

AI-Ready America

A two-day workshop conducted by SeedAI to develop recommendations on accelerating AI diffusion, access, and adoption for all Americans.

This material is based upon work supported by the U.S. National Science Foundation under Award No. 2608403 and the Alfred P. Sloan Foundation under Award No. G-2025-79265

The opinions, findings, and conclusions or recommendations expressed in this material are informed by the contributions and perspectives of the workshop participants, and do not necessarily reflect the views of the U.S. National Science Foundation or the Alfred P. Sloan Foundation.

TABLE OF CONTENTS

1. Executive Summary	2–4
2. Workshop Participants	5–6
3. Introduction	6–7
4. Methodology	7
5. Workshop Summary	
5.1 Session One: State and Local Coordination	7–13
5.2 Session Two: Nation-Scale Efforts	13–16
5.3 Session Three: Domain-Specific Strategies for AI Readiness	16–21
5.4 Session Four: AI Literacy and Learning Pathways	21–25
5.5 Conclusion	25
6. Statement of Principles	25–27
7. Roadmap for Future Action	27–28
8. Acknowledgements	29

1. Executive Summary

Artificial intelligence (AI) is poised to significantly transform all sectors of the economy and society in the United States, and in many cases is already doing so. Yet despite some promising early efforts, the programmatic infrastructure required for nationwide AI readiness remains underdeveloped. With support from the U.S. National Science Foundation (NSF) and the Alfred P. Sloan Foundation, SeedAI convened approximately 100 experts for the two-day **AI-Ready America Workshop** to examine how the United States can strengthen the institutional infrastructure needed for AI diffusion, access, and adoption. Sessions were organized around four interrelated areas: state and local coordination, nation-scale efforts, domain-specific strategies for AI readiness, and AI literacy and learning pathways.

Across sessions, participants consistently identified a set of interrelated challenges. The U.S. already possesses distributed institutional infrastructure capable of supporting AI diffusion at scale, including the Cooperative Extension System, Small Business Development Centers (SBDCs), libraries, universities and community colleges, and professional associations. But these networks lack the staffing, training, and sustained resources necessary to integrate AI expertise into their existing missions. The practitioners and intermediaries who staff these institutions, rather than end users, emerged as the highest-leverage investment target for national AI readiness, yet professional development for this workforce remains critically under-resourced. Effective programs exist across sectors and regions but remain largely isolated beyond their immediate networks, and the absence of coordination infrastructure, shared vocabulary, and common standards reinforces fragmentation.

Participants also converged on what works. AI adoption efforts that begin with defined community or institutional challenges consistently outperform those that start with the technology itself. Trust, earned through sustained local presence, demonstrated neutrality, and peer relationships, is a foundational condition for adoption and cannot be manufactured through communications campaigns or top-down mandates. And without intentional program design, the benefits of AI adoption will concentrate among those already best positioned to access them. Broad access must be an architectural choice, not an afterthought.

The findings and recommendations from each of the workshop's four sessions are summarized below:

<i>Session One: State and Local Coordination</i>	
Findings	Recommendations
<ol style="list-style-type: none"> 1. Existing trusted networks provide infrastructure for AI diffusion 2. Problem-first approaches drive more effective adoption 3. Trust is better earned through local legitimacy, not mandates 4. Professional development is critically under-resourced 5. Broad access requires intentional design and implementation 6. Cross-sector coordination is urgently needed 7. Workforce development should span all economic sectors 	<ol style="list-style-type: none"> 1. Fund capacity within existing trusted networks 2. Institutionalize problem-to-solution pathways and facilitator training 3. Build regional coordination capacity 4. Resource state and local readiness planning
<i>Session Two: Nation-Scale Efforts</i>	
Findings	Recommendations
<ol style="list-style-type: none"> 1. Broad access should be treated as a structural prerequisite of any national AI strategy 2. Successful national technology adoption follows a proven three-part formula 3. Intermediary “super nodes” are a high-leverage investment target 4. National standards and shared vocabulary are valuable coordination opportunities 	<ol style="list-style-type: none"> 1. Build and grow AI research and compute commons 2. Establish shared definitions and voluntary certification pathways 3. Implement national AI readiness indicators

<p>5. Public-private partnerships work when structured around use cases rather than tool adoption</p>	
<p><i>Session Three: Domain-Specific Strategies</i></p>	
<p>Findings</p>	<p>Recommendations</p>
<ol style="list-style-type: none"> 1. AI readiness looks different in every domain—strategies must reflect that 2. Data infrastructure is a critical domain-specific bottleneck 3. Regulatory frameworks can create asymmetric barriers across sectors 4. Front-line practitioners are adopting ahead of their institutions 5. Domain-specific models reveal what works—and what doesn't 	<ol style="list-style-type: none"> 1. Establish coordinating infrastructure within and across sectors 2. Fund domain-specific data infrastructure and governance 3. Support institutional catch-up to front-line adoption 4. Equip existing diffusion infrastructure for domain-specific needs 5. Commission sector-specific regulatory reviews
<p><i>Session Four: AI Literacy and Learning Pathways</i></p>	
<p>Findings</p>	<p>Recommendations</p>
<ol style="list-style-type: none"> 1. Train-the-trainer models are needed at the national level 2. AI credentials proliferate without quality control or labor market alignment 3. Digital divide and trust barriers limit reach 4. Existing curricula are disconnected from practical utility 5. Traditional curriculum cycles cannot keep pace with AI's rate of change 	<ol style="list-style-type: none"> 1. Establish a national AI educator training consortium 2. Develop a framework for AI literacy credentials 3. Equip community colleges and the Cooperative Extension System with AI teaching capacity 4. Support trust-centered delivery for under-resourced communities 5. Build mechanisms for continuous AI curriculum renewal

Recommendations are directed to federal, state, and local government stakeholders; philanthropic foundations; civil society and community organizations; educational institutions; and industry partners.

2. Workshop Participants

The AI-Ready America Workshop brought together approximately 100 participants from federal agencies, state and local governments, industry, philanthropy, higher education, and civil society. We are grateful to all participants for the time and expertise they contributed to this effort. The following individuals participated in the workshop:

Workshop Organizers: **Austin Carson**, SeedAI; **Marina Meyjes**, SeedAI; **Josh New**, SeedAI; **Anna Rulloda**, SeedAI; **Stuart Styron**, SeedAI

Workshop Participants: **Bethany Abbate**, Software & Information Industry Association; **Elizabeth Albro**, Institute of Education Sciences; **Sheriff Almakki**, Americans for Responsible Innovation; **Allen Antoine**, The Texas Advanced Computing Center (TACC); **Taylor Barkley**, Abundance Institute; **Chaitan Baru**, U.S. National Science Foundation; **Francine Berman**, UMass Amherst; **Amanda Bickerstaff**, AI for Education; **Laura Biven**, Jefferson Labs; **Steve Brown**, U.S. National Science Foundation AI Institutes Virtual Organization; **Adam Browning**, Washington Leadership Academy; **Frances Carter-Johnson**, South Big Data Regional Hub; **Daniel Castro**, Information Technology & Innovation Foundation; **Megan Catterton**, ICF (U.S. National Science Foundation Contractor); **B Cavello**, Aspen Digital; **An-Me Chung**, New America Foundation; **Fay Cobb Payton**, Rutgers University–Newark; **Christophe Combemale**, Carnegie Mellon University; **Mary Crowe**, U.S. National Science Foundation; **Jack Cumming**, Former Deputy Chief of Staff for the White House Office of Science and Technology Policy; **Arthur Daemrich**, Arizona State University; **Jessica Daluz Hill**, Google; **Krista D'Amelio**, Code.org; **Tess DeBlanc-Knowles**, Atlantic Council Technology Programs; **Jordan DiMaggio**, UPCEA - The Online and Professional Education Association; **Anjelica Dortch**, Independent Community Bankers of America; **Sarah Dunton**, U.S. National Science Foundation; **Harrison Durland**, Department of Labor; **Jack Furth**, National Small Business Association; **Nardos Ghebreab**, Beyond100K; **Erwin Gianchandani**, U.S. National Science Foundation; **Shaina Glass**, Computer Science Teachers Association; **Rubella Goswami**, National Institute of Food and Agriculture; **Venu Govindaraju**, National AI Institute for Exceptional Education; **Josh Greenberg**, Alfred P. Sloan Foundation; **Hope Hartman**, Larimer Small Business Development Centers; **Evan Heit**, U.S. National Science Foundation; **Ilana Herold**, ICF (National Science Foundation Contractor); **Melissa Hopkins**, Johns Hopkins Center for Health Security; **Michael Hout**, New Mexico State University; **Rachael Houston-Carter**, Accenture; **Florence Hudson**, Northeast Big Data Innovation Hub, Data Science Institute, Columbia University; **Nicholas Ivory**, U.S. Small Business Administration; **Ashley Jeffrey**, Washington Leadership Academy; **Chad Jenkins**, University of Michigan; **Shalin Jyotishi**, New America Foundation; **Chad Lane**, University of Illinois, Urbana-Champaign; **Anna Lenhart**, George Washington University; **Chauncy Lennon**, Lumina Foundation; **Kevin Logan**, National Association for

