities.

13. Appendices

The following appendices provide supporting documentation and detailed implementation guidance for Virginia Tech's AI framework:

- **Appendix A: Complete AI Readiness Assessment** - Detailed Virginia Tech AI Readiness Assessment using the Digital Education Council (DEC) Ten Dimension framework, including methodology and results by dimension
- **Appendix B: Policy Gap Analysis** - Complete findings from review of 168 Virginia Tech policies across nine organizational divisions, including detailed methodology and assessment framework
- **Appendix C: ChatGPT Edu Pilot Results** - Complete pilot program evaluation including methodology, participant demographics, quantitative results, qualitative findings, accessibility assessment, and strategic recommendations for institutional deployment

These appendices preserve the detailed analysis, methodology, and implementation guidance that supports the recommendations presented in the main report while enabling the report to maintain appropriate length and executive focus.

Note: Additional supporting materials including the Virginia Tech AI Landscape and Higher Education Context, Detailed Ethical AI Framework Implementation, AI Governance Implementation, Global AI Faculty Survey Results, and University AI Working Group Membership are available in the supplementary document "Implementation Guides and Background Research."

Appendix A: Preliminary AI Readiness Assessment

This appendix provides the Virginia Tech AI Readiness Assessment using the Digital Education Council (DEC) Ten Dimension AI Readiness Assessment framework.

A.1 Assessment Methodology

This assessment leverages the Digital Education Council (DEC) Ten Dimension AI Readiness Assessment, a structured evaluation tool developed by the DEC's Thematic Working Group on AI and Education. Virginia Tech was a contributor to the development of this assessment, with Dale Pike from Virginia Tech serving as a delegate on the DEC 2025 Artificial Intelligence Working Group. The framework is organized around ten interconnected dimensions, each with four levels of readiness progressing from "Emerging" to "Mature," acknowledging that institutions may

advance at different rates across various dimensions while emphasizing the interconnectedness of progress.

Academic Departments

Virginia Tech's AI excellence extends across multiple academic departments, with Computer Science and Electrical and Computer Engineering serving as primary anchors while fostering extensive interdisciplinary collaboration.

The **Department of Computer Science** houses core AI expertise through faculty like Naren Ramakrishnan (Sanghani Center Director) and partnerships including the Amazon-Virginia Tech Initiative for Efficient and Robust Machine Learning. The department emphasizes both technical advancement and societal impact, with dedicated leadership in AI ethics and responsible development.

The **Bradley Department of Electrical and Computer Engineering** contributes unique strengths in AI-communications convergence and hardware-software integration. ECE's Machine Learning Major within Computer Engineering represents one of the nation's few specialized undergraduate AI programs. Faculty like Walid Saad pioneer AI integration in 6G wireless systems and edge computing applications.

Beyond these core departments, AI research thrives across Virginia Tech's nine colleges. The Virginia Tech Transportation Institute applies AI to autonomous vehicle safety, the College of Veterinary Medicine achieves 95% accuracy in AI-powered cancer detection, and Business Information Technology examines enterprise AI adoption patterns. This distributed model–coordinated through the Sanghani Center's 36 faculty members from multiple departments–ensures AI innovation emerges from disciplinary intersections while maintaining consistent governance and ethical frameworks across the university.

Research Centers and Institutes

Sanghani Center for Artificial Intelligence and Data Analytics

Established in 2015 (originally as the Discovery Analytics Center), the Sanghani Center serves as the intellectual hub for AI research at Virginia Tech. The center brings together 36 faculty and over 200 graduate students focused on AI and data analytics research. Key features include:

- 127 active research projects as of 2025
- Extensive AI computing cluster with 24 high-end GPU nodes
- Dual presence in Blacksburg and at the Innovation Campus in Alexandria
- Focus areas spanning machine learning, natural language processing, computer vision, and AI for social impact

The center's interdisciplinary approach connects researchers from computer science, engineering, statistics, business, and other fields to tackle complex AI challenges.

