MISSISSIPPI

Statewide AI Framework

Statewide Priorities and Learning Progression

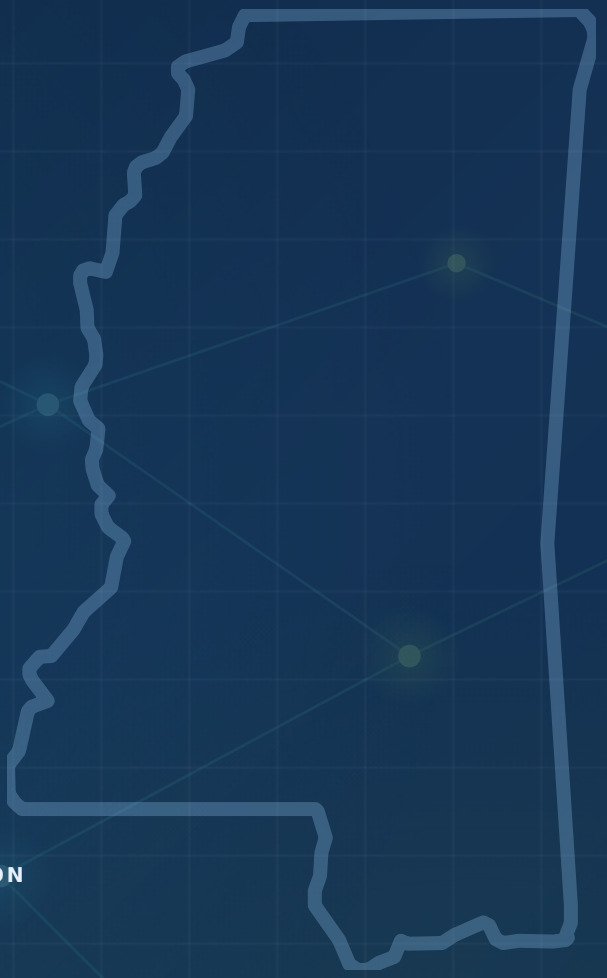

5 STATEWIDE PRIORITIES

11 COMPETENCY DOMAINS

8 LEVELS OF PROGRESSION

K-12 THROUGH THE WORKFORCE

01
K-12
Foundational awareness and practice

02
Postsecondary
Applied and technical fluency

03
Workforce
Professional application and accountability

By the AI Workforce Readiness Council, a subcommittee of the State Workforce Investment Board
Produced by the Mississippi Artificial Intelligence Network (MAIN) · Hosted by AccelerateMS

AUGUST 2026 EDITION

MAIN
MISSISSIPPI ARTIFICIAL INTELLIGENCE NETWORK

ACCELERATE MS

FIVE STATEWIDE PRIORITIES

01 **Literacy & Access**

02 **Ethics & Critical Thinking**

03 **Privacy & Governance**

04 **Workforce Readiness**

05 **Strategic Alignment**

LEARNING PROGRESSION

1 **K-12**

Foundational Awareness & Practice

2 **POSTSECONDARY**

Applied & Technical Fluency

3 **WORKFORCE**

Professional Application & Accountability

CONTENTS

What This Document Is

This is a competency framework. It describes what a Mississippian should be able to do with artificial intelligence at eight points along a path running from elementary school to senior career leadership. It is a strategic point of alignment and a leadership-oriented guide for educators, postsecondary institutions, public agencies, and industry partners.

It is organized in five parts. Part One sets the statewide priorities and the definitions used throughout. Part Two situates the framework in Mississippi law, standards, and the federal policy environment. Part Three presents the guiding principles, the eleven competency domains, and the Progression Mapping. Part Four crosswalks the domains to state, federal, and international reference instruments. Part Five covers governance and use.

What changed in this edition. The April 2026 edition established the statewide priorities, the eleven domains, and a first learning progression. This edition keeps that architecture and extends it. Part Two is new, setting out a Mississippi legal and policy foundation that covers enacted statute, the state's adopted computer science standards, the governing bodies and their current posture, and the federal environment. Factual and policy claims are now grounded in primary sources, each linking to the issuing body's own published record.

The Progression Mapping has been rewritten so that competencies describe observable capability rather than named technologies, and so that the levels after high school function as entry points rather than steps in a single sequence. Part Four adds four crosswalks to state, federal, and international reference instruments, and Part Five adds governance, measurement, and a defined review cadence built on indicators the state already collects. Several domains were extended as well, taking in how machine learning differs from human learning, intellectual property and training data, environmental and energy impacts, policy and regulatory literacy, and emerging security risks including agentic systems and synthetic media.

> **What this document does not do.** It does not set state standards, required assessments, or proficiency cut scores. It is not legislation, a requirement, or a mandate. It creates no obligation for any district, college, university, agency, or employer. It does not endorse or recommend any vendor, product, platform, or service. It is offered as a shared reference, and adoption is voluntary.

PART ONE

Mississippi's Statewide AI Priorities

The priorities below establish a shared foundation for how artificial intelligence is understood, taught, and applied across Mississippi. They define the vision, priorities, and guiding principles that shape AI education, workforce readiness, and responsible AI use, and they serve as the foundation for the AI skills, definitions, and learning progression that follow in this framework.

Mississippi's intent is to be a place where AI is not only adopted, but created, applied, and advanced. By investing in people, skills, and responsible innovation, Mississippi can foster an environment where educators, entrepreneurs, researchers, and workers contribute to new ideas, new industries, and new economic opportunities. The goal is a distinctly Mississippi approach that uses local strengths while competing nationally and globally.

These priorities provide clear statewide direction while preserving flexibility for institutions, educators, employers, and agencies to implement AI initiatives in ways that reflect their missions and local contexts. Together, they position Mississippi to proactively shape AI adoption in ways that expand opportunity, support innovation, and strengthen long-term economic resilience for communities across the state.

Definitions Used in This Framework

Artificial intelligence. This framework uses the definition adopted into the Mississippi Code by HB 1723 (2026 Regular Session, Chapter 311, Laws of 2026, effective 1 July 2026), which adopts the language of the federal definition at 15 U.S.C. § 9401(3):

> "The term 'artificial intelligence' means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action."

AI literacy. This framework adopts the definition finalized by the U.S. Department of Education on 13 April 2026 (91 FR 18774, effective 13 May 2026):

> "Artificial intelligence (AI) literacy means the technical knowledge, durable skills, civic awareness and future ready attitudes, including AI related ethical reasoning, critical social inquiry, interdisciplinary problem-solving, and creativity, required to thrive in a world influenced by AI. It enables learners to engage, create with, manage, and design AI, while critically evaluating its benefits, risks, and implications."

How a model is obtained and run. This framework distinguishes systems reached only as a hosted service from systems whose weights are, in the words of the U.S. Department of Commerce, "open to the public by allowing users to download these weights from the Internet or through other mechanisms", commonly called open-weight or open-source models. Each arrangement carries advantages the other does not. A hosted service supplies access to the most capable systems available, requires no infrastructure or maintenance, and leaves safety testing, security patching, and support with the vendor. Weights an organization runs itself keep data in its own environment, permit inspection and adaptation, and reduce dependence on a single supplier's pricing and continued service.

Postsecondary attainment and credential of value. Where this framework refers to credentials, it uses the definitions established by SB 2524 (2026), Chapter 483, Laws of 2026: postsecondary attainment means a workforce certification, an industry credential, or an academic degree leading to an occupation in the public or private sector, and a credential of value is affiliated with a private sector priority occupation, as designated by the state's Office of Workforce Development, leading to a high-growth, high-wage occupation.

Priority 1: AI Literacy and Access for All

AI learning should be available to all Mississippians, regardless of geography, background, or socioeconomic resources. The statewide approach should expand opportunity, reduce barriers to participation, and avoid creating "AI haves and have-nots," while ensuring that every learner has access to foundational AI literacy.

This priority recognizes that broad access and baseline understanding are essential for equitable participation in an AI-enabled society.

WHAT THIS MEANS IN PRACTICE

- Foundational AI literacy should be broadly available, not reserved for specialized or advanced programs.
- Access should include rural communities, under-resourced schools, adult learners, and incumbent workers.
- Access includes learners with disabilities and learners still acquiring English. Where an AI tool becomes part of instruction, assessment, or a public-facing service, it should be usable by those learners, and existing obligations to them continue to apply.
- Every learner should have a baseline understanding of what AI is, how it is used, how to evaluate AI outputs, and when AI should not be used.
- Advanced pathways in development, data, and AI-enabled systems should be available for learners with interest and aptitude, without making specialization a requirement for participation.
- Access should account for employer size as well as geography. National figures from the Census Business Trends and Outlook Survey put artificial intelligence use at 37 percent among firms with 250 or more employees and under 20 percent among firms with four or fewer, and show use rising among firms with at least 20 employees while holding steady below that size. Where this pattern holds in Mississippi, workers at small firms are less likely to encounter these tools at work and more dependent on public education and training to reach baseline literacy.

Priority 2: Ethics, Responsibility, and Critical Thinking

AI education and use should be hands-on and anchored in ethical reasoning and responsible practice, including privacy, acceptable use, bias awareness, and critical evaluation. Mississippi should prepare learners to be responsible users and leaders, not passive consumers of AI systems.

This priority ensures that AI adoption is guided by human judgment, integrity, and awareness of societal impact.

WHAT THIS MEANS IN PRACTICE

- Ethical reasoning is integrated into how AI is taught and used, rather than treated as a separate or optional topic.
- Learners should understand AI limitations, misuse risks, and the importance of validating and questioning AI-generated outputs.
- Responsible use includes awareness of bias, fairness, transparency, and appropriate contexts for AI use.

Priority 3: Data Privacy, Security, and Responsible Governance

AI education and adoption should emphasize the protection of data, secure use of AI systems, and awareness of governance and safety considerations. Learners and practitioners should understand how AI systems rely on data, networks, and digital infrastructure, and how misuse or poor practices can introduce risk.

This priority reinforces trust, safety, and resilience in AI-enabled environments.

WHAT THIS MEANS IN PRACTICE

- Learners should understand basic principles of data privacy, consent, and responsible data handling.
- Awareness of cybersecurity risks related to AI, including data leakage, misuse, and adversarial threats, should be incorporated into AI learning.
- Responsible governance includes understanding user responsibility, acceptable use, and the need for human oversight in AI-enabled systems.