Community College Entrepreneurship; **William Mapp**, Morgan State University; **Karl Martin**, UW-Madison Division of Extension; **Maria Marzullo**, Association of Public and Land-grant Universities; **Todd McCracken**, National Small Business Association; **Jaci McCune**, Expanding Computing Education Pathways Alliance; **Danielle S. McNamara**, Arizona State University; **Cierra Mitchell**, Department of Labor; **Xavier Monroe**, U.S. National Science Foundation; **Richard Montez**, Hispanic Association of Colleges and Universities; **James L. Moore III**, U.S. National Science Foundation; **Charlie Moskowitz**, Accenture; **Steven Moss**, National Security Commission on Emerging Biotechnology; **Alvaro J. Muñoz**, Association of Public and Land-grant Universities; **Nathaniel Putnam**, U.S. Small Business Administration; **Elizabeth Newbury**, American University School of Communication; **Tobi Olaiya**, Salesforce; **Greg Peterson**, U.S. National Science Foundation; **Courtney Pollack**, Institute of Education Sciences; **Becca Portman**, Patrick McGovern Foundation; **Ann Quiroz Gates**, University of Texas El Paso; **Jeremy Roschelle**, Digital Promise; **AJ Segal**, Scale AI; **Jennifer Shieh**, U.S. Small Business Administration; **Becky Smerdon**, Teach for America; **John Soroushian**, Americans for Responsible Innovation; **Taylor Stockton**, Department of Labor; **Cody Stone**, Montana State University Extension; **Meme Styles**, MEASURE; **K. Tighe**, Laude Institute; **Tim Toomey**, Accenture LearnVantage; **Eric Tucker**, The Study Group; **Nicol Turner Lee**, Center for Technology Innovation, Brookings Institution; **Grant Van Eaton**, Teach for America; **Tim VanReken**, Headwaters Tech Hub; **Margie Vela**, Amazon Web Services, Machine Learning University; **Shaowen Wang**, University of Illinois Urbana-Champaign; **Talitha Washington**, Center for Applied Data Science and Analytics, Howard University; **Aaron Weibe**, Extension Foundation; **Samuel Wells**, Autio Strategies; **Emma Westerman**, U.S. National Science Foundation; **Julia Wynn**, Code.org; **Holly Yanco**, UMass Amherst; **Ellen Zegura**, U.S. National Science Foundation; **Michal Ziv-El**, U.S. National Science Foundation

3. Introduction

Artificial intelligence is reshaping economic activity, public services, and institutional operations across the United States. Federal executive orders, bipartisan legislative proposals, and agency-level strategies have established AI as a national priority. While promising coordination efforts are beginning to take shape, the programmatic infrastructure required to translate these advances into meaningful improvements in how people live, learn, and work remains uneven and underdeveloped relative to the pace of change.

The AI-Ready America Workshop, conducted by SeedAI in Washington, D.C., on January 29-30, 2026 with support from the U.S. National Science Foundation and the Alfred P. Sloan Foundation, convened stakeholders from across the country to examine this gap. SeedAI is a Washington, D.C.-based nonprofit focused on national AI readiness. SeedAI operates at the intersection of AI policy and practical application, spanning hands-on AI literacy and adoption efforts, educational programs for government staff and officials, and strategic agenda-setting to inform policy initiatives.

The workshop brought together approximately 100 experts from federal agencies, state and local governments, industry, philanthropy, higher education, and civil society. Sessions were organized

around four interrelated areas: state and local coordination (Session I); nation-scale efforts (Session II); domain-specific strategies for AI readiness (Session III); and AI literacy and learning pathways (Session IV). Each session combined lightning talks from practitioners and researchers with facilitated working group discussions conducted under the Chatham House Rule, with participants from different sectors assigned to each group.

This report documents findings and recommendations from each session, a Statement of Principles distilled from the workshop's cross-cutting themes, and a Roadmap for Future Action identifying priority areas for AI readiness efforts.

4. Methodology

The workshop was designed to surface expert knowledge and cross-sector perspectives through a structured, multi-format process. Participants were selected through a deliberate recruitment process designed to ensure broad representation across geography, institution type, and sector. Invitations prioritized practitioners with direct implementation experience alongside senior officials. A core design principle of the workshop was to avoid over-reliance on frequently convened voices and overrepresentation from coastal regions.

The workshop was conducted in a hybrid format, with both in-person and virtual participation. Sessions combined expert lightning talks and keynote presentations with facilitated small-group discussions organized around predefined thematic areas. Participants from different sectors were deliberately assigned to each table to promote cross-disciplinary dialogue, and working group discussions were conducted under the Chatham House Rule to encourage candid exchange. After the workshop, participants were invited to fill out a survey to provide additional input.

The report draws on multiple data sources: transcripts of all table discussions, presentation materials from lightning talks and keynotes, and participant survey responses. Findings were developed through thematic analysis of these sources, identifying recurring themes, points of convergence and tension, and actionable priorities across sessions. Recommendations were derived from areas where multiple sessions independently surfaced consistent conclusions, grounded in practitioner testimony, cross-sector discussion, and supporting evidence cited in presentations.

A draft report was reviewed by a select group of workshop participants and external subject-matter experts, whose comments and inputs were incorporated into the final document.

5. Workshop Summary

5.1. Session One: State and Local Coordination

Session Overview

Some of the most promising AI adoption work is happening at the state and local level, where proximity to communities, businesses, and institutions allows for tailored approaches. This

session explored what is working in regions that have made progress, what coordination infrastructure is needed, and how successful models can be adapted for places with fewer resources or less existing momentum.

Findings

1. Existing Trusted Networks Provide Infrastructure for AI Diffusion

A consistent finding emerged across the session: the U.S. already possesses some of the distributed institutional infrastructure necessary for AI diffusion at the state and local level. The Cooperative Extension System has maintained a presence in nearly every U.S. county for more than one hundred years; the Small Business Development Centers (SBDC) network operates through 63 state and territorial systems with over 1,000 centers nationwide; and libraries, community colleges, 4-H clubs, and Boys & Girls Clubs serve communities that any new initiative would require years to reach. These institutions share an established community trust, built through decades of local presence and deliberate neutrality. For instance, Cooperative Extension agents employed by land-grant universities to serve the citizens of their state are not tied to corporate interests, which significantly shapes community willingness to engage.

Models already demonstrating this approach include the Discovery Partners Institute in Chicago (a joint venture of the University of Illinois System, the City of Chicago, and the state),¹ the Colorado SBDC network,² and Amazon Web Services (AWS) collaboration with specialized institutions.³ In Austin, Texas, Measure’s “Community in the Loop” initiative, developed in partnership with the City of Austin and the Austin AI Alliance, extends this model to municipal AI governance.⁴ By convening residents, technologists, and public officials in structured dialogue, it positions community trust as operational infrastructure for responsible AI adoption.

These networks are well positioned but under-resourced, lacking the staff capacity, training, and equipment to integrate AI expertise into their existing missions. This plays out most visibly on the ground. For example, educators frequently report being overwhelmed by the volume of available AI tools with little guidance on where to begin—underscoring the need for trusted, community-connected intermediaries that can separate signal from noise and translate capability into practice.