Institute for Advanced Computing

The new 82,000-square-foot building (expected completion in 2026) will consolidate and expand Virginia Tech's computational research infrastructure. This facility will house research programs in:

- Advanced computing architectures and systems
- Cybersecurity and secure AI
- AI and data analytics applications
- Center for Advanced Innovation in Agriculture

The institute represents a significant investment in Virginia Tech's computational future, providing advanced facilities for AI research and development.

Commonwealth Cyber Initiative (CCI)

Virginia Tech serves as the Southwest Virginia hub for this statewide initiative investing over $25 million annually in cybersecurity research and workforce development. CCI's work increasingly focuses on the intersection of AI and cybersecurity, including:

- AI-powered threat detection and response systems
- Security of AI systems and algorithms
- Privacy-preserving machine learning techniques
- Multiple faculty positions and graduate fellowships in AI/ML security applications

Center for Human-Computer Interaction

This interdisciplinary center brings together 58 faculty affiliates studying AI's impact on human interaction with technology. The center addresses critical questions about how humans and AI systems can work together effectively, with key focus areas including:

- Ethical AI design principles and implementation
- Human-AI collaboration frameworks
- AI applications in accessibility and inclusive design
- User experience research for AI-powered systems

Recent Developments and Strategic Positioning

Major Investments and Recognition

In 2021, Virginia Tech received a $10 million gift from Mehul and Hema Sanghani, leading to the renaming of the Discovery Analytics Center in their honor. The gift supports:

- Faculty recruiting in AI and data analytics
- Graduate fellowships and research funding
- A scholars program for underrepresented minorities pursuing graduate degrees in artificial intelligence

AI as Institutional Research Frontier

AI has been designated as one of four institutional research frontiers at Virginia Tech, alongside health sciences and technology, security, and social and economic well-being. The AI frontier theme, "AI for Intelligence Augmentation," reflects the university's philosophy that AI should enhance rather than replace human capabilities. This designation brings:

- Priority funding for AI research initiatives
- Strategic hiring in AI-related fields
- Enhanced infrastructure investments
- Cross-college collaboration incentives

Innovation Campus and Regional Leadership

Virginia Tech's Innovation Campus in Alexandria, Virginia, strategically positions Virginia Tech to engage with federal agencies, policy makers, and industry partners in the national capital region. The campus focuses on graduate education and research in computer science and computer engineering, with AI as a central theme. This expansion enables Virginia Tech to:

- Bridge academic research with policy applications
- Partner with government agencies on AI initiatives
- Attract industry collaborations and talent
- Influence national AI strategy and implementation

This comprehensive AI ecosystem demonstrates Virginia Tech's substantial, well-established AI programs that support the "Established" rating for Research and Innovation Leadership identified in the readiness assessment. The combination of strong academic departments, specialized research centers, strategic investments, and regional partnerships positions Virginia Tech as a leader in AI research, education, and responsible implementation.

A.2 Adoption Patterns and Current Challenges/Successes

While specific numbers on AI tool adoption across all functions are still being gathered, initial insights from pilot programs and policy reviews offer a glimpse into current patterns and challenges:

Teaching and Learning

The university encourages both faculty and students to experiment with AI tools to explore their potential in learning, teaching, and daily tasks. Guidelines emphasize human oversight, fact-checking, and disclosing AI use in coursework. Faculty are encouraged to update syllabi to clarify expectations regarding AI tool use. The ongoing Generative AI Pilot Program, scheduled for design and recruitment in Fall 2024 and implementation in Spring 2025, aims to identify and evaluate use cases of generative AI in teaching, research, and administrative tasks.

Research

AI tools are being explored for various research applications, with guidelines emphasizing ethical considerations, data protection, and intellectual property. The partnership with Children's National Hospital highlights a specific success in applying AI to pediatric health research.

Policies like “No. 13000 – Policy on Intellectual Property” and “No. 13020 – Policy on Misconduct in Research” are being reviewed for AI-related implications, recognizing the evolving challenges posed by AI technologies in research.

Administrative Functions

AI is being considered for many potential uses in areas such as internal controls (Policy No. 3010) and contract workflows (Policy No. 3015). However, there is a clear need for supplemental guidance to ensure ethical and transparent integration of AI in these processes. The Generative AI Pilot Program also includes administrative users to explore use cases and gather feedback on cost and support requirements for AI tool implementation.

