MISSISSIPPI STATEWIDE AI FRAMEWORK

Statewide Priorities and Learning Progression

K-12 → Postsecondary → Workforce

A Publication of the AI Workforce Readiness Council

Produced and Hosted by the Mississippi Artificial Intelligence Network (MAIN)

April 2026

https://mainms.org/ai-workforce-readiness-council/

Foreword

A Letter from the Chair of the AI Workforce Readiness Council

Mississippi is stepping boldly into the age of artificial intelligence, not by chasing trends, but by building on the strengths that have long defined our state: resilience, collaboration, and an unwavering commitment to our people.

Through the Mississippi Artificial Intelligence Network (MAIN), Mississippi's coordinated statewide AI initiative and the first of its kind in the nation, our state is building a strong foundation for an AI-enabled future. MAIN brings together educators, industry partners, researchers, and policymakers around a shared belief that every Mississippian should have the opportunity to develop the knowledge, skills, and confidence needed to succeed in a rapidly changing world.

At its core, this statewide framework reflects a clear conviction: the purpose of AI is not to replace human judgment, but to strengthen it. Mississippi's learners and workers deserve more than exposure to tools. They deserve the critical thinking, ethical grounding, and adaptability required to lead in a world where AI is increasingly embedded across every sector.

This work is the product of meaningful collaboration across education, industry, and public service. I am deeply grateful to the AI Workforce Readiness Council, and the many partners whose expertise and commitment helped make this possible.

Mississippi has an opportunity not only to respond to this moment, but to help shape it. This statewide framework represents an important step forward in that effort.

Together, we are building something that matters.

With respect and purpose,

Kollin Napier

Dr. Kollin Napier
Chair, AI Workforce Readiness Council
Director, Mississippi Artificial Intelligence Network (MAIN)
https://mainms.org/

Table of Contents

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 not to replicate another region's model, but to build a distinctly Mississippi approach that leverages 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.

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.
- 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.

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.

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.

Together, these priorities set the direction for everything that follows. Part Two presents the AI Learning Progression: its purpose and scope, the Guiding Principles that shape every section, the AI Literacy Pillars that organize the skills into a clear structure, and the AI Learning Progression Domains and Progression Mapping that translate these priorities into concrete skills, stage-specific descriptors, and a continuous pathway from K-12 through postsecondary and into the workforce.

PART TWO

AI Learning Progression

The AI Learning Progression maps the agreed AI skill domains into a comprehensive 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)

The AI Workforce Readiness Council was established via the Mississippi AI Talent Accelerator Program (MAI-TAP), announced by Governor Tate Reeves on June 12, 2025. The Council brings together educators, researchers, industry leaders, and public-sector representatives to ensure Mississippi's AI talent pipeline is aligned, equitable, and forward-looking.

This is a living document. Because AI is advancing so rapidly, this framework will require consistent review, updates, and refinement as needed. It is not legislation, a requirement, or a mandate. Rather, it serves as a strategic point of alignment and leadership-oriented guide for educators, postsecondary institutions, and industry partners across Mississippi.

This Document:

- Translates each AI Skill Domain into stage-specific performance descriptors.
- Shows how expectations deepen from foundational understanding to applied use to professional accountability.
- Supports alignment discussions by making handoff expectations visible across education and workforce transitions.
- Provides a reviewable matrix to identify gaps, redundancy, and unclear progression language.
- Defines what must be addressed and establishes shared language and scope.
- Provides the input layer for sequencing and readiness checkpoints.

This Document Does Not:

- Set state standards, required assessments, or proficiency cut scores.
- Prescribe courses, curriculum, lesson plans, tools, platforms, or vendors.
- Establish policy, endorsements, compliance requirements, or accountability structures.
- Replace detailed program design; it is a shared progression reference for alignment work.
- Define curriculum standards or specify courses, tools, platforms, or vendors.
- Set policy, accountability, or endorsement requirements.

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.

AI Literacy Pillars

To make the 11 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 + cross-cutting lens

AI Learning Progression Domains

The following eleven domains define the full set of AI-related skills, concepts, and capability areas that must be accounted for when building a K-12 → Postsecondary → Workforce AI learning progression in Mississippi.