2. Problem-First Approaches Drive More Effective Adoption

Broad consensus held that state and local AI adoption efforts that begin with defined community or institutional challenges generate more durable engagement than those that start with the technology itself. As an example, participants acknowledged New America's Design Labs, which convene researchers, policymakers, and practitioners to define problems before evaluating technological solutions.⁵ A university partnership in Belchertown, Massachusetts helped local

¹<https://dpi.illinois.edu/>

²<https://sbdc.colorado.gov/>

³<https://www.aboutamazon.com/news/aws/aws-launches-new-ai-program-for-community-colleges-msis-hbcus>

⁴<https://www.austintexas.gov/news/city-austin-hosts-community-loop-engage-residents-ai-accountability>

⁵<https://www.newamerica.org/insights/what-makes-education-a-public-good/>

officials identify suitable municipal challenges prior to recommending AI applications.⁶ The New Mexico State University's AI Institute collaborated with cattle farmers to surface geotag-based tracking as an operational need before introducing AI concepts.⁷ Participants emphasized the corresponding risk of technology-forward strategies—"solutions in search of problems"—particularly amid rapid proliferation of AI tools deployed without clear use cases.

3. Trust Is Better Earned Through Local Legitimacy, Not Mandates

Trust emerged as a foundational prerequisite for state and local AI adoption, with participants explicit about the mechanisms through which it develops: sustained local presence, demonstrated neutrality, and peer relationships. Participants expressed that communications campaigns or top-down mandates cannot meaningfully compensate for the absence of these trust-building mechanisms. For example, community banks adopt AI when peer institutions demonstrate success. Farmers engage with Cooperative Extension agents because they are not tied to commercial interests. Educators adopt new practices when they observe colleagues in comparable settings succeeding. This peer-to-peer dynamic was the most frequently cited adoption driver among workshop participants. Code.org's model illustrates how this can operate at scale: teachers training teachers within established professional communities.⁸ The broader conclusion was that trust deficits are most effectively bridged through local intermediaries with well-established relationships.

4. Professional Development Is Critically Under-Resourced

At the state and local level, insufficient funding for sustained, compensated professional development emerged as a major barrier. Teachers lack protected and compensated time for professional learning; Cooperative Extension agents and SBDC counselors require AI training themselves before they can support others yet lack capacity for both roles; and faculty at universities and community colleges, as well as other smaller institutions, lack resources to engage with rapidly evolving technologies.

Working models exist. For example, Code.org's state-level partnerships integrate professional development into existing state networks and the Cooperative Extension's three-tiered funding model—federal, state, and county—is specifically designed to sustain a permanent educator presence at the local level.⁹ Participants proposed additional mechanisms including fellowships, micro-grants for peer observations, and summer institutes modeled on existing NSF programs. Effective professional development is sustained, compensated, and embedded in evidence-based practice and process.

5. Broad Access Requires Intentional Design and Implementation

Without intentional design, the default outcome is uneven AI adoption. The current landscape already demonstrates this: well-resourced school networks such as the Knowledge is Power

⁶https://www.belchertown.org/AgendaCenter/ViewFile/Minutes/_12112025-1326

⁷<https://citizenportal.ai/articles/6131834/nmsu-presents-ai-institute-projects-precision-ranching-extension-chatbot-and-4h-curriculum-updates>

⁸<https://code.org/en-US>

⁹<https://www.nifa.usda.gov/grants/programs/capacity-grants/smith-lever-act-capacity-grant>

Program (KIPP) maintain dedicated curriculum teams for AI integration,¹⁰ while underfunded schools often lack capacity to begin. Rural communities face compounding barriers—broadband gaps, scarcity of local AI expertise, limited capacity to navigate public funding, and geographic distance from peer learning networks—that will deepen without deliberate countervailing investment.

The programs making progress on access share common design principles. AWS's Machine Learning University provides free nine-month faculty upskilling cohorts at HBCUs and community colleges to provide access to industry-aligned tools and curriculum for advanced AI and machine learning.¹¹ The HBCU AI Conference and Training Summit, hosted annually at Huston-Tillotson University, demonstrates intentional design for expanding AI access by centering the nation's HBCUs as innovation hubs rather than peripheral participants.¹² Through sponsored student access, faculty upskilling, and cross-sector partnerships, the summit reduces financial and network barriers while strengthening long-term institutional capacity within historically under-resourced ecosystems.

The Computing Alliance of Hispanic-Serving Institutions (CAHSI) connects over 70 two-year and four-year colleges to expand participation across a broad range of institutions in computing, and is now piloting embedded AI ethics curricula across member institutions.¹³ In each case, broad reach was an architectural choice, achieved through distributed placement, institutional partnerships, sustained engagement, and multilingual delivery, rather than an afterthought.

6. Cross-Sector Coordination Is Urgently Needed

The absence of coordination infrastructure remains a significant barrier to effective AI diffusion at the state and local level. The current landscape is fragmented and siloed, with thousands of education and workforce programs operating across states with no shared framework for assessing what works, for whom, and under what conditions. The result is not only duplication of effort, but inconsistent messaging, uneven program quality, and an irregular spread of knowledge across regions and institutions. Practitioners struggle to identify credible models, peer learning is ad hoc rather than structured, and promising practices fail to scale beyond isolated pilots.

State and local AI efforts are primarily fragmented in two ways: sectors work in isolation from each other (education, workforce development, small business, and local government rarely coordinate), and effective models within sectors remain largely invisible to peers who could adopt them. Participants described discovering successful state programs only through accidental connections at convenings like the AI-Ready America Workshop.

Emerging models show what coordination infrastructure can enable: the Center for Civic Futures convenes state chief AI officer cohorts through monthly virtual roundtables;¹⁴ the William and Flora Hewlett Foundation has funded cross-community convenings connecting stakeholders

¹⁰<https://www.kipp.org/>

¹¹<https://www.aboutamazon.com/news/aws/aws-launches-new-ai-program-for-community-colleges-msis-hbcus>

¹²<https://htu.edu/academics/college-school/university-college/hbcu-ai-conference-training-summit-2026/>

¹³<https://cahsi.utep.edu/about/>

¹⁴<https://www.rockefellerfoundation.org/news/rockefeller-foundation-and-center-for-civic-futures-launch-ai-readiness-project-to-build-state-capacity-for-responsible-ai/>

across regions.¹⁵ However, these efforts remain largely episodic. Effective coordination requires structured mechanisms connecting stakeholders to resources at the right time. Without intentional investment, promising initiatives will remain isolated.

7. Workforce Development Should Span All Economic Sectors

Participants identified a mismatch between the breadth of AI's workforce impact and the structure of current readiness efforts. AI is reshaping work across skilled trades, agriculture, healthcare, manufacturing, and services. While AI training resources have expanded rapidly from technology companies, universities, and industry groups, they remain fragmented and unevenly accessed.

The disparity in understanding, adoption, and applied skills outside the technology sector poses additional challenges. Much of the available content is generic, disconnected from day-to-day workflows, or difficult for non-technical workers to translate into practical use.

Reliable, consistent adoption is most likely when AI literacy is embedded in occupation-specific contexts. State and local institutions, such as universities and community colleges, apprenticeship programs, industry associations, unions, and trade organizations, are well positioned to integrate AI training into existing vocational and professional pathways, aligning skills development with real workplace demands rather than abstract technical competencies.

Recommendations

R1. Fund Capacity Within Existing Trusted Networks

Investment in state and local AI adoption should prioritize sustained capacity—staff, training, and equipment—within trusted networks. These include Cooperative Extension, SBDCs, libraries, technology access focused nonprofit organizations, community colleges, universities, and allied youth-serving organizations. Federal agencies should lead with multi-year capacity funding modeled on the Cooperative Extension System's approach, in which federal dollars are dedicated to maintaining educators on the ground. Philanthropic foundations should support the convenings and cross-community coordination needed to connect these networks, enabling shared resources, collective learning, and cross-sector problem solving. State governments should match these efforts by allocating for capacity within their own systems. This should include teacher stipends, substitute coverage, and protected time for professional learning. Industry partnerships can provide resources and expertise but should be structured with governance guardrails against vendor dependence. Across all sources, investment should include compensated professional development for both intermediaries and the practitioners they support.

R2. Institutionalize Problem-to-Solution Pathways and Facilitator Training

States and local partners should invest in problem-to-solution frameworks that begin with clearly defining a concrete local challenge and only then evaluating whether and how AI is an

¹⁵<https://hewlett.org/programs/cyber/>

appropriate response. Facilitators across sectors should be trained to lead this process and this training should be embedded in local professional communities and designed for peer learning, directly reducing the risk of technology-forward deployment.