Challenges

A significant challenge is the need for clear and concrete guidance on AI usage, as highlighted in stakeholder feedback on Responsible AI Principles. There’s also a demand for clearer governance structures, enforcement mechanisms, and better integration with existing compliance frameworks. Data privacy and security remain key concerns, with strict guidelines against using AI tools with sensitive, confidential, or regulated data. The university also faces the challenge of limited formal agreements with most external AI tool providers, meaning standard university security and privacy protections may not apply.

Successes

The establishment of dedicated AI centers like the Sanghani Center and the new Institute for Advanced Computing demonstrates a commitment to and investment in AI. The university’s proactive approach in developing Responsible AI Principles and conducting pilot programs indicates a strong effort to explore AI’s benefits while mitigating risks. The positive feedback on the broad applicability of the principles and their connection to the university’s core values is also a success.

A.3 Higher Education AI Landscape

The global higher education landscape has undergone a significant transformation in AI adoption from 2022 to 2025, shifting from cautious experimentation to integration across teaching, research, and administrative functions. This market has more than doubled from $2.5 billion in 2022 to $5.88 billion in 2024. Universities are now focused on harnessing AI’s potential while navigating implementation obstacles.

General Trends in AI Adoption

Since the late 2022 launch of ChatGPT, AI adoption has dramatically accelerated. By 2024, AI adoption reached 72-78% globally, with 49% of institutions viewing AI as a strategic priority. In 2025, this focus intensified to 57%, with teaching and learning becoming the top functional area for AI implementation. Surveys in late 2024 revealed that 84% of surveyed faculty and staff reported using AI in their work or personal lives, a 32 percentage point increase from the previous year, and 93% expect to expand AI use in their work over the next two years. This growth is fueled by optimism that AI can enhance efficiency and personalize education, though concerns about ethics and privacy persist.

A.2 Detailed Assessment Results by Dimension

A.2.1 Strategic Alignment (Developing)

Virginia Tech demonstrates strong foundational elements for AI strategic alignment but requires enhanced integration across institutional planning and external partnerships.

Strengths:

- AI designated as one of four institutional research frontiers
- Clear research focus with dedicated centers and institutes
- Strong technical infrastructure supporting AI initiatives

Gaps:

- Lack of major dedicated funding commitments for institution-wide AI initiatives
- Minimal collaboration with local Virginia businesses and government agencies
- Need for better integration between research excellence and broader institutional AI adoption

A.2.2 Institutional Governance (Emerging)

Current governance structures are in development phase with significant opportunities for improvement.

Current State:

- No formal AI policies across any domain
- Undefined decision-making authority for AI initiatives
- Ad hoc risk assessment processes
- Limited formal collaboration between units
- Temporary working group structure

Immediate Needs:

- Formal AI governance framework
- Clear accountability structures
- risk assessment processes
- Permanent coordination mechanisms

A.2.3 Operational Readiness (Emerging)

Technical capabilities exist but require support infrastructure for broader adoption.

Infrastructure Strengths:

- Advanced Research Computing cluster with 208 GPUs
- Plans for additional NVIDIA H200 GPUs
- Active pilot programs with ChatGPT Edu and M365 Copilot

- Security-focused approach to AI platform testing

Support System Gaps:

- Absence of support systems and training programs
- No formal procedures for assessing, testing, or deploying new AI technologies
- Limited tracking systems for usage, costs, or effectiveness
- Financial sustainability concerns with current AI tool business models
- Restricted vendor relationships beyond Microsoft

A.2.4 AI Literacy and Ethical Use (Developing)

Strong ethical foundation exists with growing educational content but requires implementation.

