In Ctrl of AI in Schools

What existing research says about safe and effective AI systems in education and the policy opportunities for Texas classrooms

Keep students safe

Support teachers

Strengthen learning



RENZO SOTO

Deputy Director, Education and Workforce Policy

November 10, 2026

TEXAS 2036

Table of Contents

Executive Summary	**1**
Introduction: AI Is Already in Texas Classrooms. How It Will Impact Students Depends on Conditions Texas Can Shape.	**5**
1. State Governance: Texas Has Built a Legal Foundation for Responsible AI Use. Now It Must Build on That Foundation for Schools.	**10**
2. Student Safety: Privacy and Developmental Risks Are Often Built into Tools Before They Reach a Classroom. Prioritizing Protections Should be the Standard.	**17**
3. Student Learning: The Evidence on AI and Student Learning Points to Genuinely Mixed Outcomes, Highlighting the Need to Govern its Use.	**26**
4. Educator AI Literacy: Educators Need Foundational AI Knowledge to Protect Students, Whether or Not They Use AI Themselves	**36**
5. Student AI Literacy: Students Can Be Educated and Empowered to Effectively Protect Their Own Safety and Learning	**46**
Conclusion: Who Is In Control of AI in Texas Schools	**57**
About This Report	**58**

Executive Summary

Artificial intelligence is already in Texas classrooms, and its use is largely unguided. In fall 2025, 86% of Texas school and district leaders reported that their teachers were using AI, and 61% reported student use.[1] However, only 29% of those leaders reported having a district policy governing student AI use, and only 20% reported one governing teacher and staff use.[2] Texas students and educators are navigating a fast-moving, highly customizable technology without consistent direction on what responsible use looks like.

At the same time, the technology has permeated local and global economies alike, and its evolving capabilities coupled with its tremendous rate of adoption suggest AI is here to stay. This means students and educators alike will still encounter this technology and its outputs regardless of active efforts to eschew it. **Thus, there is an urgent need to adapt Texas schools and the systems supporting them so that student safety and learning remain paramount**.

Texas students and educators are navigating a fast-moving, highly customizable technology without consistent direction on what responsible use looks like.

The Role of the State

This report finds that **AI's effects on students can be shaped by human decisions**, which state policy can help steer. The research shows that whether AI's risks are prevented and its potential benefits realized tracks with conditions that can be identified and governed: how the tool is designed, how well students and educators understand it, and how schools and the state govern its use.

[1] R. Torbey et al., "Leading the Way in K-12 Education: Recommendations of the Texas AI in Education Task Force," June 17, 2026.
[2] Ibid.

AI in Schools: By the Numbers

82%
of U.S. public school teachers report no formal guidance on applying AI to their work

8 in 10
U.S. students report that no teacher has taught them how to use AI for schoolwork

84.5%
of Texas voters rate protecting student data and privacy as extremely or very important if schools use AI

Sources: Gallup & Walton Family Foundation Educator Research, "Teaching for Tomorrow: Closing the Expectations Gap," 2026. RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind: Findings from the RAND Survey Panels," Sept. 30, 2025. Texas 2036, "10th Texas Voter Poll," Aug. 22–26, 2026. Conducted by Baselice & Associates; 1,369 completed interviews; effective sample size after weighting: 1,001; margin of sampling error: ±3.1 percentage points.

It is important to understand that this report does not contemplate whether AI systems should be used in public schools at all. The evidence on AI's impact on student learning and educator efficacy does not yet offer a clear direction on whether it definitively helps or harms.[3] However, this has not stopped the rate at which its use is growing across the state.

Amid this growing adoption, this report acknowledges and focuses on the fact that critical questions around the safety and efficacy of these tools remain. As such, it attempts to answer the second-order question to inform state policy action: now that AI is in Texas public schools, how do we make sure that these systems are safe and effective?

Texas has already built the legal foundation for this work. Three laws passed by the 89th Legislature in 2025 — House Bill 149, Senate Bill 1964, and HB 3512 — set responsible governance as the baseline for AI use across the public sector, including school districts.

However, these laws were written for a broad range of state agencies and local government entities, not for the classroom, where the stakes for students' safety, development and learning are distinct and higher. Closing that gap is the state's role.

[3] Throughout this report, there are references to artificial intelligence and generative artificial intelligence systems and tools. Specific references to "generative artificial intelligence" or "generative AI" are used to describe studies, reports and surveys with fidelity to the verbiage used within those documents.

Early Insights from Existing Research

SAFETY

On safety, many of the risks students face are built into an AI tool before it ever reaches a classroom, such as how it handles student data and security and how it affects students' development and well-being. A combination of proactive safety-forward design choices, active monitoring, and regular evaluation can help schools address these risks.