R3. Build Regional Coordination Capacity

Federal agencies, philanthropic foundations, and state governments should invest in regionally governed coordination hubs that leverage the physical presence and community trust of Cooperative Extension networks, community colleges, and public universities with two explicit functions. First, cross-sector convening and support for coordination: bringing together practitioners from education, workforce, small business, and local government within a region to identify shared challenges and avoid duplicative efforts. Second, within-sector matchmaking: actively curating and connecting proven models to practitioners in other jurisdictions through facilitated peer learning, regular convenings, and dedicated staff whose role is to know what is working and connect people to it. The Center for Civic Futures' state chief AI officer cohort models the cross-sector function;¹⁶ the Extension Foundation's county agent network models the within-sector function. Critically, participants emphasized that both consistently collapse without sustained, dedicated funding. Coordination infrastructure cannot run on grant cycles alone.

R4. Resource State and Local Readiness Planning

States and localities need structured processes to assess baseline conditions—connectivity, institutional capacity, workforce readiness, existing diffusion channels—before selecting interventions. A structured planning process, paired with technical assistance and cross-state learning, can help jurisdictions identify targeted AI-readiness gaps rather than defaulting to one-size-fits-all strategies. Readiness planning should incorporate two commitments from the outset. First, access-by-design: distributed delivery models, multilingual content, physical access points with technology endpoints—such as laptops and WiFi hotspots in libraries—for under-resourced institutions or digital deserts. Second, workforce relevance across sectors: training pathways reaching workers beyond high-tech roles through community colleges, apprenticeship programs, unions, and industry groups.

Pennsylvania and New Jersey illustrate early state-level approaches to operationalizing AI readiness strategies. Pennsylvania established a Generative AI Governing Board by executive order,¹⁷ partnered with Carnegie Mellon University and OpenAI on a first-in-the-nation government AI pilot,¹⁸ and launched AI literacy training for state employees.¹⁹ New Jersey paired its AI tax credit program for large-scale investments with a \$20 million NJ AI Hub Fund, backed by the New Jersey Economic Development Authority and CoreWeave, to support startups affiliated with the state's AI Strategic Innovation Center.²⁰

¹⁶<https://civicfutures.org/>

¹⁷<https://www.pa.gov/agencies/oa/programs/information-technology2/gen-ai>

¹⁸[https://www.cmu.edu/news/stories/archives/2025/march/shapiro-unveils-ai-pilot-program-results-at-carnegie-mello](https://www.cmu.edu/news/stories/archives/2025/march/shapiro-unveils-ai-pilot-program-results-at-carnegie-mellon)

¹⁹<https://www.pa.gov/agencies/oa/newsroom/agreement-innovateus-to-train-responsible-generative-artificial->

²⁰<https://www.njeda.gov/njeda-and-coreweave-announce-creation-of-20m-ai-hub-fund-to-support-innovative-startups/>

5.2 Session Two: Nation-Scale Efforts

Session Overview

Making America AI-ready requires more than a collection of local efforts. This session examined what kinds of national coordination, shared infrastructure, and cross-sector partnerships are needed to accelerate AI diffusion, and how existing federal programs and established networks can be leveraged to reach every community.

Findings

1. Broad Access Should Be Treated As a Structural Prerequisite of Any National AI Strategy

Without deliberate design, AI diffusion will predictably reflect and potentially intensify existing geographic, economic, and institutional variation in access. Public trust in AI remains limited, with only 32 percent of Americans reporting confidence in the technology, and global survey data indicates that Americans lag behind many peer nations in generative AI usage.²¹

Meanwhile, advanced adoption is concentrated in higher-income urban and suburban regions,²² while rural communities face compounding constraints, including limited broadband infrastructure and fewer locally embedded technical resources. Risks also emerge within the systems themselves: models trained on unrepresentative data can degrade performance.

These outcomes are not inevitable. The National Student Data Corps illustrates that broad reach is achievable when access is embedded as a core architectural principle. By engaging over 1,350 institutions spanning varied geographies, the program reached more than 20,000 participants distributed across the United States, rather than just wealthy urban hubs. This program demonstrates how intentional infrastructure design can create opportunity for all Americans.²³

2. Successful National Technology Adoption Follows a Proven Three-Part Formula

A consistent lesson emerged across presentations and discussions: successful national AI readiness requires federal investment, trusted intermediary institutions, and locally responsive implementation.

Evidence from multiple domains reinforces this. During COVID-19, the Cooperative Extension System leveraged its county agent network to support vaccine outreach through the EXCITE program. This CDC-funded initiative channeled nearly \$10 million through the United States Department of Agriculture (USDA) to 72 land-grant universities, reaching twenty million people in communities often underserved by traditional, centralized channels.²⁴ Similarly, the \$42.45 billion Broadband Equity, Access, and Deployment (BEAD) Program applies this same

²¹<https://www.edelman.com/trust/2025/trust-barometer/report-tech-sector>

²²<https://www.microsoft.com/en-us/corporate-responsibility/topics/ai-economy-institute/reports/global-ai-adoption-2025/>

²³<https://nebigdatahub.org/hub-community-engagement/>

²⁴<https://www.nifa.usda.gov/grants/programs/extension-covid-immunization-training-education-excite>

three-part structure to closing the digital divide. It pairs federal funding and national guardrails with state-level planning and local delivery rather than treating infrastructure investment as sufficient on its own.²⁵ System design—not funding levels alone—determines national reach.

3. Intermediary “super nodes” are a high-leverage investment target

The primary bottleneck to nation-scale AI adoption is not end users but practitioners and intermediaries—“super nodes”—that shape how others engage with technology: teachers rather than students, nurses rather than patients, small business advisors rather than business owners. These intermediaries function as multipliers of impact and trust.

The pattern holds across sectors. Surveys of school system leaders show higher rates of AI use for administrative tasks than for classroom instruction.²⁶ This gap is largely rooted in human capacity and institutional support rather than technological limitations. In small businesses, adoption is expanding,²⁷ however many business owners report limited grasp of how AI tools function or what risks they pose.²⁸ This suggests a gap in the advisory and technical assistance infrastructure that these businesses depend upon. In manufacturing, despite its central role in economic reshoring, 72 percent of U.S. manufacturers cite outdated technology as a barrier to hiring. Thousands of jobs go unfilled for lack of digital and AI skills—pointing to a need for intermediaries who can bridge the gap between available technology and workforce readiness.²⁹

In each case, the limiting factor is the capacity of the intermediaries. Yet these intermediaries consistently receive fewer targeted resources and less strategic investment than the end users they serve. Closing this gap—through sustained investment in training, capacity-building, and institutional support—represents one of the highest-leverage strategies available for achieving nation-scale AI readiness.

4. National Standards and Shared Vocabulary Are Valuable Coordination Opportunities

The absence of shared standards and vocabulary for AI readiness is an underexplored barrier to nation-scale coordination. Participants emphasized that “AI literacy,” “AI awareness,” and “AI readiness” are used interchangeably, creating confusion and impeding institutional coordination.

The NIST Cybersecurity Workforce Framework gave public and private sectors a shared vocabulary for roles and skills, now in use across both sectors and translated into five languages.³⁰ The K-12 standards record is similarly instructive: the Computer Science Teachers Association standards and the Next Generation Science Standards, developed by a 26-state consortium, began as voluntary frameworks and became the foundation of state policy. The CS standards had been adopted by only six states in 2017 and now cover nearly all of them.³¹

²⁵<https://broadbandusa.ntia.gov/funding-programs/broadband-equity-access-and-deployment-bead-program>

²⁶<https://bellwether.org/publications/surveying-artificial-intelligence-and-schools/>

²⁷<https://usmsystems.com/small-business-ai-adoption-statistics/>

²⁸<https://advocacy.sba.gov/wp-content/uploads/2025/09/Research-Spotlight-AI-in-Business-Small-Firms-Closing-In-092425.pdf>

²⁹<https://hexagon.com/company/newsroom/press-releases/2025/us-manufacturers-say-outdated-technology-is-hurting-hiring>

³⁰<https://www.nist.gov/itl/applied-cybersecurity/nice/nice-framework-resource-center>

³¹<https://csteachers.org/k12standards/revision/>

AI does not yet benefit from an equivalent, nationally recognized coordinating framework. State-level AI guidance documents and education-focused resource hubs represent meaningful early progress, cataloging best practices, providing implementation guidance, and defining baseline competencies.³² Yet these efforts remain decentralized, unevenly distributed, and variably adopted. They do not yet function as a shared national reference point that enables consistent terminology or comparable metrics across jurisdictions.