Educational Content Development:

- Over 20 specialized training modules in development
- Content covers research applications, teaching integration, and workflow enhancement
- Addresses practical academic needs
- Structured learning design with video and assessment components

Implementation Challenges:

- No formal assessment of current AI literacy levels across campus
- Limited programming for undergraduate and graduate students
- Gap between sophisticated ethical principles and implementation
- Some cultural resistance regarding AI urgency among faculty and administration

A.3 Assessment Summary and Strategic Implications

Virginia Tech's AI readiness assessment reveals a distinctive institutional model characterized by leveraging exceptional technical capabilities as a foundation for transformation. Unlike institutions that adopt AI reactively, Virginia Tech has strategically positioned AI research as a cornerstone for campus-wide integration while developing necessary governance, ethical principles, and operational systems.

The assessment demonstrates that Virginia Tech's research excellence in AI (Established level) provides a unique advantage that can be leveraged to accelerate advancement across other readiness dimensions. The institution's approach emphasizes responsible integration while building on existing strengths in research, infrastructure, and educational innovation.

Appendix B: Policy Gap Analysis

This appendix provides detailed findings from the review of 168 Virginia Tech policies across nine organizational divisions to identify AI-related gaps and revision needs.

B.1 Methodology Details

The policy gap analysis employed a evaluation approach with the following components:

B.1.1 Scope and Coverage

- Total Policies Reviewed: 168 across nine organizational divisions
- Review Period: Six months with iterative evaluation cycles
- Subject Matter Experts: Assigned based on operational expertise and domain familiarity

B.1.2 Assessment Framework

Reviewers evaluated policies using standardized criteria focusing on:

- AI tool integration potential and current coverage
- Data governance requirements for AI applications
- Ethical use standards and bias considerations
- Operational oversight mechanisms for AI systems
- Adequacy of existing language for current and future AI developments

B.1.3 Categorization System

Policies were classified into three categories:

- No Action Required: Existing coverage adequate for AI applications
- Guidance Needed: Supplemental guidance beneficial but formal revision unnecessary
- Formal Revision Suggested: Policy language requires updating for AI integration

B.2 Detailed Findings by Domain

B.2.1 General University (19 policies)

Overview: Foundational policies with minimal direct AI impact Key Finding: Most policies require no action Notable Exception: Policy on Harassment, Discrimination, and Sexual Assault (1025) may need alignment with AI bias considerations

Specific Policy Analysis:

- University Mission and Goals: No AI-specific updates needed
- Code of Conduct: Existing framework adequate for AI applications
- Anti-discrimination policies: May benefit from AI bias prevention guidance

B.2.2 Academic Affairs (24 policies)

Overview: Significant AI impact potential, particularly in academic integrity Critical Gap Identified: Undergraduate Honor Code (Policy 6000) requires formal review for additional guidance and/or revision of the policy.

Detailed Policy Assessment:

- Undergraduate Honor Code (6000): Requires Formal Review

        - Current language insufficient for AI-assisted academic work
        - Need clear definitions of acceptable AI use in assignments
        - Attribution requirements for AI-generated content
        - Detection and adjudication procedures for AI misconduct
    - Academic Standards policies: May benefit from AI assessment guidance
    - Curriculum policies: Could incorporate AI literacy requirements
    - Faculty evaluation: May need AI use disclosure considerations

B.2.3 Research (35 policies)

Overview: Substantial AI integration opportunities with complex intellectual property implications Critical Gaps: Two policies require formal revision

Detailed Policy Assessment:

- Policy on Intellectual Property (13000): Requires Formal Revision
    - AI-generated works ownership and attribution unclear
    - Collaboration agreements need AI tool usage clauses
    - Patent and copyright implications for AI-assisted research
    - Revenue sharing considerations for AI-enhanced innovations
- Policy on Misconduct in Research (13020): Requires Supporting Standards, Procedures, and Guidance
    - LLM citation standards and requirements
    - AI data fabrication and falsification definitions
    - Plagiarism detection in AI-assisted research
    - Disclosure requirements for AI tool usage in publications
- IRB standards and procedures: May need AI research ethics guidance
- Data management policies: Could benefit from AI-specific protocols

B.2.4 Student Affairs (18 policies)

Overview: Moderate AI impact with specific considerations for student services and media Critical Gap: Visual Media Policy requires attention

Detailed Policy Assessment:

- Visual Media Policy (8205): Requires Formal Revision
    - AI-generated imagery and deepfake considerations
    - Consent and privacy implications for AI-enhanced media
    - Attribution requirements for AI-created visual content
    - Student safety considerations with AI-manipulated media
- Student conduct policies: May benefit from AI usage guidelines
- Privacy policies: Could incorporate AI data processing considerations

B.2.5 Human Resources (22 policies)

Overview: Substantial AI integration opportunities across recruitment, evaluation, and management processes Assessment: Review needed but no immediate formal revisions required

AI Integration Opportunities:

- **Recruitment and hiring**: AI-assisted screening and evaluation tools
- **Position description development**: AI-enhanced job analysis and writing
- **Salary determination**: AI-powered market analysis and equity assessment
- **Performance evaluation**: AI-supported review processes with human oversight
- **Training and development**: AI-personalized learning and skill assessment

Guidance Needed Areas:

- Bias prevention in AI-assisted hiring processes
- Privacy protection in AI-enhanced HR data processing
- Transparency requirements for AI-supported employment decisions
- Human oversight standards for AI-assisted HR functions

B.2.6 Information Technology (15 policies)

Overview: Critical domain requiring AI governance integration Assessment: Multiple policies need supplemental guidance

Key Areas for Enhancement:

- Acceptable Use Policies: Integration of AI Responsible Use guidelines
- Data Security Standards: AI-specific security and privacy requirements
- Customer Data Usage: AI application protocols and restrictions
- System Access Control: AI tool authentication and authorization
- Incident Response: AI-related security incident procedures

B.2.7 Finance and Administration (12 policies)

Overview: Limited direct AI impact with some operational enhancement opportunities Assessment: Most policies adequate with minor guidance needs

Potential Enhancements:

- Procurement policies: AI tool evaluation and selection criteria
- Financial management: AI-assisted budgeting and analysis guidelines
- Contract management: AI service provider agreement standards

B.2.8 Safety and Security (12 policies)

Overview: Moderate AI impact with specific surveillance and safety considerations Assessment: Some policies may benefit from AI-specific guidance

Areas for Consideration:

- University Safety and Security (5615) and Safety and Security Camera Surveillance (5617): May need AI-enhanced surveillance guidelines
- Emergency response: AI-assisted communication and coordination protocols
- Access control: AI-powered security system management
- Incident investigation: AI tool usage in security analysis

B.2.9 Outreach and Engagement (11 policies)

Overview: Minimal current AI impact with future potential for enhancement Assessment: Most policies currently adequate

Future Considerations:

- Public communication: AI-assisted content creation and distribution
- Community partnerships: AI collaboration and data sharing agreements
- Extension services: AI-enhanced service delivery and analysis

B.3 Priority Implementation Recommendations

B.3.1 Immediate Actions (0-90 days)

1. Academic Integrity Framework: Begin formal review of Undergraduate Honor Code (6000)
2. Research Integrity Standards: Initiate updates to Research Misconduct Policy (13020)
3. Intellectual Property Clarification: Start review of IP Policy (13000) for AI implications
4. IT Governance Integration: Develop AI Responsible Use guidelines for integration

B.3.2 Near-term Enhancements (6-12 months)

1. Human Resources Review: Complete evaluation of all HR policies
2. Security and Surveillance Guidance: Create supplemental guidance for AI-powered systems
3. Data Governance Standards: Establish AI-specific data handling protocols
4. Student Affairs Policy Updates: Complete Visual Media Policy revision and related updates

B.3.3 Ongoing Monitoring (12+ months)

1. Regular Review Cycles: Implement AI policy review schedule
2. Cross-functional Coordination: Establish consistent AI governance standards
3. Stakeholder Engagement: Maintain ongoing policy owner involvement in AI reviews
4. Integration with practical implementation guides: Incorporate policy updates into role-based guidance

B.4 Stakeholder Engagement Framework

Policy revision processes require coordinated engagement across multiple stakeholder groups:

B.4.1 Academic Leadership

- Faculty Senate involvement in academic integrity policy updates
- College-level dean engagement for curriculum policy considerations
- Academic department input on discipline-specific AI applications

B.4.2 Administrative Leadership

- IT governance committee participation in technology policy updates
- HR leadership involvement in employment-related policy revisions
- Legal counsel consultation for compliance and risk considerations

B.4.3 Student and Staff Representatives

- Student government input on policies affecting academic experience
- Staff council participation in operational policy considerations
- Union representative involvement where applicable

B.5 Implementation Timeline and Resource Requirements

B.5.1 Resource Allocation

- Policy revision working groups for each critical area
- Legal review and compliance assessment resources
- Communication and training materials development
- Ongoing monitoring and evaluation systems

B.5.2 Success Metrics

- Completion rates for priority policy revisions
- Stakeholder satisfaction with revised policy clarity
- Compliance rates with new AI-related policy requirements
- Reduction in AI-related policy interpretation requests

Appendix C: ChatGPT Edu Pilot Results

C.1 Pilot Overview and Methodology

The ChatGPT Edu pilot program operated from January 16 to May 16, 2025, providing controlled evaluation of institutional AI deployment with 425 full-time employees and graduate assistants. The pilot employed evaluation methodology including pre- and post-deployment surveys, usage analytics, and structured feedback collection to assess tool effectiveness, user adoption, and institutional integration requirements.

C.2 Participant Demographics and Representation

Pilot participants represented diverse university constituencies: 41% AP Faculty, 21% Tenured and Tenure-Track Instructional Faculty, 12% Staff, with representation across all colleges. The top

five participating units were College of Engineering (20%), Executive VP and Provost (10%), Pamplin College of Business (9%), Agriculture and Life Sciences (8%), and Liberal Arts and Human Sciences (8%). Primary usage areas included administrative/operations (50%), teaching and learning, and research applications.

C.3 Quantitative Results and Impact Measurement

C.3.1 Usage and Engagement Metrics

- Active User Rate: 78% weekly active users throughout pilot period
- Usage Volume: Average 34 messages per user per week (median 11 messages)
- Sustained Engagement: 425 total active participants with consistent week-to-week usage

C.3.2 Productivity and Effectiveness Results

- Productivity Increase: 94% of respondents reported increased productivity
- Task Efficiency: 98% completed tasks faster with less effort
- Time Savings: 39% saved 3+ hours weekly, 38% saved 1-3 hours weekly
- Work Quality: 95% reported improved work quality
- Idea Generation: 86% found AI helpful for generating new ideas
- Challenge Resolution: 90% reported help overcoming work challenges

C.3.3 User Satisfaction and Adoption

- Net Promoter Score: 55 (excellent level)
- Promoters: 64% of participants
- Detractors: Only 9% of participants
- Integration Ease: 81% found ChatGPT Edu easy to integrate into work routines
- Value Assessment: 77% considered the tool worth $20/month cost

C.4 Qualitative Findings and Use Case Validation

C.4.1 Most Beneficial Applications

1. Efficiency and Productivity Enhancement (33% of responses): Significant time savings in routine tasks, meeting preparation, and administrative workflows
2. Brainstorming and Idea Generation (11% of responses): Creative collaboration, problem-solving, and alternative approach development
3. Writing and Editing Support (10% of responses): Content creation, editing assistance, and communication enhancement

C.4.2 Validated Use Cases by Domain

- Coding and Programming: Debugging, code generation, optimization, and technical troubleshooting
- Writing and Communication: Email drafting, report writing, editing, and content refinement

- Research and Analysis: Literature review support, data analysis, and synthesis activities
- Administrative Operations: Meeting notes, document creation, and workflow optimization

C.4.3 Primary Challenges Identified

1. Accuracy and Hallucination Concerns (31% of responses): Need for fact-checking and verification
2. Prompting Complexity (9% of responses): Time investment in developing effective prompts
3. Dependency and Skill Attrition Concerns (6% of responses): Risk of over-reliance on AI tools
4. Privacy and Security Limitations (6% of responses): Inability to use with high-risk data