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

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?”

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
- Responsible data handling in education and work contexts

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?”

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.”

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."

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.”

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

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: "AI drafts parent letters about school hurricane closure; staff edits and approves before sending."

Postsecondary: "AI triages support tickets in a coastal emergency management simulation; high-risk escalated to a human."

Workforce: "AI triages support tickets in coastal emergency management; high-risk escalated to a human with agentic hand-off logged."

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

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: “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?”

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."

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

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."

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, 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."

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."

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, creating a continuous pathway from elementary school through senior career leadership.

AI 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 distinguish rules-based tools from learning-based AI by sorting everyday examples such as spell-check versus image generators; they avoid magical thinking by documenting one data-driven reason before trusting outputs; ready to compare AI types next.	Learners can explain machine learning as pattern recognition versus fixed rules by labeling spam filters or recommendation lists; they recognize limits by testing one edge case and recording the error; ready to differentiate generative from predictive outputs next.	Learners can differentiate generative AI from predictive systems by testing prompts and observing plausible versus labeled outputs; they avoid overconfidence by documenting expected errors; ready to integrate uncertainty into workflows next.	Learners can map AI system behaviors to data-driven versus rule-based designs by tracing simple model outputs in technical projects; they recognize uncertainty by testing edge cases; ready to evaluate system-level risks next.	Learners can contrast AI architectures with traditional software by diagramming data-to-output flows in introductory courses; they avoid blind trust by citing model limitations in assignments; ready to apply risk-management vocabulary next.	Learners can critique generative versus classification systems in domain projects by analyzing failure modes; they recognize uncertainty by quantifying confidence scores; ready to oversee integration in professional contexts next.	Learners can explain AI foundations to teams by referencing real project outputs; they avoid magical thinking by requiring data-source checks; ready to lead governance discussions next.	Learners can audit organizational AI understanding by leading cross-team reviews; they recognize persistent limits by updating policies for new capabilities; accountable for enterprise-wide conceptual alignment by leading annual AI literacy audits.
2. Data Literacy and Data Stewardship	Learners can identify data sources and quality by sorting clean versus messy example datasets in class activities; they avoid bias by noticing missing groups; ready to check privacy next.	Learners can evaluate representativeness and sampling by revising a school survey dataset; they recognize consent issues by labeling personal information; ready to apply redaction routines next.	Learners can document data lineage and quality checks before AI tasks by creating annotated datasets; they avoid sensitive-data leaks by practicing de-identification; ready to govern data in workflows next.	Learners can classify data for AI use by applying privacy and quality protocols in credential projects; they recognize sampling bias by testing outcomes; ready to integrate into system design next.	Learners can assess data stewardship impacts on model behavior by tracing inputs in course exercises; they avoid consent violations by citing regulations; ready to manage high-stakes data next.	Learners can design data pipelines with bias audits for capstone projects; they recognize representativeness gaps by documenting mitigations; ready to enforce stewardship in career roles next.	Learners can redact and classify workplace data before AI input by following sector protocols; they avoid leakage by logging decisions; ready to advise on organizational data policy next.	Learners can establish enterprise data-governance standards by leading audits; they recognize systemic bias by requiring impact assessments; accountable for data decisions across teams.
3. Critical Evaluation of AI Outputs	Learners can verify AI answers against classroom facts by cross-checking simple outputs; they avoid accepting drafts as final by asking “show your sources”; ready to spot hallucinations next.	Learners can label hallucinations and unsupported claims by comparing AI text to trusted references; they recognize confidence versus correctness by flagging polished errors; ready to triangulate sources next.	Learners can treat outputs as drafts by revising AI content with primary-source evidence in projects; they avoid misuse by documenting verification steps; ready to apply in complex tasks next.	Learners can critique AI outputs against professional standards by running verification checklists; they recognize bias by documenting representativeness gaps across three sources; ready to embed in workflows next.	Learners can evaluate outputs for relevance and accuracy by annotating limitations in assignments; they avoid overreliance by requiring human rationale; ready to scale to system-level review next.	Learners can lead peer reviews of AI-generated content by applying domain-specific rubrics; they recognize failure modes by quantifying uncertainties; ready to institutionalize evaluation next.	Learners can verify AI outputs in daily tasks by cross-referencing official sources; they avoid automation bias by mandating documentation; ready to train others next.	Learners can establish organizational evaluation protocols by auditing high-stakes outputs; they recognize systemic risks by requiring escalation logs; accountable for enterprise verification culture.
4. Human-AI Interaction (Prompting and Iteration)	Learners can ask clear questions by adding audience and constraints to simple prompts; they avoid vague results by comparing two versions; ready to iterate next.	Learners can refine prompts with purpose and tone by revising outputs in group tasks; they recognize prompt influence by explaining changes; ready to add success criteria next.	Learners can iterate prompts with constraints and compare results by documenting tradeoffs in projects; they avoid poor outputs by testing variations; ready to map human-in-the-loop checkpoints in end-to-end workflows next.	Learners can craft audience-specific prompts for technical tasks by iterating in credential simulations with effectiveness measurement; they recognize effectiveness by comparing quantitative improvements; ready to scale to system integration next.	Learners can design iterative prompting protocols for course assignments by specifying rubrics; they avoid suboptimal results by logging comparisons; ready to apply in research next.	Learners can optimize prompts for complex multidisciplinary projects by leading team iterations; they recognize constraints by documenting context; ready to embed in professional systems next.	Learners can prompt AI within role-specific workflows by iterating under deadlines; they avoid inefficiency by comparing outputs against metrics; ready to mentor teams next.	Learners can define organizational prompting standards by auditing enterprise use; they recognize scalability limits by updating guidelines; accountable for human-AI interface governance.