5. Public-Private Partnerships Work When Structured Around Use Cases Rather than Tool Adoption

Partnerships organized around public-interest objectives, such as literacy, economic resilience, and measurable access outcomes, operate differently from those centered on tool adoption. The Small Business Digital Alliance, a public-private partnership between the U.S. Small Business Administration (SBA) and national technology companies, frames engagement around AI literacy and practical business use cases rather than specific products or platforms.³³ Funded counselors work directly with small businesses to identify where AI could address operational challenges, enabling adoption grounded in business needs rather than vendor offerings.

By contrast, partnerships structured around specific platforms can introduce lock-in risks. Participants described technology companies attaching exclusivity conditions to institutional partnerships, such as providing funded equipment or tools on the condition that institutions not partner with competitors. In workforce and education contexts, vendor lock-in extends to what people learn: training built around a specific product may not transfer to other systems, and the product itself may not be available in two years.

For nation-scale coordination, these examples underscore that partnerships grounded in use cases and measurable public outcomes are more durable than those oriented toward specific tool adoption, particularly in preserving institutional independence and avoiding commercial lock-in.

Recommendations

R1. Build and Grow AI Research and Compute Commons

Federal investment should support a shared AI compute and research infrastructure accessible to community colleges, universities, and civil society organizations. The National Artificial Intelligence Research Resource (NAIRR) pilot demonstrates how coordinated access to compute, data, models, and training resources can broaden participation in AI research and workforce development.³⁴ Earlier distributed computing initiatives, such as the Extreme Science and Engineering Discovery Environment (XSEDE), illustrate how federally supported infrastructure can expand access to advanced computing capacity across a diverse research ecosystem.³⁵

³²<https://www.aiforeducation.io/ai-resources/state-ai-guidance>;
<https://drive.google.com/file/d/1p1x-UFVec0fwXOBv2B17j65wiD8KYSvQ/view>

³³<https://www.smallbusinessdigitalalliance.com/>

³⁴<https://nairrpilot.org/>

³⁵<https://www.xsede.org/>

Expanding federal efforts and establishing new regionally-focused AI research and compute commons should help significantly reduce barriers to AI research, testing, evaluation, and adoption. These compute commons should leverage state and regional hubs embedded within trusted public institutions including land-grant universities and their Extension networks, regional public universities, community colleges, and SBDCs with federal seed funding structured to catalyze matching state and local investment. This infrastructure should integrate computing capacity, technical assistance, workforce training, and governance mechanisms that prioritize access for under-resourced regions and institutions.

R2. Establish Shared Definitions and Voluntary Certification Pathways

NSF and other relevant federal agencies including the Office of Science and Technology Policy (OSTP), the Department of Labor (DOL), and Department of Education should convene a multi-stakeholder initiative establishing shared definitions, curriculum standards, and voluntary certification pathways for AI literacy. This effort should follow the model of CSTA (Computer Science Teachers Association) standards and Next Generation Science Standards (NGSS): community-developed, peer-reviewed, and iteratively updated. It should deliver actionable definitions of “AI literacy,” “AI awareness,” and “AI readiness” across K–12, higher education, workforce development, and small business contexts.

R3. Implement National AI Readiness Indicators

A federal requirement should be established for biennial reporting on AI adoption, talent development, and access metrics, modeled on NSF's Science and Engineering Indicators.³⁶ Reporting should track adoption rates by sector and firm size; talent pipeline metrics disaggregated by demographic and regional factors; workforce displacement and transition support; and benefit distribution across communities.

5.3. Session Three: Domain-Specific Strategies for AI Readiness

Session Overview

Different domains have developed distinctive approaches to AI readiness based on unique constraints, user needs, and institutional contexts. This session drew on practitioner experiences across several fields to surface practical insights and identify lessons that might translate across sectors.

Findings

1. AI Readiness Looks Different in Every Domain—Strategies Must Reflect That

The session’s central insight is that “AI readiness” means fundamentally different things depending on sector context. In healthcare, readiness centers on liability frameworks and clinical accountability. For example, practitioners need clarity about responsibility when an AI-supported

³⁶<https://nces.nsf.gov/indicators>

recommendation contributes to patient harm. In banking, readiness is shaped by factors like navigating regulatory compliance requirements that differ sharply between large institutions and community banks. In agriculture, readiness often begins with foundational data practices, such as data cleaning, interoperability, and provenance tracking, before advanced machine learning tools can be responsibly deployed.

In K–12 education, readiness requires building coherent integration across computer science, data science, AI literacy, and digital literacy.³⁷ Computer science and data science remain foundational; AI literacy builds on—not replaces—these domains. Many districts are still scaling durable computer science and data capacity, so readiness depends on strengthening this continuum rather than introducing AI as a disconnected add-on.

For small businesses, readiness is often far more practical and resource-constrained: whether an owner has the time, technical guidance, and trusted support to determine if an AI tool meaningfully improves operations. The barrier is not conceptual alignment but bandwidth, risk tolerance, and return on investment. And because small businesses span every industry and domain, they face a compounding challenge: navigating constraints common to resource-limited organizations while also grappling with questions specific to their sector. The concerns of a family-owned pharmacy look nothing like those of a small engineering firm or a local service provider. An established network of peers that have experience with AI adoption is invaluable to those looking to integrate AI into their work.

One-size-fits-all approaches to AI readiness will fail because starting conditions, institutional constraints, and definitions of success vary across every domain. Effective strategies must be context-specific and aligned to clearly defined outcomes.

2. Data Infrastructure Is a Critical Domain-Specific Bottleneck

Practitioners across sectors identified fragmented data infrastructure and inconsistent governance norms as a binding constraint on AI readiness. In agriculture, data cleaning and standardization across heterogeneous land-grant institutions remain under-resourced. Silos between research, teaching, and extension functions compound interoperability challenges. In education, the Family Educational Rights and Privacy Act (FERPA) and the Children’s Online Privacy Protection Act (COPPA) create compliance obligations that are interpreted and implemented differently across states and districts, complicating consistent operations for national ed-tech providers. In banking, intellectual property (IP) concerns and competitive dynamics inhibit data sharing.³⁸ Small businesses lack the resources and expertise for basic data preparation.

The trajectory of AlphaFold in structural biology illustrates what becomes possible when this foundation is in place.³⁹ Its breakthrough performance was only possible due to decades of standardized data collection and open data infrastructure, most notably the Protein Data Bank.⁴⁰ The Open Knowledge Network, an NSF-funded effort to create standardized interoperability

³⁷<https://www.newamerica.org/education-policy/briefs/foundational-skills-digital-literacy-in-the-age-of-ai-analysis-and-voices-from-the-field/>

³⁸<https://www.bis.org/publ/bcbs239.htm>

³⁹<https://alphafold.ebi.ac.uk/>

⁴⁰<https://www.rcsb.org>

between knowledge graphs for broad data leverage, represents an emerging cross-sector approach to building similar foundations. But data governance frameworks should be tailored to each sector's regulatory environment; no single standard will work.

3. Regulatory Frameworks Can Create Asymmetric Barriers Across Sectors

Varied regulatory environments significantly influence AI adoption in different ways across domains. In banking, compliance requirements designed for major institutions impose disproportionate burdens on community banks, while bank examiners themselves often lack sufficient understanding of AI to evaluate its use. The absence of chief AI officers at most community banks increases the difficulty of developing governance frameworks with limited staff. In healthcare, unresolved liability questions create institutional hesitation: when an AI-supported clinical recommendation goes wrong, accountability pathways remain unclear. In agriculture, regulations often constrain autonomous equipment deployment despite technology readiness. In education, the proliferation of state-level AI guidance—while critical for local education agencies—can create national fragmentation rather than coherence. For example, participants pointed to differing interpretations of FERPA across jurisdictions. Federal procurement rules further restrict agencies from timely adoption of commercial AI tools.

Participants were broadly aligned on this diagnosis but divided on solutions. The "guardrails vs. leashes" framing recurred: regulations should move with technology, but cannot be written quickly enough, and deregulation could remove protections against real harms in domains where trust is already low. Some pointed to emerging approaches, such as principle-based guidance, regulatory sandboxes, and independent evaluation models like the Illinois Learning Technology Center,⁴¹ but no consensus formed on the right framework, the right level of governance, or the right pace of reform. This remains one of the session's most clearly identified challenges and least resolved.