C.5 Accessibility and Equity Assessment

Accessibility review revealed mixed results: basic chat functionality (prompt entry and response) met accessibility standards, but advanced features presented significant barriers for users relying on assistive technologies. Key findings include:

- Screen reader users cannot identify many buttons or interface changes
- Keyboard-only users cannot access most functionality beyond basic chat
- Users with visual impairments face challenges with poor contrast and control boundaries
- Voice control users have difficulty targeting controls due to labeling issues

C.6 Cost-Benefit Analysis and ROI Considerations

C.6.1 Financial Overview

- **External ChatGPT Spending**: $23,122 across departments over a 12-month period. This figure does not include personal spending on ChatGPT accounts, which was not quantifiable.
- **Recommended Investment**: $120,000 annually for 1,000 ChatGPT Edu licenses, plus variable cost for credit-based access to advanced features and models.
- **User Value Perception**: 77% of survey respondents indicated that ChatGPT Edu was worth $20/month, the license price at the pilot's launch–signaling strong perceived value at current market cost.

C.6.2 ROI Analysis Caveats and Observations

Note: A true Return on Investment (ROI) calculation was *not* conducted during the pilot. The pilot did not track pre/post performance benchmarks, monetize time savings, or measure organizational outcomes tied to usage.

However, qualitative indicators suggest a strong cost-effectiveness profile:

Time Savings:

- 77% of respondents reported saving at least 1 hour per week.
- 39% reported saving 3 or more hours per week.

Cost Efficiency:

- Based on usage data, the estimated annual cost per user (licenses + credits) is approximately $50 (under pricing model afforded during initial pilot).
- Total estimated cost for 425 participants over 12 months: ~$112,000.
- Credit cost modeling indicates a potential annual range of $84,000-$196,000 depending on usage patterns.

High Engagement and Satisfaction:

- 78% of participants were active weekly.
- Net Promoter Score (NPS): 55 – considered excellent.

C.6.3 Summary

While the pilot did not generate a numerical ROI, it does provide strong evidence of:

- Perceived productivity benefits
- Positive user engagement
- Favorable cost-to-value perception

Conclusion: The pilot supports a compelling case for continued, scoped investment in ChatGPT Edu, with future work needed to quantify direct ROI through rigorous outcome tracking.

C.7 Implementation Lessons and Strategic Recommendations

C.7.1 Security and Integration Considerations

- Conservative security approach during pilot limited integration capabilities but ensured data protection
- Need for enhanced tool integration (Microsoft OneDrive, Outlook, macOS apps) to maximize productivity
- Comprehensive onboarding process critical for user adoption and effectiveness

C.7.2 Training and Support Requirements

- Mandatory 30-minute training provided foundation, but participants requested advanced implementation guidance
- Role-specific training needs identified across different user constituencies
- Ongoing support mechanisms (office hours, help desk) proved effective for issue resolution

C.7.3 Scaling and Sustainability Factors

- Strong user adoption and satisfaction support institutional deployment
- Need for accessibility improvements before full-scale implementation
- Credit-based pricing model requires careful cost management and usage monitoring

C.8 Implications for Institutional AI Adoption

The ChatGPT Edu pilot provides empirical validation for institutional AI investment while identifying critical success factors for deployment. High user satisfaction, documented productivity gains, and sustained engagement demonstrate institutional readiness for AI integration. However, accessibility challenges, integration limitations, and training needs require attention to ensure equitable and effective implementation across the university community.

C.8.1 Recommendations for Institutional Deployment

1. Proceed with recommended investment for 1,000 ChatGPT Edu licenses
2. Implement a parallel pilot of a platform that builds on Virginia Tech's Azure infrastructure to provide broad access to frontier AI models
3. Prioritize accessibility improvements before full-scale deployment
4. Develop training programs addressing advanced implementation needs
5. Enhance integration capabilities to maximize productivity benefits
6. Establish cost monitoring for credit-based features and usage patterns

These pilot results provide an essential empirical foundation for Virginia Tech's responsible AI adoption strategy, validating both the institutional readiness for AI integration and the specific requirements for implementation across teaching, research, and administrative functions.