5. Algorithmic Thinking and Problem Decomposition	Learners can break tasks into steps by sequencing daily routines with if-then decisions; they avoid skipping human parts by marking review points; ready to identify AI opportunities next.	Learners can decompose problems into workflows by mapping “if data then AI” in school projects; they recognize automatable steps by testing simple scripts; ready to translate real problems next.	Learners can translate real-world goals into decision-point workflows by labeling human versus AI tasks; they avoid over-automation by mandating human sign-off on high-stakes decisions; ready to design systems next.	Learners can decompose technical processes by creating escalation paths in applied projects; they recognize appropriate AI use by documenting tradeoffs; ready to integrate oversight next.	Learners can map problems to AI-assisted workflows by identifying decision gates in coursework; they avoid automation bias by justifying human steps; ready to evaluate architectures next.	Learners can lead problem decomposition for capstone systems by specifying human-in-the-loop checkpoints, including decomposing problems for safe delegation to agentic systems; they recognize workflow risks by testing; ready to apply in leadership next.	Learners can translate job tasks into AI-supported flows by building review routines, including decomposing problems for safe delegation to agentic systems; they avoid inappropriate automation by preserving accountability; ready to scale team processes next.	Learners can redesign enterprise workflows by auditing decomposition practices, including decomposing problems for safe delegation to agentic systems; they recognize high-stakes gaps by mandating oversight; accountable for organizational algorithmic maturity.
6. AI-Enabled Workflow and System Design	Learners can map simple AI tasks with human review by drawing end-to-end steps; they avoid standalone use by adding checkpoints; they can distinguish single-prompt tools from simple multi-step agent behaviors (e.g., “the AI researched then drafted”); ready to integrate processes next.	Learners can design basic workflows with escalation by documenting school-group projects; they recognize system risks by noting data flow; they can add human approval gates to simple agentic sequences; ready to add architecture awareness next.	Learners can build human-in-the-loop workflows by specifying review gates in projects; they avoid integration failures by testing fallbacks; they can prototype basic multi-step agentic workflows (e.g., AI triages coastal flood alerts then escalates to human); ready to apply in professional contexts next.	Learners can integrate AI into existing processes by mapping checkpoints for credential projects; they recognize architecture elements by using lifecycle terms; they can incorporate agentic components while documenting oversight requirements; ready to manage risks next.	Learners can design workflows with oversight by incorporating evaluation in coursework; they avoid system-level gaps by documenting risks; they can contrast traditional vs. agentic architectures; ready to specialize next.	Learners can architect complex AI-enabled systems by leading team designs with monitoring; they recognize drift by planning audits; they can design and test multi-agent systems with monitoring for goal drift and inter-agent risks; ready to govern in careers next.	Learners can implement role-aware workflows by embedding human review and applying basic architecture concepts such as inputs/outputs, monitoring, and drift in daily operations; they avoid escalation failures by logging incidents; they can deploy and monitor sector-specific agentic workflows (e.g., precision-ag irrigation agent or shipbuilding quality agent) under supervision; ready to advise leadership next.	Learners can establish enterprise system-design standards by auditing integrated risks; they recognize organizational vulnerabilities by requiring red-team testing; they lead governance of multi-agent orchestration, including persistent autonomy, tool-use policies, and escalation protocols across Mississippi sectors; accountable for scalable oversight.