Priority 4: Workforce Readiness and Flexibility

AI education should connect to Mississippi's industries and workforce needs while preparing learners to be adaptable as tools, roles, and processes change. The focus should be on durable skills that transfer across occupations and evolve over time, rather than training for a single tool or moment in time.

This priority supports a skilled workforce capable of operating in diverse and changing work environments.

WHAT THIS MEANS IN PRACTICE

- AI-related skills should translate across sectors and adapt to changing workplace demands.
- Employer partnerships should inform relevance and real-world context without narrowing education to short-term vendor or product training.
- Advanced depth and specialization should be available for roles that require it, while foundational readiness remains broadly applicable.
- Workforce readiness includes incumbent workers and career changers who will not re-enter the K-12 pipeline. Entry into this progression should be possible at any level without prerequisite coursework from an earlier one.
- Credential recommendations should distinguish independently assessed certifications from vendor learning products, and should specify a minimum instructional dosage rather than treating all short-term training as equivalent.

Priority 5: Strategic Alignment with Statewide Initiatives and Measurable Outcomes

The statewide AI approach should align with Mississippi's broader priorities, including education, workforce development, government modernization, and economic growth. Clear, shared outcomes support coordination, accountability, and continuous improvement across sectors.

This priority ensures coherence and sustainability without imposing rigid compliance structures.

WHAT THIS MEANS IN PRACTICE

- Shared priorities help reduce fragmentation and duplication across institutions and initiatives.
- Outcomes should be measurable and transparent, supporting evaluation and improvement without turning the priorities into a compliance instrument.

- Measurement should draw on indicators the state already collects rather than new instruments: credential attainment reportable under Perkins V and the Workforce Innovation and Opportunity Act, wage outcomes through the existing unemployment insurance wage-record match, and completions in courses qualifying under Miss. Code § 37-13-215.
- Employer-side and demand-side conditions should be tracked using recurring public sources rather than one-time surveys.

Together, these priorities set the direction for what follows. Part Two situates them in the Mississippi law, standards, and federal policy that already govern this ground. Part Three then presents the AI Learning Progression: its purpose and scope, the Guiding Principles, the AI Literacy Pillars, and the Domains and Progression Mapping that translate these priorities into stage-specific descriptors and a pathway that runs from K-12 through postsecondary and into the workforce.

PART TWO

Mississippi Legal and Policy Foundation

This framework is not law and creates no obligation. It operates inside a Mississippi legal and policy environment that already addresses artificial intelligence, and inside a federal environment that shapes what a state instrument of this kind can usefully be. Both are set out here so the framework can be read against them.

Statute

Miss. Code § 37-13-207 (Mississippi Computer Science and Cyber Education Equality Act, HB 633, 2021; Laws 2021, ch. 389, § 4, effective 1 July 2021) directs the State Department of Education to implement K-12 computer science curriculum "which includes instruction in, but not limited to" eight named subjects: "(a) Computational thinking; (b) Problem solving; (c) Programming; (d) Cyber security; (e) Data science; (f) Robotics; **(g) Artificial intelligence and machine learning;** and (h) Other computer science and cyber-related content." Artificial intelligence and machine learning have been statutorily required K-12 content in Mississippi since 2021.

Miss. Code § 37-13-215, created by SECTION 29 of SB 2294 (2026 Regular Session), Chapter 479, Laws of 2026, approved by the Governor 8 April 2026, effective 1 July 2026:

> "(1) This section shall be known and may be cited as the 'Mississippi Future Innovators Act.'
>
> (2) Beginning with the entering ninth-grade class of 2029-2030, a public high school student shall, before graduation, be required to earn one (1) unit of credit in a high school computer science course, or one (1) unit of credit in an industry-aligned career and technical education (CTE) with embedded computer science course providing instruction in the foundations of computer science. Either credit unit must be approved by the State Board of Education. This computer science requirement shall not increase the number of state-required credits for graduation.
>
> (4) The computer science course or courses fulfilling this graduation requirement must include instruction on the fundamental concepts of emerging computer science technologies, such as artificial intelligence (AI), including an understanding of what AI is, how it works, and its impact on society."

HB 1723 (2026), Chapter 311, Laws of 2026, approved 9 March 2026, effective 1 July 2026, codified in Chapter 3, Title 1, Miss. Code of 1972, supplies the statutory definition of artificial intelligence used throughout this framework. It creates no obligation and restricts nothing.

SB 2524 (2026), Chapter 483, Laws of 2026, approved 8 April 2026, effective 1 July 2026, establishes the Mississippi Postsecondary Attainment Council and supplies the definitions of postsecondary attainment and credential of value adopted above. The act provides that the council "shall not displace any existing, codified governing or coordinating bodies or their responsibilities."

Standards and Curriculum

The **2026 Mississippi College- and Career-Readiness Standards for Computer Science** were adopted by the Mississippi State Board of Education in two unanimous actions, on 29 April 2026 (approval to begin the Administrative Procedures Act process) and 18 June 2026 (approval to revise, having cleared that process without public comments). They carry a required implementation year of 2026-2027 and supersede the 2018 standards, which contained no artificial intelligence content.

The 2026 standards comprise 321 standards adapted from the Computer Science Teachers Association PK-12 Computer Science Standards, Revised 2026. They are organized around five concepts (Computing & Society, Algorithms & Design, Programming, Data & Analysis, Systems & Security), four pillars expressed as twelve practices, and seven dispositions, with discrete grade-level standards for Kindergarten through grade 6, a Secondary Foundational band for grades 7-12, and two tiers of High School Specialty Standards across six specialty areas.

The standards state that "AI-related content is distributed across the five concepts instead of being a discrete concept." Machine Learning is a named subconcept under Algorithms & Design with a progression running from Kindergarten through grade 12, Artificial Intelligence is one of the six high school specialty areas with 27 standards across two tiers, and roughly forty standards name artificial intelligence or machine learning in the standard statement itself.

The **Exploring Computer Science** career and technical curriculum framework, approved by the State Board of Education on 16 January 2025 (course code 000283, one Carnegie unit, 140 hours, entered as freshmen), contains eight units, one of which is Artificial Intelligence, alongside Data Science and Computing and Cybersecurity units.

Relationship to the 2026 Computer Science Standards

The Mississippi Department of Education's standards are the operative instrument for K-12 artificial intelligence content, and this Council endorses them as such. They carry implementation force and they specify that content in more detail than a voluntary framework should attempt. Crosswalk 4 in Part Four maps each of this framework's eleven domains to where the corresponding content sits in the standards.

The 2026 standards reach grade 12. This framework carries the same competencies onward through the CC/Technical, University, Early Career, and Mid-to-Senior levels, which no Mississippi instrument currently addresses. Together the two describe the full path from elementary school to senior career leadership.

A note on terminology. This framework's four AI Literacy Pillars and the 2026 standards' four Pillars are different structures that share a word. The standards' pillars are Ethics & Social Responsibility, Inclusive Collaboration, Computational Thinking, and Human-Centered Design. This framework's pillars are organizing themes for its eleven domains and are a different structure from the standards' practices. Where a district is building scope and sequence, the standards' terminology governs. Grade banding also differs. The standards carry discrete grade-level standards from Kindergarten through grade 6 and a Secondary Foundational band for grades 7-12; this framework uses K-5, 6-8, and 9-12. Grade 6 is the only grade the two place differently: this framework groups it with grades 7 and 8, while the Department groups it with the grade-level standards below it. Placement decisions should follow the Department's banding.

Governing Bodies and Current Posture

The AI Workforce Readiness Council is an advisory subcommittee of the State Workforce Investment Board, with representation from the Mississippi Department of Education, the Mississippi Community College Board, the Institutions of Higher Learning, MAIN, and industry. It does not have authority over them. Every alignment described in this framework is a statement of shared direction, not a mandate.

The Council was established in June 2025 through the State Workforce Investment Board and AccelerateMS, following the Board's direction that AccelerateMS establish a council to assess and expand curricula to include artificial intelligence and machine learning. The Mississippi AI Talent Accelerator Program announcement of 12 June 2025 publicized the Council and assigned it a coordinating role across the funded institutions.

The Board returned to artificial intelligence in its recommendations of 4 September 2025, which call for a Mississippi Workforce Pell Readiness Task Force drawn from AccelerateMS, the Institutions of Higher Learning, the Mississippi Community College Board, the Mississippi Development Authority, the Mississippi Department of Employment Security, the Mississippi Department of Education, and industry. Among the items that task force is asked to recommend is "how AI and digital skills may be incorporated, potentially through AI apprenticeships or bootcamps." The same recommendations direct annual fall meetings with employers, beginning in 2025, to capture competency gaps employers identify. This framework is available to either effort as a reference for what those skills consist of.