4. Front-Line Practitioners Are Adopting Ahead of Their Institutions

Across domains, individual practitioners are moving faster than the institutions responsible for governing, training, and supporting them. Mechanics use ChatGPT for repair manuals before their managers are aware of it. Teachers engage with AI because students already use it, yet receive limited institutional support or training. Nonprofits often use unauthorized tools, or "shadow tech," because they cannot afford enterprise licenses but recognize the productivity gains. Table discussions noted that nonprofits are actually ahead of philanthropy in adoption, while funders themselves lack capacity to evaluate AI-related grants. This institutional lag extends into the knowledge pipeline itself: in higher education, faculty often lack the training to guide graduate students on AI, even as the PhD pathway shifts toward shorter-term, industry-oriented work.

This front-line-first dynamic has domain-specific implications. In healthcare, the "human in the loop" principle emerged as essential.⁴² Practitioners should accept AI-assisted decision-making only when a human clinician retains authority and accountability. At Johns Hopkins University,

⁴¹<https://www.ltcillinois.org/>

⁴²<https://pmc.ncbi.nlm.nih.gov/articles/PMC10280340/>

researchers drew a critical distinction between explainability (making outputs interpretable to users) and interpretability (understanding model internals), arguing the former is achievable and necessary for clinical trust.⁴³ In agriculture, ExtensionBot, an LLM built on over 360,000 curated extension publications with human-in-the-loop design, illustrates what domain-appropriate AI tools look like when practitioners drive development.⁴⁴ In industry, the Small Business Digital Alliance frames engagement around practical use cases, with funded counselors meeting owners where they are.⁴⁵ The pattern that emerged across sectors is that adoption succeeds when it is practitioner-led and institutionally supported, but fails when institutions are absent from the process entirely.

5. Domain-Specific Models Reveal What Works—and What Doesn't

The session surfaced concrete models that demonstrate both effective and ineffective approaches to domain-specific AI readiness. In New York, Empire AI—a ten-institution consortium focused on GPU supercomputing—represents a unique, large-scale compute-first strategy for supporting AI research infrastructure.⁴⁶ The Massachusetts AI Hub demonstrates a slightly different approach. Rather than centering primarily on compute, the Hub is designed to pair expanded access to computing and shared data with support for research collaboration, startup growth, and small business adoption, alongside workforce development aligned with industry needs and a focus on delivering broad public benefit.⁴⁷

Effective models share common design principles across sectors. In governance contexts, structured community participation was central to building legitimacy. The City of Austin's AI Governance Resolution, passed unanimously in 2025 following "Community in the Loop" engagement sessions, illustrates how civic participation can support trust in AI policy.⁴⁸ The Harvard Berkman Klein Center's participatory AI governance initiatives—including policy research clinics that integrate community voice into AI system design—illustrate how principles such as these can be institutionalized.⁴⁹

In education, local actors are creating coordination mechanisms in the absence of clear system-wide guidance. In Washington, D.C., the Washington Leadership Academy has launched the DC AI Collaborative, a network to share implementation strategies and advocate for district-level policy clarity.⁵⁰ The effort reflects both rising demand for coordination and the ad hoc nature of current approaches.

In industry, effective entry points begin with defined operational needs. In banking, fraud detection surfaced as a clear value proposition for community institutions, yet the infrastructure and expertise required to deploy it remain concentrated: three core service providers serve more

⁴³<https://extension.org/2025/04/02/extensionbot-highlighted-in-national-cast-report-on-ai-in-agriculture/#:~:text=The%20Extension%20Foundation%20is%20proud,with%20local%20practices%20and%20conditions>

⁴⁴<https://extension.org/tools/extbot/>

⁴⁵<https://www.smallbusinessdigitalalliance.com/>

⁴⁶<https://www.empireai.edu/>

⁴⁷<https://aihub.masstech.org/>

⁴⁸<https://austin-ai.org/2025/06/06/austin55/>

⁴⁹<https://cyber.harvard.edu/topics/ethics-and-governance-ai>

⁵⁰<https://dccharters.org/blog/bringing-ai-equity-to-dcs-classrooms-how-wla-is-leading-the-way>

than 70 percent of depository institutions, shaping AI access for smaller banks.⁵¹ In agriculture, precision agriculture and field-level data analysis offer comparable entry points aligned with existing workflows.⁵²

Recommendations

R1. Establish Coordinating Infrastructure Within and Across Sectors

Federal agencies, foundations, and professional associations should create or strengthen coordinating bodies with dedicated funding and staffing to convene stakeholders, facilitate knowledge sharing, support peer learning networks, and commission sector-specific AI readiness research. The Massachusetts AI Hub, NSF AI Institutes, Big Data Hubs, and Extension system demonstrate viable models. The key design principle is sustained funding with intermediaries that have practitioner credibility. Coordination should be structured to surface what is working in specific domains and enable those lessons to travel.

R2. Fund Domain-Specific Data Infrastructure and Governance

NSF and federal partners should support the development of sector-specific data standards, governance frameworks, and data management practices aligned with each sector's regulatory environment as opposed to generic, one-size-fits-all initiatives. Priority investments should include FERPA-compliant data-sharing frameworks for education; privacy-preserving standards for healthcare; data provenance and cleaning capacity for agriculture and Extension; and practical data management support for small businesses.

AlphaFold's success demonstrates the payoff when robust data infrastructure precedes AI deployment, and the Open Knowledge Network offers an emerging cross-sector model. Under-resourced institutions require dedicated capacity-building in data management before AI tools can be deployed effectively.

R3. Support Institutional Catch-Up to Front-Line Adoption

Across sectors, practitioners are adopting AI tools faster than institutions can govern, train, or support them, thereby creating liability exposure, quality risks, and ungoverned "shadow tech" usage. Federal agencies and foundations should fund institutions to close this gap through domain-appropriate mechanisms: in healthcare, establishing clear accountability frameworks and the "doctor in the loop" governance principle; in education, providing implementation support rather than just guidance documents; in nonprofits, subsidizing enterprise-grade tools with implementation support to replace unauthorized workarounds; in small business, expanding intermediary models like the Small Business Digital Alliance that meet practitioners where they are. Trust-building should be embedded in these efforts through community participation in AI governance and transparency about how AI systems work and who is accountable when they fail.

R4. Equip Existing Diffusion Infrastructure for Domain-Specific Needs

⁵¹<https://www.cgcatalyst.com/thought-leadership/commentary/ai-developments-and-community-bank-adaptations/>

⁵²<https://www.gao.gov/products/gao-24-105962>

The institutions best positioned to deliver AI readiness with broad access and at scale, such as universities, community colleges, Extension offices, and SBDCs, currently operate as general-purpose diffusion infrastructure. To be effective in an AI context, they need domain-specific resources, tools, and expertise that go beyond general AI literacy, backed by nation-scale AI competency frameworks and content.

Federal agencies and foundations should fund these institutions to build capacity tailored to the sectors they serve: fraud detection and AI governance for community banks, precision agriculture technologies, tools and data management for rural Extension offices, automation and supply chain optimization for small manufacturers, healthcare decision-support for community health networks. Without this domain-specific investment, diffusion infrastructure will default to least-common-denominator programming that fails to meet the actual readiness needs of practitioners identified across sectors in this session.

R5. Commission Sector-Specific Regulatory Reviews

Regulatory misalignment of existing frameworks for AI was one of the most consistently cited barriers in this session, but participants lacked consensus on ideal solutions. This reflects the genuine complexity of reforming frameworks like FERPA and the Health Insurance Portability and Accountability Act (HIPAA) that serve real protective functions even as they impede beneficial innovation.

State and federal agencies should conduct sector-specific regulatory reviews assessing where existing frameworks create friction without corresponding protection, where interpretive guidance (rather than statutory change) could resolve ambiguity, and where sandbox or pilot approaches could generate evidence about workable models. These reviews should involve both regulators and the domain practitioners who experience regulatory barriers firsthand: bank examiners alongside community bankers, school data officers alongside ed-tech providers, agricultural technology developers alongside Cooperative Extension agents.