7. Ethical Reasoning and Responsible AI Use	Learners can identify fair versus unfair AI examples by discussing classroom cases; they avoid inappropriate use by practicing disclosure statements; ready to apply fairness checks next.	Learners can weigh appropriate versus inappropriate uses by debating bias scenarios; they recognize integrity risks by citing transparency needs; ready to document decisions next.	Learners can apply ethical principles by disclosing AI assistance in assignments; they avoid high-stakes misuse by referencing fairness criteria; ready to integrate accountability next.	Learners can evaluate fairness and bias in technical contexts by documenting mitigation steps; they recognize privacy norms by following sector rules; ready to embed in systems next.	Learners can analyze ethical tradeoffs by applying transparency frameworks in coursework; they avoid integrity violations by requiring disclosure; ready to handle complex cases next.	Learners can lead ethical reviews for capstone projects by balancing accountability and explainability; they recognize bias impacts by conducting impact assessments; ready to enforce in careers next.	Learners can apply sector-specific ethics by disclosing AI use in deliverables; they avoid inappropriate cases by consulting guidelines; ready to train teams next.	Learners can establish organizational ethical governance by auditing fairness and transparency; they recognize systemic risks by updating policies; accountable for enterprise integrity.
8. Human Agency, Judgment, and Restraint	Learners can decide when AI helps versus when to think alone by marking “human-only” steps; they avoid automation bias by explaining choices; ready to preserve accountability next.	Learners can reflect on dependency by journaling AI overuse in tasks; they recognize human-only decisions by requiring personal rationale; ready to apply in higher stakes next.	Learners can maintain oversight by requiring human judgment before AI-influenced outcomes; they avoid bias by documenting independent reasoning; ready to scale restraint next.	Learners can identify restraint points in technical workflows by specifying no-AI zones; they recognize accountability by signing off decisions; ready to embed in roles next.	Learners can exercise judgment in academic contexts by justifying AI limits; they avoid overreliance by reflecting on learning goals; ready to apply professionally next.	Learners can lead agency-preserving protocols in projects by defining human-decision triggers, including when to pause, override, or decommission misaligned agentic behaviors; they recognize dependency risks by auditing usage; ready to mentor next.	Learners can enforce restraint in job tasks by requiring documented human rationale, including when to pause, override, or decommission misaligned agentic behaviors; they avoid automation bias by escalating high-stakes cases; ready to shape culture next.	Learners can institutionalize human-agency norms by leading policy reviews, including when to pause, override, or decommission misaligned agentic behaviors; they recognize enterprise dependency by requiring training; accountable for organizational judgment culture.