Body	Current posture
Mississippi AI Regulation Task Force	Created by SB 2426 (2025). Report One to the Legislature, 13 January 2026, identified 232 active or planned state agency AI projects in FY2025, up from 63 in FY2024. It makes two recommendations to the Legislature: that the state adopt a statutory definition of artificial intelligence matching 15 U.S.C. § 9401(3), which the Legislature did the following session in HB 1723, and that agencies disclose clearly when citizens are interacting with conversational AI systems. The report separately describes stackable AI education pathways across technical, undergraduate, and graduate levels as work already under way, and identifies MAIN as the statewide implementation and measurement backbone. The Task Force reports annually and dissolves 31 December 2027.
Department of Information Technology Services	The Acceptable Use Policy for Artificial Intelligence (PSG-100-03, approved 25 November 2025 under Executive Order 1584) requires approval for AI use touching personally identifiable or nonpublic information, mandatory training, and human review of outputs before action, and prohibits artificial intelligence as sole final decision-maker. Its stated scope is ITS operations, so it does not by its terms cover other agencies. All state agencies are bound by the Enterprise Security Policy (PSG-100-09). The Mississippi State Personnel Board's Artificial Intelligence (AI) Use for HR Policy, posted 7 April 2026, offers an optional template that agencies may incorporate into their own AI use policies.
Institutions of Higher Learning	The Board of Trustees governs Mississippi's public universities, which cover the University (Fr-So) and University (Jr-Sr) levels of this progression. Campus-level artificial intelligence programming is developing across the system, including a Bachelor of Science in Artificial Intelligence launched in 2024 and a Master of Science in Artificial Intelligence launched in 2025 at Mississippi State University and the work funded through the Mississippi AI Talent Accelerator Program. This framework offers a common reference for that work.
Mississippi Community College Board	The Board develops and maintains the statewide career and technical curriculum frameworks used by Mississippi's community colleges, which cover the CC/Technical level of this progression. Adjacent frameworks include Data Analytics Technology, Network Security Technology (Cyber Security), and Extended Reality (XR) Courses and Certificates. The Council is available to support the framework process wherever artificial intelligence content is under consideration.
Workforce Pell	Final rule published 19 May 2026, effective 1 July 2026. Eligible programs run at least 8 and fewer than 15 weeks, carry 150 to 599 clock hours or the credit-hour equivalent, and must annually show 70 percent completion within 150 percent of normal time and 70 percent of completers employed in the second quarter after exit, with published tuition and fees not exceeding value-added earnings. The Governor, consulting the State Workforce Investment Board, identifies qualifying programs; AccelerateMS began accepting institutional applications 1 July 2026 and reviews them in consultation with the State Workforce Investment Board.

Federal Policy Environment

Executive Order 14277, "Advancing Artificial Intelligence Education for American Youth" (23 April 2025, 90 FR 17519). The order establishes a White House Task Force on AI Education and the Presidential AI Challenge, and directs the Secretary of Education to prioritize artificial intelligence in teacher professional development and the Secretary of Labor to expand AI-related Registered Apprenticeships. For institutions asking what federal funds may pay for, the operative document is the Secretary of Education's guidance of 22 July 2025, *Guidance on the Use of Federal Grant Funds to Improve Education Outcomes Using Artificial Intelligence*, which addresses the use of federal formula and discretionary grant funds for artificial intelligence consistent with existing program requirements.

America's AI Action Plan (July 2025). Under "Train a Skilled Workforce for AI Infrastructure," the Plan directs the Departments of Labor and Commerce to convene stakeholders to "develop or identify national skill frameworks and competency models" for AI infrastructure occupations, and states that these frameworks "would provide voluntary guidance that may inform curriculum design, credential development, and alignment of workforce investments." The Plan further directs early career exposure and pre-apprenticeships for middle and high school students, Department of Education guidance to state and local career and technical education systems, Department of Energy partnership with community and technical colleges, and pilots "carried out by states and workforce intermediaries" under the Workforce Innovation and Opportunity Act. States are named as the executing party.

National Policy Framework for Artificial Intelligence: Legislative Recommendations (20 March 2026). Section VI recommends that "Congress should use non-regulatory methods to ensure that existing education programs and workforce training and support programs, including apprenticeships, affirmatively incorporate AI training," and that Congress "bolster capabilities at land-grant institutions" for technical assistance, demonstration projects, and AI youth development programs.

Executive Order 14365, "Ensuring a National Policy Framework for Artificial Intelligence" (11 December 2025, 90 FR 58499), addresses the relationship between federal and state artificial intelligence policy, including conditions attached to certain federal funds. This framework sets no state requirement on developers or employers, creates no preemption exposure, and is designed to remain useful regardless of how that relationship develops. Its recommendations are directed to what the state itself controls: its schools, colleges, agencies, and procurement.

U.S. Department of Labor. Training and Employment Guidance Letter 03-25 (26 August 2025) encourages state workforce agencies and boards to use WIOA Title I Youth, Adult, and Dislocated Worker funds to help participants develop artificial intelligence skills. Training and Employment Notice 07-25 (13 February 2026) issues the Department's AI Literacy Framework, with five content areas and seven delivery principles, and states that "State and local agencies can use this framework to advance AI literacy across the public workforce and education systems." Crosswalk 1 maps this framework's domains to it. The Department announced the AI in Registered Apprenticeship Innovation Portal on 29 April 2026, which supplies industry-specific AI skill modules, National Guideline Standards for AI-integrated programs, and options for adding artificial intelligence content to existing apprenticeships through On-the-Job Learning and Related Technical Instruction.

President Trump's Cyber Strategy for America (March 2026). Pillar 6 calls for a cyber talent pipeline that is "pragmatic and accessible, reconciling and taking advantage of existing avenues within academia, vocational and technical schools, corporations, and venture capital opportunities, to educate and train our existing cyber workforce across industries and occupations." Pillar 5 treats artificial intelligence security and agentic-AI defense as core cyber work. Both support the inclusion and scope of Domain 11.

Executive Order 14409, "Promoting Advanced Artificial Intelligence Innovation and Security" (2 June 2026, 91 FR 34565), addresses federal cybersecurity, the review of frontier models before release, and criminal enforcement. Section 2 directs the Department of Homeland Security, through the Cybersecurity and Infrastructure Security Agency, to "facilitate access to cybersecurity tools and services including, where appropriate, covered frontier models for agencies, State and local authorities, and operators of critical infrastructure such as rural hospitals, community banks, and local utilities," and directs the Office of Personnel Management to expand federal cybersecurity specialist hiring and placement pathways. Section 3 states that nothing in the order authorizes "a mandatory governmental licensing, preclearance, or permitting requirement for the development, publication, release, or distribution of new AI models, including frontier models." Section 4 directs enforcement against anyone who uses artificial intelligence, including AI agents, to gain unlawful access to computer systems. The order bears on the scope of Domain 11 and on the definition of how a model is obtained and run.

Science: A New Golden Age (Office of Science and Technology Policy, July 2026). The report recommends requiring universities and community colleges to embed practical technical training and externships into STEM curricula, extending registered apprenticeships into science and technology fields, and scaling Workforce Pell. Among its recommendations is to "[a]dopt pay-for-performance funding models, scale Workforce Pell Grants, and back community colleges as regional hubs of scientific and technical talent." The report's emphasis on community colleges as sites of applied technical training corresponds to the weight this framework gives the CC/Technical level of the progression.

PART THREE

AI Learning Progression

The AI Learning Progression maps the agreed AI skill domains into a learning progression that articulates what learners should be able to do in each domain across three stages:

- K-12 (Foundational Awareness & Practice)
- Postsecondary (Applied & Technical Fluency)
- Workforce (Professional Application & Accountability)

This is a living document. Because the technology, the labor market, and the surrounding law are all moving quickly, this framework is reviewed on a fixed annual cycle each January and revised out of cycle whenever a defined trigger occurs.

Guiding Principles

- The goal is not to replace human judgment but to amplify it.
- Focus on durable, transferable skills that endure as tools evolve.
- Ground all learning in Mississippi's economic and community realities (agriculture, coastal resilience, advanced manufacturing, healthcare, hospitality, and public service).
- Prioritize equity, safety, accountability, and responsible innovation in every context.
- Recommend, do not mandate. This framework's usefulness depends on being adoptable by institutions with different missions and different governing bodies.
- Cite what can be verified, and say plainly what cannot. Where an evidence base does not exist, name the gap instead of filling it.
- Contribute to the evidence base, not only cite it. Where this framework names a gap, Mississippi institutions are positioned to help close it, and findings generated here can inform the next edition.
- Default to enabling. Caution in this framework attaches to consequential decisions, not to learning.

AI Literacy Pillars

To make the eleven domains more memorable and actionable, they are organized under four interconnected AI Literacy Pillars. These pillars show how the skills work together in real use and serve as a one-page overview for educators, administrators, and industry partners.

AI Understanding & Mindset
Domains 1, 5, 8, 9

Responsible Data & Evaluation Practices
Domains 2, 3, 7, 11

Effective Human-AI Collaboration
Domains 4, 6

Contextual Application & Adaptation
Domain 10 plus cross-cutting application

Each domain section below is marked with the color of the pillar it belongs to.

The North Star: Human Agency, Judgment, and Restraint

Every other domain in this framework exists to support one capability. A Mississippian should be able to decide, deliberately and with reasons, when to use these systems and when not to.

This framework is organized around competent judgment, of which restraint is a part. Judgment also means recognizing when to proceed. A person who never uses these systems has not exercised judgment either. The distinction matters in practice.

- **Domains 1 through 6** build the technical understanding required to judge well. Someone who cannot tell a generated claim from a verified one, or trace where training data came from, has no basis for restraint. They have only caution or credulity.
- **Domain 7** supplies the ethical reasoning that turns a judgment into a defensible one.
- **Domains 9 and 10** make the judgment communicable and transferable across roles and sectors.
- **Domain 11** protects the conditions under which any of it can be trusted.
- **Domain 8** is what the other ten are for.

Two instruments in this framework operationalize it. The High-Stakes Human-Review Triggers name the categories where this framework calls for a human decision regardless of how good the output looks. The Progression Mapping carries restraint through all eight levels, from a fifth grader marking which steps a person must do, to an executive setting the organizational standard for when a system is paused, overridden, or retired.

Mississippi's public sector has already written this into enforceable policy. ITS PSG-100-03 prohibits artificial intelligence as sole final decision-maker and prohibits substituting it for human judgment in ethical or sensitive decisions. The state's 2026 computer science standards carry it into K-12 through S1.AIN.10, which asks students to promote safeguards that ensure meaningful human involvement in decision-making.

The purpose of artificial intelligence is to strengthen human judgment.

AI Learning Progression Domains

The following eleven domains set out the AI-related skills, concepts, and capability areas this framework treats as necessary when building a K-12 → Postsecondary → Workforce AI learning progression in Mississippi.

The Mississippi examples in each domain are illustrative scenarios drawn from the state's industries and public services. They show what a competency looks like in context and do not describe documented deployments.

AI UNDERSTANDING & MINDSET

Domain 1. AI Foundations and Conceptual Understanding

DEFINITION

Understanding what AI is and is not, how it differs from traditional software, and how learning-based systems generate outputs from data rather than rules alone.