5.4. Session Four: AI Literacy and Learning Pathways

Session Overview

Building an AI-ready America means equipping not just students but workers, small business owners, local leaders, and community members with the knowledge to use these tools effectively. This session focused on what it takes to build sustainable training infrastructure across these populations, who is positioned to train the trainers, and how to design learning pathways that remain relevant as the technology evolves.

Findings

1. Train-the-Trainer Models are Needed at the National Level

The infrastructure required to cultivate AI educators across the country does not yet exist. AI educators are largely self-taught or learn through local professional development opportunities.

In K-12, formal training is sparse and often reliant on motivated individuals seeking it out independently to bring it back to their schools. The result is a champion-dependent model rather than an institutional one. This dynamic places the burden of both technological adoption and change management on individuals, rather than on school or district leadership. The pattern extends beyond the K-12 system—participants described how workforce trainers, community college faculty, and Cooperative Extension agents are generally learning AI tools ad hoc, without support or recognized credentials for what they are teaching.

A lightning talk delivered by the Executive Director of the Center for Applied Data Science and Analytics at Howard University illustrated the potential of institutional approaches: having an executive-level, university-wide council to coordinate and implement AI strategies across the university can accelerate the implementation of AI in curriculum, upskilling of staff and faculty, and innovation in research. These efforts help to inform AI literacy at a national scale through the NSF-funded Research Coordination Network on Assessing and Predicting Jobs Outcomes in AI that brings forth much-needed conversations connecting curriculum to AI jobs.⁵³ While these efforts are commendable, more needs to be done at colleges and universities nation-wide to ensure that we have AI-ready talent to fuel the workforce.

2. AI Credentials Proliferate Without Quality Control or Labor Market Alignment

AI literacy credentials, such as industry certificates and micro-credentials, are proliferating without standardized quality assurance, clear learning outcomes, or recognized labor market value. Many programs are abstract and inapplicable without value propositions tied to the job market.

The cybersecurity sector demonstrates how an analogous field has addressed this gap. CyberSeek.org tracks jobs by state and maps skill gaps in real time;⁵⁴ the National Initiative for Cybersecurity Education (NICE) Framework maps skills to workforce roles;⁵⁵ CyberCorps Scholarship for Service links education funding to federal employment.⁵⁶ The banking sector is also developing meaningful certificates. The tension between flexible credentials and stable, quality frameworks remains unresolved—but the cybersecurity precedent demonstrates that resolution is achievable.

3. Digital Divide and Trust Barriers Limit Reach

Current AI literacy and learning efforts do not reach significant populations, and the barriers are structural. Roughly one in ten U.S. adults do not own a smartphone, and an additional 15 percent are without home broadband.⁵⁷ Without reliable Internet access, AI tools are simply out of reach.

The credibility of the messenger determines whether AI literacy efforts gain any traction at all. The implication is that scaling AI education requires identifying and supporting people who already hold community trust, such as educators, faith leaders, local organizers, and Cooperative

⁵³<https://www.aijobsrncn.org/>

⁵⁴<https://www.cyberseek.org/>

⁵⁵<https://niccs.cisa.gov/tools/nice-framework>

⁵⁶<https://sfs.opm.gov/>

⁵⁷<https://www.pewresearch.org/internet/2024/01/31/americans-use-of-mobile-technology-and-home-broadband/>

Extension agents. AI literacy programs should equip them with AI knowledge, rather than parachute in technical experts into unfamiliar settings. Programs that skip this step underperform and reinforce the skepticism they were meant to address.

4. Existing Curricula Are Disconnected from Practical Utility

Across learning pathways, some existing AI curricula emphasize explaining AI abstractly rather than why it matters in the context of learners' work and lives. Participants consistently noted that much of the available AI training is too generic and unconnected to learners' actual workflows or sector-specific challenges. A National Institute of Food and Agriculture (NIFA) presentation illustrated how the Cooperative Extension System naturally begins from community needs rather than technology. The National AI Institute for Exceptional Education offers an additional model, where learners with social-emotional development needs across K–12, shape both how AI tools are designed and the literacy educators need to deploy them responsibly.

5. Traditional Curriculum Cycles Cannot Keep Pace with AI's Rate of Change

Even where AI curricula exist, they are built on development and approval timelines that are oftentimes fundamentally mismatched with the speed at which AI tools and capabilities evolve. Participants noted that programs risk teaching frameworks and tools that are already outdated by the time students encounter them. The problem is structural: conventional curriculum governance treats content as stable, but AI content is not. Several participants pointed to modular, rapidly updatable approaches as a necessary alternative—short units that can be revised or replaced independently without redesigning entire programs. Graduate student-led workshops and industry advisory partnerships were cited as mechanisms for keeping content current without requiring permanent faculty to retool continuously. Without a shift toward curriculum architectures designed for change, AI education will remain perpetually behind the technology it aims to teach.

Recommendations

R1. Establish a National AI Educator Training Consortium

Federal agencies should fund a consortium connecting K-12, higher education, community colleges, universities, workforce boards, and community organizations around AI educator development. The NICE Framework provides a template for mapping skills to workforce roles. The consortium should include summer professional development with rapid curriculum update mechanisms, for example, incorporating graduate student-led workshop models and culturally embedded educator approaches.

R2. Develop a Framework for AI Literacy Credentials

Accreditation bodies, higher education, and industry partners should collaborate on a tiered credential structure spanning completion certificates, micro-credentials, and degrees, with learning outcomes mapped to labor market needs. The framework should include independent quality assurance mechanisms, regular curriculum refresh cycles, and employer recognition

requirements. Cybersecurity's professionalized credentialing demonstrates that meaningful credentials require both educational rigor and labor market alignment.

At the same time, federal entities such as NSF and the Department of Education should support foundational research to establish standards for K–12 AI education and developmentally appropriate instructional approaches. At present, there is little evidence-based guidance on what K–12 AI literacy should meaningfully entail. Advancing research in this area is essential to ensure that emerging standards are pedagogically sound, age-appropriate, and grounded in cognitive and educational science rather than market momentum alone.

R3. Equip Community Colleges and the Cooperative Extension System with AI Teaching Capacity

Earlier sections of this report identify community colleges and the Cooperative Extension System as primary delivery infrastructure for AI diffusion. The barrier is not whether to use these systems but whether they are AI-ready. Most lack faculty with AI expertise, current curricula, or resources to develop either. Federal agencies should fund AI faculty lines, curriculum development grants, and structured employer partnerships at community colleges, using models like the Alamo Community Colleges District, which integrates two-year degrees, 2+2 university pathways, and on-the-job training.⁵⁸ USDA and NSF should jointly fund AI upskilling for county-level agents so that the system's existing reach translates into AI literacy delivery.

R4. Support Trust-Centered Delivery for Under-Resourced Communities

AI literacy programs require trusted intermediaries. Federal agencies and foundations should support models that identify and equip people who already hold community credibility, such as educators from the communities they serve, faith leaders, local organizers, and Cooperative Extension agents, rather than defaulting to outside technical experts. NSF and philanthropic partners should support pilot programs testing community-driven AI literacy models with evaluation frameworks that measure trust and sustained engagement, not just enrollment.

R5. Build Mechanisms for Continuous AI Curriculum Renewal

AI capabilities are changing faster than traditional curriculum development cycles can accommodate. Programs designed on two- to four-year revision timelines risk teaching tools and frameworks that are already obsolete. Education institutions and credentialing bodies should adopt modular curriculum architectures that are composed of short, self-contained units that can be swapped, updated, or retired independently without redesigning entire programs.

At the same time, policymakers should invest in rigorous research on how AI is affecting learning outcomes, cognitive development, skill acquisition, and instructional practice. Curriculum modernization should not rely on assumption alone. As frameworks and training models are deployed, federal agencies should fund parallel research efforts to evaluate what improves learning, what undermines foundational skill development, and where guardrails are needed.

⁵⁸<https://www.alamo.edu/>

NSF and the Department of Education should support the development of shared, openly licensed, evidence-based curriculum modules with built-in revision mechanisms and embedded evaluation processes. Federally funded AI education programs should be required to include plans for ongoing content updates, continuous assessment, and public reporting of outcomes. Graduate student-led workshop models and industry advisory boards offer two mechanisms for keeping content current without placing the full burden on permanent faculty.