9. Communication and AI Literacy Across Audiences	Learners can explain AI simply by presenting to classmates with examples; they avoid hype by stating limits; ready to add risks next.	Learners can translate concepts for peers by creating plain-language summaries; they recognize audience needs by adjusting explanations; ready to include implications next.	Learners can communicate risks and limitations by supporting informed choices in group presentations; they avoid jargon by testing understanding; ready to adapt for professionals next.	Learners can explain AI uses to non-technical stakeholders by documenting guardrails; they recognize cross-audience gaps by gathering feedback; ready to translate in teams next.	Learners can deliver cross-disciplinary explanations by presenting to mixed groups; they avoid miscommunication by citing evidence; ready to support decisions next.	Learners can lead literacy sessions for diverse audiences by tailoring risk communications; they recognize translation challenges by revising one explanation after audience comprehension checks; ready to influence careers next.	Learners can brief teams on AI implications by using plain language with examples; they avoid confusion by checking comprehension; ready to advise leadership next.	Learners can shape organizational communication strategies by auditing audience understanding and influencing policy; they recognize literacy gaps by measuring outcomes; accountable for enterprise AI literacy.
10. Sector and Pathway Awareness	Learners can describe basic AI uses across simple roles by comparing school examples; they avoid generic views by noting differences; ready to map transitions next.	Learners can connect AI skills to career clusters by exploring case studies; they recognize transition points by discussing middle-to-high school shifts; ready to choose pathways next.	Learners can align projects to pathway expectations by explaining transfer versus workforce constraints; they avoid mismatch by referencing adaptability; ready to specialize next.	Learners can map credentials to sector workflows by documenting direct-to-workforce constraints; they recognize compliance differences by adapting use cases; ready to deepen theory or application next.	Learners can distinguish transfer-oriented theory from applied pathways by comparing coursework; they avoid misalignment by identifying readiness checkpoints; ready to specialize next.	Learners can analyze pathway-specific AI demands by leading capstone alignments while incorporating direct-to-workforce compliance constraints such as data classification, record retention, and auditability; they recognize evolution needs by planning adaptability; ready to enter workforce next.	Learners can adapt AI skills to role-specific constraints by following sector rules; they avoid pathway gaps by updating based on job feedback; ready to mentor next.	Learners can lead pathway and sector strategy by aligning talent development; they recognize cross-sector shifts by updating governance; accountable for organizational adaptability.
11. Cyber Security and Safety	Learners can protect basic data by never pasting personal information into AI tools; they avoid leakage by asking permission first; they learn that some AI tools can act autonomously and must be monitored; ready to recognize suspicious requests next.	Learners can spot prompt injection attempts by ignoring override instructions; they recognize deepfake patterns by verifying with known sources; they understand that agentic tools can chain multiple actions; ready to follow rules next.	Learners can apply acceptable-use policies by classifying data before AI input and verifying urgent requests via separate channels while experimenting with AI-assisted detection and always validating outputs; they avoid social engineering by reporting deepfakes; they practice safe interaction with simple agentic behaviors; ready to integrate secure handling next.	Learners can secure AI tools in technical settings by redacting sensitive data, detecting adversarial inputs, and escalating prompt-injection attempts per protocol; they recognize AI-enabled phishing by confirming identities; they identify risks from persistent agent actions and tool misuse; ready to support detection functions next.	Learners can protect confidentiality and integrity by documenting data-classification routines and testing for model misuse in coursework; they avoid data leakage by using approved environments only; they analyze emerging agent-specific risks (tool misuse, excessive agency, autonomous escalation); ready to analyze threats next.	Learners can implement organizational controls by designing verification routines, monitoring for deepfakes and poisoned inputs, and using AI for basic detection while validating outputs; they recognize user-responsibility gaps by auditing compliance; they evaluate security of multi-step agentic systems; ready to govern next.	Learners can enforce secure practices by redacting workplace data, responding to prompt injection (including indirect variants) and memory-poisoning attempts per policy, and verifying AI-assisted threats like automated attacks; they classify data sensitivity and log escalations per protocol; they monitor and report risks from deployed AI agents (persistent memory, cascading actions); ready to advise leadership next.	Learners can establish enterprise cybersecurity governance by auditing AI-related risks (including memory poisoning, supply-chain poisoning, and agentic threats), mandating data-protection training, and requiring human oversight; they recognize systemic vulnerabilities by requiring incident reporting; they set standards for agentic AI risks (excessive agency, tool-abuse vectors, long-running autonomy) across organizational and sector contexts; accountable for organizational safety culture.