INCLUDES

- AI vs. automation vs. traditional software
- Machine learning as pattern recognition
- Generative AI vs. predictive/classification systems
- Limits, uncertainty, and error
- The distinction between how people learn and how these systems learn, and what that difference implies about where each is reliable
- How a system is obtained and run, from a hosted service to weights an organization downloads and operates itself, and what each arrangement implies for cost, control, verification, and where data travels

WHY THIS DOMAIN EXISTS

Without this, all other AI use is either magical thinking or blind trust.

LEARNERS CAN

Learners can explain the difference between a rules-based tool (calculator, if-then script) and a learning-based system (spam filter). They can say what GenAI is doing (predicting plausible text/images) versus a classifier (labeling). They expect mistakes: "AI can sound confident and still be wrong."

MISSISSIPPI EXAMPLES

- K-12: "Is a pilot-assist system on a Mississippi River towboat 'AI' or just automation?"
- Postsecondary: "Why does a precision-ag yield model in a Delta soybean field fail on edge cases?"
- Workforce: "When should we not trust an AI summary of weld-inspection data in shipbuilding without checking?"

RESPONSIBLE DATA & EVALUATION PRACTICES

Domain 2. Data Literacy and Data Stewardship

DEFINITION

Understanding how data is collected, represented, evaluated, protected, and governed, and how data choices shape AI behavior.

INCLUDES

- Data sources and data quality
- Bias, representativeness, and sampling
- Privacy, consent, and sensitive data, including the federal rules that govern education records under FERPA and health information under HIPAA
- Responsible data handling in education and work contexts
- Intellectual property, attribution, and the use of creators' work as training data

WHY THIS DOMAIN EXISTS

AI outcomes are inseparable from data decisions.

LEARNERS CAN

Learners can identify data inputs and how bad/biased data changes outcomes. They can describe representativeness: “If our examples only come from one district, the tool may fail statewide.” They apply privacy rules: don’t paste student records, SSNs, patient info, or confidential HR data into AI tools.

MISSISSIPPI EXAMPLES

- K-12: “Why would a Delta cotton irrigation dataset from field sensors need cleaning before training an AI model?”
- Postsecondary: “How do you document data lineage for a Gulf Coast shrimp-harvest supply-chain AI project?”
- Workforce: “How do we redact, de-identify, and classify patient records before using AI in a Mississippi hospital under HIPAA?”

RESPONSIBLE DATA & EVALUATION PRACTICES

Domain 3. Critical Evaluation of AI Outputs

DEFINITION

The ability to assess AI-generated outputs for accuracy, bias, relevance, and appropriateness, and to revise or reject outputs using human judgment.

INCLUDES

- Verification against trusted sources
- Identifying hallucinations and unsupported claims
- Understanding confidence vs. correctness
- Treating outputs as drafts, not answers

WHY THIS DOMAIN EXISTS

This is the primary safeguard against misuse and overreliance.

LEARNERS CAN

Learners treat AI as a first draft and routinely check claims. They can label failure modes: hallucination, outdated info, missing context, fabricated citations. They can triangulate using primary sources or trusted references.

MISSISSIPPI EXAMPLES

- K-12: “Show me where this AI claim about hurricane evacuation routes is supported by current state emergency management data.”
- Postsecondary: “Verify an AI-generated analysis of Mississippi River flooding patterns against official streamgage data.”
- Workforce: “Verify AI-generated policy recommendations against official coastal resilience dashboards before using in a coastal resilience report.”

EFFECTIVE HUMAN-AI COLLABORATION

Domain 4. Human-AI Interaction (Prompting and Iteration)

DEFINITION

The ability to effectively communicate intent, context, and constraints to AI systems and to iteratively refine inputs to improve outputs.

INCLUDES

- Asking clear, purpose-driven questions
- Providing audience, tone, and constraints
- Iteration and comparison of outputs
- Understanding how prompts shape results

WHY THIS DOMAIN EXISTS

This is the practical interface between humans and AI across all roles.

LEARNERS CAN

Learners can write prompts that specify task + audience + constraints + success criteria. They iterate: “Try again with fewer assumptions” or “give two options and tradeoffs.” They compare outputs and choose based on rubric.

MISSISSIPPI EXAMPLES

- K-12: “Draft a 150-word public service announcement about Delta precision-ag water conservation for local farmers, neutral tone, cite only trusted extension service data.”
- Postsecondary: “Draft a lab report outline on AI-assisted coastal erosion modeling using IMRaD format, no fabricated citations.”
- Workforce: “Iterate an AI-generated quality-inspection checklist for shipbuilding weld seams until it meets OSHA and industry standards.”

AI UNDERSTANDING & MINDSET

Domain 5. Algorithmic Thinking and Problem Decomposition

DEFINITION

Breaking tasks or problems into structured steps, identifying decision points, and determining where AI assistance is appropriate.

INCLUDES

- Step-by-step reasoning
- If-then logic and sequencing
- Identifying automatable vs. human tasks
- Translating real-world problems into workflows

WHY THIS DOMAIN EXISTS

This connects AI use to broader computational and systems thinking.

LEARNERS CAN

Learners can break a messy goal into steps and decide where AI helps versus where humans must decide. They can create decision points and checks.

MISSISSIPPI EXAMPLES

- K-12: “If the output affects student safety during a hurricane drill, then require human review.”
- Postsecondary: “Convert ‘optimize irrigation in a Sunflower County soybean field’ into data collection → AI pattern analysis → human farmer approval → evaluation.”
- Workforce: “Convert ‘improve weld quality in shipbuilding’ into sensor data → AI defect detection → human inspector safety review → final sign-off.”

EFFECTIVE HUMAN-AI COLLABORATION

Domain 6. AI-Enabled Workflow and System Design

DEFINITION

Understanding how AI tools operate within larger systems that include people, processes, data, and other technologies, and how to design workflows with oversight.

INCLUDES

- Human-in-the-loop design
- Integration with existing processes
- Checkpoints, escalation, and review
- Awareness of system-level risks
- Awareness of AI system architecture for effective integration and oversight
- Multi-step agentic workflows where AI systems act autonomously across tools
- Where a workflow's data physically goes, who controls the model making each decision, and what changes when that model runs on infrastructure the organization operates

WHY THIS DOMAIN EXISTS

Most AI use is embedded in systems, not standalone tools.

LEARNERS CAN

Learners can map an AI workflow end-to-end: inputs, model/tool, outputs, human review, logging, escalation, storage. They design guardrails: approvals, audit trails, red-team testing, fallback procedures.

MISSISSIPPI EXAMPLES

- K-12 institutional practice: "AI drafts parent letters about school hurricane closure; staff edits and approves before sending."
- Postsecondary: "In a hospitality management course, AI forecasts staffing for a Gulf Coast hotel ahead of a storm closure; a manager approves the schedule before it is posted."
- Workforce: "AI triages support tickets in coastal emergency management; high-risk escalated to a human with agentic hand-off logged."

RESPONSIBLE DATA & EVALUATION PRACTICES

Domain 7. Ethical Reasoning and Responsible AI Use

DEFINITION

Applying ethical principles and professional norms to AI use, including fairness, accountability, transparency, privacy, and integrity.

INCLUDES

- Fairness and bias considerations
- Academic and professional integrity
- Transparency and explainability
- Appropriate vs. inappropriate use cases
- Environmental and energy impacts of large-scale adoption, including compute, water, and grid demand

WHY THIS DOMAIN EXISTS

Ethics constrains how every other skill is applied.

LEARNERS CAN

Learners can identify “should we” questions, not just “can we.” They recognize high-stakes contexts and fairness risks. They practice disclosure: “This was AI-assisted.” They can argue appropriate use cases (brainstorming, drafting, accessibility supports) vs inappropriate (doing graded work dishonestly, replacing required human judgment, generating fake evidence).

MISSISSIPPI EXAMPLES

- K-12 institutional practice: “Should an AI tool recommend student discipline actions in a Mississippi public school, or must a counselor always decide?”
- Postsecondary: “Is it ethical to use AI to generate reflections on a healthcare ethics case study involving real Mississippi hospital patient data?”
- Workforce: “How do we explain AI’s role to stakeholders in a Delta precision-agriculture co-op project while protecting farmer privacy?”

AI UNDERSTANDING & MINDSET

Domain 8. Human Agency, Judgment, and Restraint

DEFINITION

Maintaining human responsibility, oversight, and intentional decision-making in AI-enabled environments, including knowing when not to use AI.

INCLUDES

- Avoiding automation bias
- Preserving accountability
- Reflecting on overuse and dependency
- Identifying human-only decisions

WHY THIS DOMAIN EXISTS

This protects learning, trust, and professional responsibility. This domain serves as the explicit North Star for the entire framework. Every other skill ultimately supports the ability to keep humans in control and make wise decisions about when AI should (or should not) be used.

LEARNERS CAN

Learners keep responsibility with humans: “AI advised; I decide.” They can name when not to use AI: sensitive decisions, unclear context, when it undermines learning goals, when it reduces trust. They spot automation bias: believing the tool because it sounds polished.

MISSISSIPPI EXAMPLES

- K-12: “Require independent student reasoning before using AI on any graded assignment about coastal resilience.”
- Postsecondary: “Require independent student reasoning before using AI on an emergency-response ethics case study.”
- Workforce: “AI advised on a shipbuilding weld repair; I decide, document the rationale, and retain accountability.”

AI UNDERSTANDING & MINDSET

Domain 9. Communication and AI Literacy Across Audiences

DEFINITION

The ability to explain AI concepts, uses, limitations, and implications clearly to technical and non-technical audiences.

INCLUDES

- Plain-language explanations
- Communicating risks and limitations
- Supporting informed decisions
- Cross-disciplinary translation
- Policy and regulatory literacy: reading, comparing, and explaining the rules that govern these systems

WHY THIS DOMAIN EXISTS

AI adoption fails without shared understanding.

LEARNERS CAN

Learners can explain AI without hype to parents, students, board members, supervisors, or coworkers. They can communicate limits: uncertainty, bias, privacy, and what controls are in place.

MISSISSIPPI EXAMPLES

- K-12: “Student explains to parents why an AI weather app can be wrong during hurricane season.”
- Postsecondary: “Student explains model evaluation to a non-technical partner on a Mississippi River flooding project.”
- Workforce: “Leader summarizes an AI pilot for a Delta precision-agriculture co-op: purpose, data used, guardrails, results, next steps, and risks.”