5.5. Conclusion

The United States possesses the requisite institutional infrastructure to support broad-based AI adoption. However, it lacks the investment, coordination, and intentional design needed to effectively leverage these resources. Across four sessions spanning state and local coordination, nation-scale efforts, domain-specific strategies, and literacy and learning pathways, workshop participants consistently identified the same structural gaps: under-resourced intermediaries, fragmented coordination, lack of shared standards, and training that is disconnected from practical application. These are capacity gaps with known precedents. The U.S. has built diffusion systems for previous generations of technology with great success. Many of the lessons from those efforts can and should be directly applied to making America AI ready.

The recommendations in this report are informed by people with direct experience implementing AI readiness efforts across sectors, institutions, and communities. They point toward a consistent operational model: sustained public investment channeled through trusted local institutions, tailored to the specific conditions of the communities and sectors they serve, and measured against outcomes that reflect whether AI adoption is reaching all Americans—not just those already positioned to benefit.

6. Statement of Principles

These principles reflect feedback from leaders in government, industry, philanthropy, education, and civil society on what it will take to build durable national infrastructure for AI diffusion, access, and adoption.

These are intended to guide federal agencies, state and local governments, philanthropic foundations, educational institutions, and industry partners as they design and fund AI readiness initiatives.

1. *Build on existing infrastructure—don't duplicate it.*

The United States already has distributed, community-embedded networks, such as the Cooperative Extension System, Small Business Development Centers (SBDCs), libraries, and universities and community colleges, with decades of trust and reach. Investment should strengthen capacity within these systems rather than create parallel structures.

2. *Start with the problem rather than the technology.*

AI adoption is most effective when programs begin with clearly defined community or institutional challenges and define AI's role in addressing them. Technology-first strategies risk producing “solutions in search of problems,” whereas problem-driven approaches generate clearer use cases, stronger buy-in, and more durable outcomes.

3. Invest in the multipliers.

Practitioners and intermediaries are the primary channels for AI diffusion and investment in them can have outsized impact on a community's AI readiness. Teachers, Cooperative Extension agents, SBDC counselors, higher education faculty, and workforce trainers shape how communities engage with AI. Their training, compensation, and sustained capacity should be a funding priority.

4. Design for broad access from the outset.

Broad access requires intentional program design, not supplemental outreach. Without deliberate mechanisms to reach under-resourced communities, AI adoption will compound existing disparities. Who benefits is determined by program architecture.

5. Treat AI adoption as a whole-economy challenge.

AI adoption is occurring across agriculture, healthcare, manufacturing, services, skilled trades, and small business—not only within the technology sector. Workforce and literacy strategies should therefore embed AI capabilities within the institutions and occupations wherever economic activity occurs, rather than concentrating investment solely in technical talent pipelines.

6. Establish shared vocabulary to enable coordination.

Shared standards enable coordinated action across sectors and regions. Common definitions, curriculum frameworks, and credential standards allow institutions to align without top-down mandates. The absence of shared vocabulary is itself a barrier to progress.

7. Structure public-private partnerships around outcomes.

Public-private partnerships should be structured around outcomes, not volume of tool adoption. Partnerships organized around literacy, economic resilience, and measurable outcomes are more sustainable than those oriented toward specific platforms.

8. Strengthen data infrastructure to more effectively scale AI.

Data infrastructure and governance are prerequisites for AI readiness. AI systems depend on clean, standardized, accessible data. Investment in data infrastructure and governance should precede or accompany investments in AI deployments.

9. Measure access, adoption, and impact systematically.

Systematic measurement of AI access, adoption, and impact is essential to informed policymaking. Regular reporting on adoption rates, talent pipeline metrics, and benefit distribution, modeled on the NSF Science and Engineering Indicators, would provide the data needed to guide resource allocation.

10. *Durable AI readiness requires sustained, multi-level investment with clearly defined roles.*

Government, philanthropic, and industry funding should prioritize building long-term institutional capacity rather than one-time project support. Federal efforts should prioritize enabling programs to mature rather than wither after short project-cycle grants. State investments should complement federal commitments to support local implementation. Philanthropic investment should prioritize coordination and strategic gap-filling that public funding cannot address quickly. Industry resources should preserve institutional independence and alignment with the public interest.

7. Roadmap for Future Action

The recommendations in this report point to five areas where follow-on action—from government, industry, academia, philanthropy, and civil society—is most needed.

1. Build the Knowledge and Workforce Pipeline

The “train-the-trainer gap” is a persistent barrier to AI diffusion across all sectors. A federally funded, cross-sector initiative should bring together teacher preparation programs, university and community college systems, the Cooperative Extension System, workforce boards, and credentialing bodies to design a national AI educator training pipeline with pathways for both formal educators and informal intermediaries. Concurrent with this, federal and state agencies should also support the development of a structured process to establish quality standards for the rapidly proliferating landscape of AI literacy credentials, with learning outcomes mapped to labor market needs and built-in mechanisms for curriculum refresh.

2. Establish Shared Definitions and Standards

The absence of shared vocabulary for AI readiness is itself a coordination barrier. A multi-stakeholder federal initiative including NSF, the Office of Science and Technology Policy (OSTP), the Department of Education (ED), and the Department of Labor (DOL) should facilitate the development of shared definitions, curriculum standards, and voluntary certification pathways for AI literacy across K-12, higher education, workforce development, and small business contexts.

3. Connect Practitioners and Institutions

State and local AI efforts remain fragmented both across sectors and within them, with effective models largely inaccessible to practitioners outside their immediate networks. Relevant federal agencies, including NSF and USDA NIFA, should develop a pilot program to experiment with

and validate models for cross-sector convening, resource matching, and knowledge exchange about how to best scale effective AI readiness programs.

NSF should develop principles and best practices for public-private partnerships that support AI readiness efforts, including guidance on intellectual property, exclusivity, platform independence, and public-interest alignment.

States should develop structured readiness planning processes, including assessing baseline connectivity, institutional capacity, and workforce readiness, paired with technical assistance and cross-state learning to guide AI readiness efforts.

4. Develop National AI Readiness Indicators

No systematic mechanisms currently exist for tracking AI adoption rates, workforce readiness, or access distribution. NSF should support research to develop a national AI readiness measurement framework. This effort should identify which indicators meaningfully capture AI readiness across sectors and geographies, what collection mechanisms are needed, and what reporting cycles would be most useful, with the goal of establishing biennial reporting aligned with NSF's Science and Engineering Indicators. The resulting framework should generate both the metrics themselves and promising practices for how institutions at different scales can assess and communicate their own readiness.

One particularly useful indicator of AI readiness identified by workshop participants is access to data and compute. NSF should support a scoping study that evaluates compute and data capacity gaps across the country, evaluates shared infrastructure models, and develops recommendations to close these gaps.

5. Align AI Governance Across Sectors

Inconsistent guidelines, fragmented policies, and a lack of clarity around AI governance were among the consistently cited barriers across domain-specific sessions, from education to industry to agriculture. Practitioners operate under patchwork-like frameworks that vary not only across sectors but within them, making confident AI adoption difficult at scale. Federal and state agencies should establish cross-sector working groups charged with developing clear, implementable governance frameworks aligned across domains. These working groups should be cross-disciplinary by design, bringing together the people actually deploying AI tools—educators, community bankers, Cooperative Extension agents, local government staff, small business operators—alongside policymakers, regulators, and technologists. Investment should support a structured process that produces practical, sector-tested governance guidance rather than abstract principles, including pilot programs that test framework alignment across two or more sectors simultaneously.

8. Acknowledgements

This report was prepared by SeedAI with support from the U.S. National Science Foundation under Award No. 2608403 and the Alfred P. Sloan Foundation under Award No. G-2025-79265. The opinions, findings, and conclusions or recommendations expressed in this material are informed by the contributions and perspectives of the workshop participants, and do not necessarily reflect the views of the U.S. National Science Foundation or the Alfred P. Sloan Foundation.

We sincerely thank the following individuals for their comments on this report: Adam Browning (Washington Leadership Academy), An-Me Chung (New America Foundation), Antonio DelGado Fornaguera (Miami Dade College), Florence Hudson (Northeast Big Data Innovation Hub, Data Science Institute, Columbia University), Christine Kirkpatrick (San Diego Supercomputer Center), Chad Lane (University of Illinois, Urbana-Champaign), Danielle S. McNamara (Arizona State University), Meme Styles (MEASURE), Talitha Washington (Center for Applied Data Science and Analytics, Howard University), and Julia Wynn (Code.org).