Authority Alignment

The AI Learning Progression is grounded in and aligned with the following national and international authorities:

- Progression emphasizes human-centered design and durable skills (critical thinking, judgment, agency) per the OECD Digital Education Outlook 2026 and OECD/EC AI Literacy Framework.
- Risk-based language (bias, transparency, accountability, human oversight) mirrors the NIST AI Risk Management Framework (1.0 with 2024 Generative AI Profile and 2025 updates) and EU AI Act high-risk system obligations beginning 2 August 2026.
- Equity, inclusion, and protection of learner rights draw explicitly from UNESCO's "AI and education: protecting the rights of learners" (2025).
- Disclosure norms, appropriate-use distinctions, and verification routines align with UNESCO's Guidance for generative AI in education and research and the 2025 "AI and education: protecting the rights of learners" report.
- Cybersecurity progression (data leakage prevention, prompt injection, deepfakes, acceptable use) reflects EU AI Act transparency rules and emerging OECD emphasis on safe AI interaction.
- Sector- and pathway-aware deepening supports the OECD Skills Outlook 2025 focus on adaptability across education-to-workforce transitions.
- K-12 foundational stage with progressive development across grade bands and increasing judgment in postsecondary/workforce contexts follows the U.S. DOL Artificial Intelligence Literacy Framework (2026) and ETS K-12 AI literacy progression models.
- Cross-audience communication and ethical reasoning integrate UNESCO teacher-competency dimensions and EU AI Act Article 4 literacy requirements.
- Overall handoff visibility and restraint against overreliance draw from the draft OECD/EC AI Literacy Framework emphasis on preserving human responsibility.

Reference List

OECD Digital Education Outlook 2026

OECD (2026), OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education, OECD Publishing, Paris. Released 19 January 2026.

OECD Skills Outlook 2025

OECD (2025), OECD Skills Outlook 2025: Building the Skills of the 21st Century for All, OECD Publishing, Paris. Released 9 December 2025.

OECD/EC AI Literacy Framework for Primary and Secondary Education

European Commission & OECD (2025), Empowering Learners for the Age of AI: An AI Literacy Framework for Primary and Secondary Education (review draft; final version expected 2026).

NIST AI Risk Management Framework

NIST AI RMF 1.0 (January 2023; under revision): Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). Includes Generative AI Profile (NIST AI 600-1, July 2024) and 2025 updates/guidance.

UNESCO Guidance for Generative AI in Education and Research

UNESCO (2023), Guidance for generative AI in education and research. Released 7 September 2023 (updated 2026).

UNESCO “AI and Education: Protecting the Rights of Learners”

UNESCO (2025), AI and education: protecting the rights of learners. Released 4 September 2025.

EU Artificial Intelligence Act

Regulation (EU) 2024/1689 (Artificial Intelligence Act). Most high-risk rules apply from 2 August 2026.

U.S. Department of Labor Artificial Intelligence Literacy Framework (2026)

U.S. Department of Labor, Employment and Training Administration (2026), The U.S. Department of Labor’s Artificial Intelligence Literacy Framework. Training and Employment Notice (TEN) 07-25 (13 February 2026).

ETS Preparing K-12 Students With AI Literacy

Chakraburty, S., Ober, T. M., & Liu, L. (2025). Preparing K–12 Students With AI Literacy: Proposed Framework, Progression, and Task Design Principles (ETS Research Report No. RR-25-14). Released November 2025.

Council and Contributing Partners

The Mississippi AI Framework was developed by working groups of the AI Workforce Readiness Council, chaired by Dr. Kollin Napier, Director of the Mississippi Artificial Intelligence Network (MAIN).

This collaborative effort draws on the expertise of representatives from the following organizations and institutions:

- AccelerateMS
- Alcorn State University
- Amazon Web Services (AWS)
- Belhaven University
- 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

This broad representation ensures the framework is practical, equitable, and directly aligned with Mississippi's educational systems, workforce needs, and economic priorities.

High-Stakes Human-Review Triggers

The following categories of AI output require mandatory human review before action is taken. These triggers apply across all domains and levels 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