CONTEXTUAL APPLICATION & ADAPTATION

Domain 10. Sector and Pathway Awareness

DEFINITION

Understanding how AI skills apply differently across education levels, career pathways, and workforce sectors (cross-cutting lens).

INCLUDES

- Sector-specific AI use cases
- Transition points (K-12 → CC → university → workforce)
- Transfer vs. direct-to-workforce pathways
- Adaptability as tools evolve

WHY THIS DOMAIN EXISTS

This grounds the learning progression in real transitions, not abstraction.

LEARNERS CAN

Learners can connect AI skills to real roles and sectors, and they understand that “using AI well” looks different in healthcare vs manufacturing vs education vs public sector. They can describe transition expectations at key checkpoints (middle school readiness vs high school vs CC credential vs university specialization vs workplace performance). They can choose a pathway: transfer-oriented (deeper theory, research, advanced math/CS) vs direct-to-workforce (applied tools, workflows, compliance, domain constraints). They stay adaptable: focus on durable skills (evaluation, data stewardship, workflow design, security) even as tools change.

MISSISSIPPI EXAMPLES

- K-12: “I compared AI use in precision agriculture (Delta soybean yield optimization) vs. coastal resilience modeling.”
- Postsecondary: “I aligned my AI project to a career cluster in health sciences, IT, advanced manufacturing, hospitality, or emergency planning.”
- Workforce: “I articulated what is allowed in my sector, like healthcare data rules under HIPAA, agriculture supply-chain auditing, or public-sector record retention for emergency response.”

RESPONSIBLE DATA & EVALUATION PRACTICES

Domain 11. Cyber Security and Safety

DEFINITION

Understanding the security risks, protections, and responsibilities associated with using and interacting with AI systems, including how AI can introduce new vulnerabilities and support cybersecurity practices.

INCLUDES

- Protecting data used by AI systems (confidentiality, integrity, availability)
- Recognizing common AI-related security risks (prompt injection including indirect variants, memory poisoning in agents, data leakage/exfiltration, model and supply-chain poisoning, adversarial inputs)
- Secure and responsible use of AI tools in educational and workplace settings
- Awareness of AI-enabled cyber threats (phishing, social engineering, deepfakes, automated attacks)
- Using AI to support cybersecurity functions (monitoring, detection, analysis)
- Understanding user responsibility and acceptable use in AI-enabled systems
- Emerging risks from AI agents (persistent, multi-step autonomous actions)

WHY THIS DOMAIN EXISTS

AI systems rely on data, networks, and digital infrastructure. Without cybersecurity awareness, AI use can expose sensitive information, increase cyber risk, and erode trust. Cybersecurity literacy ensures AI is used safely and responsibly across education and workforce contexts.

LEARNERS CAN

Learners can operate AI tools safely: don’t paste secrets or sensitive data; recognize prompt injection and memory-poisoning attempts; use only approved environments; follow acceptable-use rules. They identify AI-amplified threat patterns: sophisticated deepfake audio/video, highly personalized phishing, and supply-chain risks in AI tools or plugins. They apply verification habits and defense-in-depth: confirm requests via known channels, validate data sources before agent ingestion, and maintain human oversight on autonomous actions.

MISSISSIPPI EXAMPLES

- K-12: “This prompt is trying to override instructions or poison memory on a hurricane preparedness AI tool; ignore and report.”
- Postsecondary: “This document or external data source an AI agent wants to use may contain indirect prompt injection or poisoned content; validate before processing.”
- Workforce: “Use AI to help draft phishing detection rules for shipbuilding or hospital systems but always test thoroughly and maintain human review of agent actions to prevent supply-chain or memory poisoning risks in critical infrastructure.”

> Domain 11 additionally covers synthetic media directed at identifiable people. Learners at every level should understand that generating or circulating a synthetic depiction of a real person without consent is a harm to that person, that this applies to classmates, teachers, and public figures alike, and that the correct response to encountering such material is to report it. At the workforce levels this extends to organizational response, including preservation of evidence, notification obligations, and support for the person depicted.

Progression Mapping

The following table presents the complete learning progression across all eleven AI skill domains and eight developmental levels. Each cell describes what learners can do at that stage. The levels after high school are entry points rather than steps in one sequence. A learner may move from high school directly into work, into a community college or technical program, or into a university, and may combine them in any order.

Skill Domain	Elementary (K-5)	Middle School (6-8)	High School (9-12)	CC / Technical Programs	University (Fr-So)	University (Jr-Sr)	Early Career (0-4 yrs)	Mid-to-Senior Career & Leadership
1. AI Foundations and Conceptual Understanding	Learners can tell rules-based tools apart from systems that learn from examples by sorting everyday cases such as spell-check versus an image generator; they check one reason drawn from the data before trusting an answer; ready to explain how a system learns from examples next.	Learners can explain machine learning as pattern recognition rather than fixed rules by labeling how a spam filter or a recommendation list works; they test one edge case and record the error; ready to separate generated content from labeled prediction next.	Learners can distinguish systems that generate new content from systems that classify or predict by testing inputs and comparing plausible output against labeled output; they document the error types they expect before relying on a result; ready to carry uncertainty into workflow decisions next.	Learners can trace whether a system's behavior comes from data or from coded rules by following model outputs through a technical project; they probe uncertainty by testing edge cases; ready to evaluate system-level risk next.	Learners can contrast learning-based architectures with conventional software by diagramming data-to-output flow in coursework; they cite documented model limitations rather than assuming reliability; ready to apply risk-management vocabulary next.	Learners can critique generative versus classification approaches in a domain project by analyzing failure modes; they quantify confidence and state what the score does not mean; ready to oversee integration in professional settings next.	Learners can explain how these systems work to colleagues using real project outputs; they require data-source checks before results are acted on; ready to lead governance discussions next.	Learners can audit how well the organization understands the systems it uses by leading cross-team reviews; they update policy as capabilities change; accountable for enterprise-wide conceptual alignment through a recurring literacy audit.
2. Data Literacy and Data Stewardship	Learners can identify where data comes from and judge basic quality by sorting clean versus messy example datasets; they notice which groups are missing; ready to check privacy next.	Learners can evaluate representativeness and sampling by revising a school survey dataset; they label personal information and explain why consent matters; ready to apply redaction routines next.	Learners can document data lineage and run quality checks before a dataset is used by creating annotated records; they practice de-identification to prevent leaks of sensitive information; ready to govern data inside workflows next.	Learners can classify data for permitted use by applying privacy and quality protocols in credential projects; they test outcomes for sampling bias; ready to integrate stewardship into system design next.	Learners can assess how stewardship decisions change model behavior by tracing inputs in course exercises; they cite the regulations that govern consent; ready to handle high-stakes data next.	Learners can design data pipelines with bias audits in capstone projects; they document representativeness gaps and the mitigations chosen; ready to enforce stewardship in a professional role next.	Learners can redact and classify workplace data before it enters any external system by following sector protocols; they log the decisions that were made; ready to advise on organizational data policy next.	Learners can establish enterprise data-governance standards by leading audits; they require impact assessments wherever systemic bias is plausible; accountable for data decisions across teams.
3. Critical Evaluation of AI Outputs	Learners can check a generated answer against classroom facts before using it; they ask where the claim came from rather than accepting a first draft; ready to spot invented details next.	Learners can label fabrications and unsupported claims by comparing generated text to trusted references; they separate confident tone from correctness; ready to triangulate across sources next.	Learners can treat any generated output as a draft by revising it against primary-source evidence; they document the verification steps taken; ready to apply this in complex tasks next.	Learners can critique outputs against professional standards using a verification checklist; they document representativeness gaps across at least three sources; ready to require a stated rationale before acting next.	Learners can evaluate relevance and accuracy by annotating limitations in assignments; they require a stated human rationale rather than deferring to the output; ready to scale review to the system level next.	Learners can lead peer review of machine-generated content using domain-specific rubrics; they name failure modes and quantify residual uncertainty; ready to institutionalize evaluation next.	Learners can verify outputs in daily work by cross-referencing official sources; they counter automation bias by requiring documentation before action; ready to train others next.	Learners can establish organizational verification protocols by auditing high-stakes outputs; they require escalation logs where systemic risk exists; accountable for the enterprise verification culture.

Skill Domain	Elementary (K-5)	Middle School (6-8)	High School (9-12)	CC / Technical Programs	University (Fr-So)	University (Jr-Sr)	Early Career (0-4 yrs)	Mid-to-Senior Career & Leadership
4. Human-AI Interaction (Prompting and Iteration)	Learners can ask a clear question by stating audience and constraints; they compare two versions to see what changed; ready to iterate next.	Learners can refine a request with purpose and tone by revising the result in group tasks; they explain how the wording changed the output; ready to add success criteria next.	Learners can iterate requests with constraints and compare results by documenting tradeoffs in a project; they test variations rather than accepting the first attempt; ready to measure whether output actually improved next.	Learners can write audience-specific instructions for technical tasks and measure whether output improved across iterations; they compare results quantitatively; ready to scale to system integration next.	Learners can design an iterative protocol for course assignments by specifying rubrics in advance; they log comparisons instead of relying on impression; ready to apply the method in research next.	Learners can optimize instructions for complex multidisciplinary projects by leading team iterations; they document the context and constraints that shaped the result; ready to embed the practice in professional systems next.	Learners can work these tools into role-specific workflows under deadline pressure; they compare output against defined metrics rather than assuming time savings; ready to mentor teams next.	Learners can define organizational standards for how staff direct these systems by auditing enterprise use; they revise guidance as limits appear at scale; accountable for governance of the human-machine interface.
5. Algorithmic Thinking and Problem Decomposition	Learners can break a task into steps by sequencing a routine with if-then decisions; they mark which steps a person must do; ready to identify where a tool could help next.	Learners can decompose a problem into a workflow by mapping which steps depend on data and which depend on judgment; they test where a simple script can take over; ready to translate real problems next.	Learners can translate a real-world goal into a decision-point workflow by labeling machine steps and human steps; they require human sign-off on high-stakes decisions rather than automating end to end; ready to design systems next.	Learners can decompose a technical process and build escalation paths into an applied project; they document the tradeoffs of automating each step; ready to integrate oversight next.	Learners can map a problem to an assisted workflow by identifying decision gates in coursework; they justify each human step instead of defaulting to the tool; ready to evaluate architectures next.	Learners can lead decomposition for a capstone system by specifying human checkpoints, including which subtasks can be safely delegated to systems that act without a prompt; they test the workflow for failure points; ready to apply this in leadership next.	Learners can translate job tasks into assisted workflows with built-in review routines, deciding case by case which subtasks are safe to delegate to systems that act on their own; they preserve accountability rather than automating judgment; ready to scale team processes next.	Learners can redesign enterprise workflows by auditing decomposition practice across teams, setting the standard for which subtasks may be delegated to systems that act on their own; they mandate oversight where high-stakes gaps appear; accountable for organizational workflow maturity.
6. AI-Enabled Workflow and System Design	Learners can draw a simple task end to end and mark where a person checks the work; they add a checkpoint rather than letting a tool finish alone; they can tell a single-request tool from one that takes several steps on its own; ready to connect steps into a process next.	Learners can design a basic workflow with an escalation step by documenting a group project; they note where data moves and where it could go wrong; they add a human approval gate to any multi-step sequence; ready to add architecture awareness next.	Learners can build a workflow with defined review gates and test what happens when a step fails; they specify who approves before action is taken; they can prototype a multi-step sequence that hands off to a person at a set threshold, such as triaging flood alerts and escalating severe cases; ready to apply this in professional contexts next.	Learners can integrate an assisted step into an existing process by mapping checkpoints in a credential project; they use lifecycle terms for inputs, outputs, monitoring, and retirement; they document the oversight required for any component that acts without a prompt; ready to manage risk next.	Learners can design a workflow with oversight and evaluation built in; they document system-level gaps rather than assuming the tool handles them; they can contrast a single-call design with one where the system plans and acts across steps; ready to specialize next.	Learners can architect complex assisted systems by leading a team design with monitoring in place; they plan audits for drift; they can design and test systems in which several automated components interact, watching for goal drift and hand-off failure; ready to govern in a career setting next.	Learners can implement role-aware workflows by embedding human review and applying architecture concepts such as inputs, outputs, monitoring, and drift in daily operations; they log incidents when escalation fails; they can deploy and monitor sector-specific automated workflows under supervision, such as an irrigation controller or a quality-inspection routine; ready to advise leadership next.	Learners can establish enterprise system-design standards by auditing integrated risk; they require red-team testing where organizational vulnerabilities are plausible; they govern orchestration of multiple automated components, including persistent autonomy, tool-use limits, and escalation protocols across sectors; accountable for oversight that scales.
7. Ethical Reasoning and Responsible AI Use	Learners can identify fair versus unfair treatment in classroom cases; they practice saying what help they used on a piece of work; ready to apply fairness checks next.	Learners can weigh appropriate against inappropriate use by debating bias scenarios; they explain why transparency protects integrity; ready to document decisions next.	Learners can apply ethical principles by disclosing what assistance they used on assignments; they reference fairness criteria before acting in high-stakes situations; ready to document bias mitigation next.	Learners can evaluate fairness and bias in a technical context and document the mitigation steps taken; they follow the privacy rules of their sector; ready to embed ethics in system design next.	Learners can analyze ethical tradeoffs using transparency frameworks in coursework; they require disclosure rather than treating it as optional; ready to handle complex cases next.	Learners can lead ethical review of a capstone project by balancing accountability and explainability; they conduct impact assessments where bias is plausible; ready to enforce standards in a career setting next.	Learners can apply sector-specific ethics by disclosing what assistance produced a deliverable; they consult published guidelines before acting in uncertain cases; ready to train teams next.	Learners can establish organizational ethical governance by auditing fairness and transparency; they update policy as systemic risks emerge; accountable for enterprise integrity.

Skill Domain	Elementary (K-5)	Middle School (6-8)	High School (9-12)	CC / Technical Programs	University (Fr-So)	University (Jr-Sr)	Early Career (0-4 yrs)	Mid-to-Senior Career & Leadership
8. Human Agency, Judgment, and Restraint	Learners can decide when a tool helps and when to think it through alone by marking human-only steps; they explain the choice rather than defaulting to the tool; ready to preserve accountability next.	Learners can reflect on dependency by tracking when they reached for a tool instead of reasoning; they require their own rationale before a decision counts; ready to apply this at higher stakes next.	Learners can maintain oversight by requiring independent human judgment before any assisted result affects an outcome; they document their own reasoning first; ready to scale restraint next.	Learners can identify restraint points in a technical workflow by specifying which steps must stay manual; they sign off on decisions and own them; ready to embed this in a role next.	Learners can exercise judgment in academic work by stating where assistance would defeat the learning goal; they reflect on overreliance rather than optimizing for speed; ready to apply this professionally next.	Learners can lead protocols that preserve human agency by defining decision triggers, including when to pause, override, or retire a system whose behavior has drifted; they audit usage for dependency risk; ready to mentor next.	Learners can enforce restraint in job tasks by requiring documented human rationale, and they know when to pause, override, or escalate a system whose behavior has drifted; they resist automation bias rather than accepting a polished result; ready to shape team culture next.	Learners can institutionalize human-agency norms by leading policy review, setting the organizational standard for when a drifting system is paused, overridden, or retired; they require training where enterprise dependency is growing; accountable for the organization's judgment culture.
9. Communication and AI Literacy Across Audiences	Learners can explain how these tools work in plain words by presenting examples to classmates; they state the limits instead of overselling; ready to add risks next.	Learners can translate technical concepts for peers by writing plain-language summaries; they adjust the explanation to what the audience already knows; ready to include implications next.	Learners can communicate risk and limitation clearly enough to support an informed choice; they drop jargon and test whether the audience understood; ready to adapt this for professional settings next.	Learners can explain a technical use case to non-technical stakeholders and document the guardrails in place; they gather feedback to find where the explanation failed; ready to translate across teams next.	Learners can deliver cross-disciplinary explanations to mixed audiences; they cite evidence rather than asserting capability; ready to lead literacy sessions next.	Learners can lead literacy sessions for diverse audiences by tailoring risk communication; they revise an explanation after a comprehension check shows it did not land; ready to influence practice in a career setting next.	Learners can brief a team on implications using plain language and concrete examples; they check comprehension rather than assuming it; ready to advise leadership next.	Learners can shape organizational communication strategy by auditing what audiences actually understand and influencing policy accordingly; they measure outcomes to find literacy gaps; accountable for enterprise-wide understanding.
10. Sector and Pathway Awareness	Learners can describe how different jobs use technology differently by comparing school examples; they notice the differences rather than generalizing; ready to map transitions next.	Learners can connect their skills to career clusters through case studies; they discuss what changes between middle and high school; ready to choose a pathway next.	Learners can align a project to pathway expectations by explaining the difference between transfer-oriented and workforce-oriented constraints; they plan for adaptability rather than a single tool; ready to map credentials to a sector next.	Learners can map credentials to sector workflows and document the compliance constraints of a direct-to-workforce role; they adapt use cases to the rules of the sector; ready to deepen theory or application next.	Learners can distinguish transfer-oriented theory from applied pathways by comparing coursework; they identify readiness checkpoints to avoid misalignment; ready to specialize next.	Learners can analyze pathway-specific demands in capstone alignment while accounting for workforce compliance constraints such as data classification, record retention, and auditability; they plan for how requirements will change; ready to enter the workforce next.	Learners can adapt their skills to role-specific constraints by following sector rules; they update their approach based on job feedback rather than assuming transfer; ready to mentor next.	Learners can lead pathway and sector strategy by aligning talent development to demand; they update governance as cross-sector requirements shift; accountable for organizational adaptability.

Skill Domain	Elementary (K-5)	Middle School (6-8)	High School (9-12)	CC / Technical Programs	University (Fr-So)	University (Jr-Sr)	Early Career (0-4 yrs)	Mid-to-Senior Career & Leadership
11. Cyber Security and Safety	Learners can protect basic information by never entering personal details into an online tool; they ask permission before sharing anything about themselves or others; they understand that some tools keep working on their own and have to be watched; ready to recognize suspicious requests next.	Learners can spot attempts to trick a tool into ignoring its instructions and refuse to pass them along; they verify a suspicious image, voice, or video against a source they already trust; they understand that some tools chain several actions together without asking again; ready to follow formal rules next.	Learners can apply acceptable-use policy by classifying information before entering it anywhere and confirming urgent requests through a separate channel; they report manipulated media rather than forwarding it; they practice safe interaction with tools that take several steps on their own; ready to integrate secure handling into technical work next.	Learners can secure tools in a technical setting by redacting sensitive information, detecting manipulated inputs, and escalating instruction-override attempts per protocol; they confirm identity before responding to an urgent request; they identify risk from tools that retain state and act across sessions; ready to support detection work next.	Learners can protect confidentiality and integrity by documenting classification routines and testing for misuse in coursework; they work only in approved environments to prevent leakage; they analyze risks specific to systems that act on their own, including excessive permissions and unchecked escalation; ready to implement controls next.	Learners can implement organizational controls by designing verification routines and monitoring for manipulated media and corrupted inputs, validating any automated detection before relying on it; they audit compliance to find gaps in user responsibility; they evaluate the security of systems that operate across multiple steps; ready to govern next.	Learners can enforce secure practice by redacting workplace information, responding to instruction-override and stored-context manipulation per policy, and verifying automated threats such as scaled intrusion attempts; they classify sensitivity and log escalations per protocol; they monitor and report risk from deployed systems that retain memory and take chained actions; ready to advise leadership next.	Learners can establish enterprise security governance by auditing these risks, including corrupted stored context, compromised supply chains, and threats from systems that act without direction; they mandate protection training, require human oversight, and require incident reporting; they set standards for excessive permissions, tool abuse, and long-running autonomy across the organization and its sectors; accountable for the safety culture.

K-12 (Foundational Awareness & Practice) **Postsecondary (Applied & Technical Fluency)**

Workforce (Professional Application & Accountability)

A standalone landscape version of the Progression Mapping is available from MAIN for printing and for use in scope and sequence work.

PART FOUR

Alignment

Alignment and Reference Points

The AI Learning Progression was developed with reference to the instruments below. Mississippi is governed by none of them except where state or federal law independently applies. They are cited so that the origin of each design choice can be traced and compared, not to assert that any external body has authority over this framework.

- Progression emphasizes human-centered design and durable skills (critical thinking, judgment, agency) per the OECD Digital Education Outlook 2026 and the OECD/EC AI Literacy Framework.
- Risk-based language (bias, transparency, accountability, human oversight) mirrors the NIST AI Risk Management Framework 1.0 with the July 2024 Generative AI Profile, currently under revision. NIST is the operative reference for this framework. The risk-tiering vocabulary of the EU AI Act is noted only because Mississippi employers with European operations are subject to it and their workforce needs reflect it; its obligations were rescheduled by Regulation (EU) 2026/1744 to 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems, with Article 50 transparency duties applying from 2 August 2026. Nothing in this framework derives authority from it.
- Equity, non-discrimination, and the protection of learner rights draw on the U.S. Department of Education's Office for Civil Rights guidance on avoiding the discriminatory use of artificial intelligence (November 2024, revised January 2025). That document has since been rescinded by the Department and remains posted for historical reference only. It is cited here as a source of this framework's design reasoning, not as current guidance. The underlying statutory obligations are unchanged: Title VI, Title IX, Section 504, and the Americans with Disabilities Act continue to apply to Mississippi institutions regardless of the tools they use. Comparable international treatments, including UNESCO's *Guidance for generative AI in education and research*, informed the disclosure and appropriate-use distinctions.
- Cybersecurity progression (data leakage prevention, prompt injection including indirect variants, deepfakes, acceptable use, and agent-specific risk) draws on the NIST Adversarial Machine Learning taxonomy (AI 100-2 E2025, final March 2025), the OWASP Top 10 for Agentic Applications 2026 (December 2025), the NIST AI Agent Standards Initiative (February 2026), and MITRE ATLAS.
- Sector- and pathway-aware deepening supports the OECD Skills Outlook 2025 focus on adaptability across education-to-workforce transitions.
- The K-12 foundational stage, with progressive development across grade bands and increasing judgment in postsecondary and workforce contexts, follows the U.S. Department of Labor AI Literacy Framework (2026) and ETS K-12 AI literacy progression models.
- Cross-audience communication and ethical reasoning reflect the communication and disclosure expectations already present in Mississippi practice, including the plain-language and human-oversight duties in ITS PSG-100-03 and the ethics standards of the 2026 Mississippi CCR Standards for Computer Science.
- Overall handoff visibility and restraint against overreliance draw from the OECD and European Commission framework *Empowering Learners for the Age of AI* (final, 17 June 2026) and its emphasis on preserving human responsibility.

Crosswalk 1. U.S. Department of Labor AI Literacy Framework

Training and Employment Notice 07-25, 13 February 2026.

DOL content area	Corresponding domains
Understand AI Principles	Domain 1, AI Foundations; Domain 2, Data Literacy
Explore AI Uses	Domain 10, Sector and Pathway Awareness; Domain 5, Algorithmic Thinking
Direct AI Effectively	Domain 4, Human-AI Interaction
Evaluate AI Outputs	Domain 3, Critical Evaluation
Use AI Responsibly	Domain 7, Ethical Reasoning; Domain 8, Human Agency, Judgment, and Restraint
Extending beyond the federal content areas	Domain 6, Workflow and System Design; Domain 9, Communication Across Audiences; Domain 11, Cyber Security and Safety

Crosswalk 2. OECD and European Commission AI Literacy Framework

Empowering Learners for the Age of AI, final, 17 June 2026.

AILit domain	Corresponding domains
Engage with AI	Domains 1, 3, 4
Create with AI	Domains 4, 5, 6
Manage AI	Domains 2, 6, 8, 11
Shape AI	Domains 7, 9, 10

The AILit framework treats ethical principles as cross-cutting rather than as a domain of their own, distributing them through the knowledge, skills, and attitudes in all four domains. Domain 7 therefore corresponds to ethical content throughout AILit, not only to Shape AI.

Crosswalk 3. Mississippi Instruments

Instrument	Relationship
Miss. Code § 37-13-215 (SB 2294, 2026)	Domains 1, 3, 7, 8. The Department of Education’s 2026 standards are the operative content interpretation; this framework carries the competencies beyond grade 12.

Instrument	Relationship
Miss. Code § 37-13-207 (HB 633, 2021)	Domains 1, 2, 5, 11. Statutory basis for K-12 coverage predating this framework.
Miss. Code Title 1, Ch. 3 (HB 1723, 2026)	Supplies the definition of artificial intelligence used throughout.
2026 MS CCR Standards for Computer Science	All eleven domains. See Crosswalk 4.
Exploring Computer Science CTE framework (2025)	Domains 1, 2, 11. An existing course-level delivery vehicle at the High School level.
ITS PSG-100-03 (2025)	Domains 8 and 11, and the High-Stakes Human-Review Triggers. The public-sector expression of this framework's oversight posture.
SB 2524 (2026), Chapter 483	No competency domain. Supplies the statutory definitions of postsecondary attainment and credential of value, and the reporting apparatus through which credentials are counted.

Crosswalk 4. 2026 MS CCR Standards for Computer Science

Domain	Where it appears in the standards	K-12 status
1. AI Foundations	ALG Machine Learning subconcept (EK.ALG.05 through E6.ALG.09; FS.ALG.13, .16, .17) · CAS Humans and Computing (E5.CAS.03, FS.CAS.06) · S1.AIN.01-03	Fully covered
2. Data Literacy and Stewardship	DAA concept in full · FS.ALG.16 · S1.AIN.06, S1.AIN.08 · S2.AIN.21-22	Fully covered
3. Critical Evaluation of Outputs	FS.ALG.15 · PRO Reading and Documenting Code; Testing and Refining Code · S1.AIN.11 · S1.SWD.06	Fully covered
4. Human-AI Interaction	PRO Program Development · S1.SWD.07 · S2.SWD.16	Partial. This framework adds an explicit prompting and iteration progression.
5. Algorithmic Thinking	ALG concept · Computational Thinking pillar, practices 6-9 · E6.ALG.09	Fully covered
6. Workflow and System Design	SAS concept · Human-Centered Design pillar · S1.AIN.13 and S2.AIN.27, design and optimize AI agents	Covered, including agents. This framework adds oversight design at the workforce levels.
7. Ethical Reasoning	Ethics and Social Responsibility pillar · CAS Impacts subconcepts · S1.AIN.09-12, .14, .16 · S2.AIN.23, .25	Extensively covered

Domain	Where it appears in the standards	K-12 status
8. Human Agency and Restraint	E4.CAS.03 · FS.CAS.02 · S1.AIN.10 · Human-Centered Design practice 12	Partial. This framework carries restraint into professional practice.
9. Communication Across Audiences	Inclusive Collaboration pillar, practice 3 · S2.AIN.26	Covered
10. Sector and Pathway Awareness	CAS Career Exploration subconcept (E2.CAS.03 through FS.CAS.07) · E6.CAS.07	Covered at K-12; this framework extends it to postsecondary and workforce.
11. Cyber Security and Safety	SAS concept · CYB specialty · S2.CYB.27	Extensively covered

Mississippi has strong K-12 coverage across all eleven domains through an instrument that carries implementation force. This framework adds an explicit prompting and iteration progression, restraint carried into professional practice, and the four levels beyond grade 12 that no Mississippi instrument reaches. Read together, the two documents describe the full path.

PART FIVE

Governance and Use

High-Stakes Human-Review Triggers

Institutions adopting this framework should require human review of the following categories of AI output before action is taken. The categories are the same at every level of the progression:

- Any AI output affecting grades, academic assessment, or credential decisions
- Any AI output influencing financial allocations, budgeting, or compensation
- Any AI output related to safety-critical operations, health, or physical risk
- Any AI output recommending discipline, hiring, termination, or employment actions
- Any AI output generating or interpreting legal, regulatory, or compliance content
- Any AI output involving sensitive, confidential, protected, or personally identifiable data

Agent-specific triggers. Where a system acts autonomously, the following are added, whether or not the output falls into a category above:

- Any autonomous action that writes to, alters, or deletes a record in a system of record
- Any system granted persistent memory across sessions, where stored context could carry forward manipulated or erroneous content
- Any system granted tool access, credentials, or the ability to initiate external transactions
- Any hand-off between automated components in a safety-critical, regulated, or student-facing process
- Any system operating without a defined stop condition or without a logged escalation path to a named human owner
- Any institutional response to synthetic media depicting an identifiable person, which should not be delegated to automated detection alone

Outside these categories, and outside any further limits an institution sets for itself, exploring, testing, and building with these systems is how capability develops.

Measurement: What Can and Cannot Be Claimed

No validated, at-scale assessment of student or workforce AI literacy suitable for state adoption currently exists. The Meta AI Literacy Scale is a self-report instrument. The AI Competency Objective Scale is a research instrument. The ETS research report cited in this framework proposes a framework, a hypothesized progression, and task design principles, and explicitly does not produce or validate an instrument. Work is under way elsewhere. ETS released Adapt AI on 11 February 2026, an assessment of artificial intelligence readiness for K-12 educators rather than for students. The OECD is developing a Media and Artificial Intelligence Literacy assessment for PISA 2029. California launched an AI literacy micro-credential in July 2026 for an initial 500 participants, awarded on completion. No state adoption of an AI literacy assessment was identified in preparing this edition. The Council will not adopt or construct a proficiency

rubric it cannot defend, and the instruments that do exist were built elsewhere for other purposes. Mississippi has the institutional range to help close the gap: research universities, community colleges operating at scale, K-12 districts implementing new computer science standards, and public agencies running real deployments. Instrument development and validation is work this state is positioned to contribute to rather than only wait for.

The state already collects the first four indicators below. The last two would require a source to be selected or a reporting practice to be established. The first indicator sits inside a single instrument rather than two: the 2024-2028 Mississippi Workforce Innovation and Opportunity Act State Plan is a Combined State Plan, and Perkins V career and technical education is a combined-plan partner program alongside the WIOA core programs.

Indicator	Source and cadence
Credential attainment	Perkins V and WIOA reporting, annual
Wage and employment outcomes	Unemployment insurance wage-record match
Course completions	Completions in courses qualifying under Miss. Code § 37-13-215, as the requirement phases in
Employer adoption	Census Business Trends and Outlook Survey, by sector and firm size, biweekly
Employer demand	Job-posting analysis of AI skill share, benchmarked nationally. Source not yet selected
Educator capacity	Participation and follow-through in state and partner professional learning. Reporting practice not yet established

These measure system activity and labor market conditions, not individual learner proficiency, and that limit should be stated wherever they are reported. On the employer-demand row, the State Workforce Investment Board has already directed work that could supply it: the fall employer meetings described above are to capture competency gaps employers identify, and the Board asked that a partnership with the Coleridge Initiative and the National Labor Exchange be examined for integrating real-time labor market data into program evaluation. Neither is a committed source for this indicator.

Governance, Ownership, and Review Cadence

Ownership

The AI Workforce Readiness Council, chaired by Dr. Kollin Napier, Chief Artificial Intelligence Officer of MAIN, and operating as a subcommittee of the State Workforce Investment Board, maintains this document. MAIN serves as the drafting body, and AccelerateMS hosts the published edition. Neither the Council nor MAIN holds regulatory authority over the Mississippi Department of Education, the Mississippi Community College Board, or the Institutions of Higher Learning. Adoption by those bodies is voluntary and at their discretion.

Review cycle

This framework is reviewed on a fixed cycle each January, timed to follow the OECD and Credential Engine publication cycles and to precede Mississippi legislative committee deadlines. Each edition carries its publication date and a record of what changed.

Revision out of cycle is triggered by any of the following:

- Enactment of Mississippi artificial intelligence or computer science legislation
- Adoption or revision of Mississippi Department of Education academic or career and technical standards touching computer science or artificial intelligence
- A new federal executive order, or a new U.S. Department of Labor or Department of Education issuance addressing AI literacy or workforce
- A change to EU AI Act application dates
- Publication of a revision to the NIST AI Risk Management Framework
- Rescission, withdrawal, or replacement of a federal guidance document cited in this framework

Source refresh cadence

Anchor source	Refresh
Bureau of Labor Statistics Employment Projections	Annual, late summer
Census Business Trends and Outlook Survey	Biweekly
Credential Engine, Counting Credentials	Periodic; most recent edition December 2025
Stanford AI Index	Annual, April
Department of Labor and Department of Education advisories	Continuous
Mississippi legislative session	Annual, January through April

Regional alignment

Mississippi is a member of the Southern Regional Education Board, whose Commission on Artificial Intelligence in Education issued twelve recommendations between January and November 2025. This framework responds to Recommendation 1, establish state AI networks, and partially to Recommendation 9, align across K-12, postsecondary, and workforce. Recommendation 11, postsecondary adoption playbooks, and Recommendation 12, evaluation frameworks, remain open. The Measurement section above begins a response to Recommendation 12.

Succession

The Mississippi AI Regulation Task Force dissolves 31 December 2027. Because it is the current legislative channel for artificial intelligence policy recommendations and identifies MAIN as the statewide implementation and measurement backbone, a successor reporting arrangement should be identified before the 2027 review.

Reference List

Every source below links to the issuing body's own published record, except for two legal citations where no official free full text is published and a recognized public legal repository is used instead. Links were verified in August 2026 and are re-verified at each annual review.

Mississippi statutes

Miss. Code § 37-13-201 et seq. (Mississippi Computer Science and Cyber Education Equality Act, HB 633, 2021; Laws 2021, ch. 389); codified text at § 37-13-207. Miss. Code § 37-13-215 (Mississippi Future Innovators Act, SB 2294, 2026; Chapter 479, Laws of 2026). HB 1723 (2026), definition of artificial intelligence (Chapter 311, Laws of 2026). SB 2524 (2026), Mississippi Postsecondary Attainment Council (Chapter 483, Laws of 2026). SB 2426 (2025), Mississippi Artificial Intelligence Regulation Task Force. Federal definition incorporated by HB 1723 at 15 U.S.C. § 9401(3).

Mississippi standards, policy, and reports

2026 Mississippi College- and Career-Readiness Standards for Computer Science, Mississippi Department of Education, revised February 2026, effective 2026-2027. State Board of Education adoption: agenda and minutes, 29 April 2026 (minutes) and 18 June 2026 (minutes). Exploring Computer Science CTE curriculum framework, approved 16 January 2025. Acceptable Use Policy for Artificial Intelligence, PSG-100-03, Mississippi Department of Information Technology Services, 25 November 2025 (program page). Mississippi Artificial Intelligence Regulation Task Force, Report One to the Legislature, 13 January 2026. State Workforce Investment Board, Approved Recommendations List, 2025. State Workforce Investment Board, Approved Credential List for 2025-2026, 20 February 2025. 2024-2028 Mississippi Workforce Innovation and Opportunity Act State Plan (Combined State Plan), AccelerateMS. State Workforce Investment Board, Recommendations, 4 September 2025. Institutions of Higher Learning, Five-Year Strategic Plan 2026-2030. Mississippi AI Talent Accelerator Program announcement, 12 June 2025. Executive Order 1584, 8 January 2025. Mississippi Statewide AI Framework, April 2026.

Computer Science Teachers Association

CSTA PK-12 Computer Science Standards, Revised 2026.

U.S. Department of Labor

Training and Employment Guidance Letter 03-25, "Encouraging the Use of Workforce Innovation and Opportunity Act Funding to Help Youth and Adults Develop Artificial Intelligence Skills," 26 August 2025. Training and Employment Notice 07-25, "The U.S. Department of Labor's Artificial Intelligence Literacy Framework," 13 February 2026. AI in Registered Apprenticeship Innovation Portal. Department announcement of the portal, 29 April 2026.

U.S. Department of Education

Final Priority and Definitions: Secretary's Supplemental Priority and Definitions on Advancing Artificial Intelligence in Education, 91 FR 18774, 13 April 2026. Office for Civil Rights, "Avoiding the Discriminatory Use of Artificial Intelligence", November 2024 (revised January 2025); rescinded by the Department and posted for historical reference. *Guidance on the Use of Federal Grant Funds to Improve Education Outcomes Using Artificial Intelligence*, Office of the Secretary, 22 July 2025.

The White House

Executive Order 14277, "Advancing Artificial Intelligence Education for American Youth," 23 April 2025, 90 FR 17519. *Winning the Race: America's AI Action Plan*, July 2025. *National Policy Framework for Artificial Intelligence: Legislative Recommendations*, 20 March 2026. Executive Order 14365, "Ensuring a National Policy Framework for Artificial Intelligence," 11 December 2025, 90 FR 58499. *President Trump's Cyber Strategy for America*, March 2026. Executive Order 14409, "Promoting Advanced Artificial Intelligence Innovation and Security," 2 June 2026, 91 FR 34565. *Science: A New Golden Age*, Office of Science and Technology Policy, July 2026.

Federal programs

Workforce Pell Grant final rule, 91 FR, 19 May 2026, effective 1 July 2026.

OECD and European Commission

OECD Digital Education Outlook 2026, 19 January 2026. *OECD Skills Outlook 2025*, 9 December 2025. *Empowering Learners for the Age of AI: An AI Literacy Framework for Primary and Secondary Education*, 17 June 2026.

UNESCO

Guidance for generative AI in education and research, 7 September 2023. *AI and education: Protecting the rights of learners*, 4 September 2025.

NIST

Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, January 2023, under revision; Generative AI Profile, NIST AI 600-1, July 2024. NIST AI 100-2 E2025, *Adversarial Machine Learning*, final March 2025. AI Agent Standards Initiative, announced 17 February 2026.

European Union

Regulation (EU) 2024/1689 (Artificial Intelligence Act), as amended by Regulation (EU) 2026/1744, in force 27 July 2026.

Security and agent risk

OWASP GenAI Security Project, *OWASP Top 10 for Agentic Applications 2026*, December 2025. MITRE ATLAS.

Research, assessment, and credentials

Chakraburty, S., Ober, T. M., & Liu, L. (2025), *Preparing K-12 Students With AI Literacy* (ETS Research Report RR-25-14), November 2025. Southern Regional Education Board, Commission on Artificial Intelligence in Education recommendations, 2025. Credential Engine, *Counting Credentials 2025*, 9 December 2025. ETS, Adapt AI, 11 February 2026. OECD, PISA 2029 Media and Artificial Intelligence Literacy assessment. California Labor and Workforce Development Agency, AI-Ready California micro-credential, 14 July 2026.

Recurring data sources

U.S. Bureau of Labor Statistics, Employment Projections. U.S. Census Bureau, Business Trends and Outlook Survey.

Authorship and Council Participation

Prepared by the Mississippi Artificial Intelligence Network for the AI Workforce Readiness Council, chaired by Dr. Kollin Napier, Chief Artificial Intelligence Officer of MAIN.

The Council brings together the organizations and institutions below. Listing indicates participation in the Council, not institutional review or endorsement of this framework.

- AccelerateMS
- Alcorn State University
- Amazon Web Services (AWS)
- Belhaven University
- C Spire
- Center for Cyber Education
- Copiah-Lincoln Community College
- Hinds Community College
- Institutions of Higher Learning
- Jackson State University
- Millsaps College
- Mississippi Artificial Intelligence Network (MAIN)
- Mississippi College
- Mississippi Community College Board
- Mississippi Department of Education
- Mississippi Department of Information Technology Services
- Mississippi Gulf Coast Community College
- Mississippi State University
- NVIDIA
- Oxford-Lafayette Economic Development Foundation
- Pearl Public School District
- Tougaloo College
- University of Mississippi
- University of Mississippi Medical Center
- University of Southern Mississippi

Mississippi Statewide AI Framework, August 2026 Edition. A publication of the AI Workforce Readiness Council, a subcommittee of the State Workforce Investment Board. Produced by the Mississippi Artificial Intelligence Network (MAIN) and hosted by AccelerateMS, the state's Office of Workforce Development. Supersedes the April 2026 edition.

mainms.org · mainms.org/ai-workforce-readiness-council · August 2026 Edition