LEARNING

On learning, the evidence is genuinely mixed. There are documented gains, harms, and sometimes both to student outcomes, and they often depend on various circumstances surrounding AI use. Where learning improved, AI use came with protections such as safeguards built into the tool, educator oversight, and teaching methods that kept students actively engaged.

TEACHERS

On AI use by teachers, there are real documented benefits in areas like efficiency and instructional improvements. Educators and how they guide instruction also help make the difference between positive and negative academic outcomes for students where AI is used. But the research base on teacher AI literacy is lacking in quantity and quality, and there are similar limitations on the long-term impact of AI use on educator practice. This elevates the importance of maintaining human judgment within instruction.

STUDENTS

On AI use by students (referred to as student AI literacy), students' ability to critically and creatively think and manage their own learning emerge as leading contributors to realizing AI's potential benefits to advance outcomes and protect against cognitive offloading. Building these skills require comprehensive, age-appropriate, and cross-curricular AI literacy efforts.

OUR APPROACH

This report does not contemplate putting AI in every classroom, nor does it contemplate keeping it out entirely. *The evidence supports neither forcing these tools on students and educators nor ignoring a technology that students are already using on their own****. It supports a third path: preparing Texas schools to engage with AI responsibly, clear-eyed about both its risks and its benefits, and anchored in the principle that where AI use is appropriate, student safety and learning come first.*** *This third path is key to ensuring Texas schools are ready to critically engage with AI as the technology and its outputs pervade society.*

Recommendations for Texas

Four conditions recur across the evidence as the levers most within the state's control. To support the responsible use of AI in Texas schools, the state should:

1	Establish a common, evidence-based baseline for responsible use, translating Texas' existing AI-governance laws into the K–12 context so every district starts from the same foundation.
2	Ensure the AI tools reaching Texas classrooms are safe, secure and effective before they ever enter the classroom.
3	Equip educators with the foundational knowledge to recognize AI's risks and respond to them, regardless of whether an educator chooses to use AI in their own practice.
4	Prepare students to engage with AI safely and critically, building the foundational skills that make them resilient in an AI-integrated world.

Introduction:

AI Is Already in Texas Classrooms. How It Will Impact Students Depends on Conditions Texas Can Shape.

The wide availability of artificial intelligence (AI) has pushed the technology into nearly every corner of daily life, and Texas has acutely felt its impact. The state had the second highest number of estimated net tech employment in 2025, and AI skills showed up in over 90,000 Texas job postings.[4] Business adoption is climbing, with 36% of Dallas Fed survey respondents reporting their use of generative AI, or a mix of generative and traditional AI, up from 20% from the prior year, and another 19% stated they plan to adopt within 12 months.[5] Texans themselves are using generative AI, with 40% of Texas adults reporting at least a somewhat often level of usage, even though 30% of respondents also reported negative feelings toward AI.[6] **AI is here to stay. Industries across the economy are racing to adapt to it. Schools are no exception.**

AI Has Moved from Novelty to Workforce Fact in Texas

2nd highest estimated net tech employment of any U.S. state, 2025	**90,000+** Texas job postings in 2025 listed AI skills	**20% → 36%** of Texas businesses use generative AI, up from the prior year Another 19% plan to adopt within 12 months	**40%** of Texas adults use generative AI at least somewhat often	**30%** of Texas adults also reported negative feelings toward AI

Sources: CompTIA, "State of the Tech Workforce 2026," Mar. 24, 2026. Federal Reserve Bank of Dallas, "Texas Firms Open to AI as Tariff Work-Around Strategy," July 17, 2025. Texas Lyceum, "The 2025 Texas Lyceum Poll," Apr. 29, 2025.

[4] CompTIA, "State of the Tech Workforce 2026," Mar. 24, 2026.
[5] Federal Reserve Bank of Dallas, "Texas Firms Open to AI as Tariff Work-Around Strategy," July 17, 2025.
[6] Texas Lyceum, "The 2025 Texas Lyceum Poll," Apr. 29, 2025.

The State of AI in Texas Classrooms

Students and teachers are using these tools now. Eighty-six percent of Texas school and district leaders surveyed in fall 2025 indicated their teachers used AI in school, and 61% indicated student usage.[7] This tracks with national trends. In 2025, 54% of students reported AI use for schoolwork and 53% of English language arts, math and science teachers reported using AI for instructional planning or teaching, each an increase of more than 15 percentage points in just one to two years.[8]

FIGURE 1

Texas School Leaders Report Significant AI Use Among Teachers and Students

Share of Texas school and district leaders reporting AI use, fall 2025



Source: R. Torbey et al., "Leading the Way in K-12 Education: Recommendations of the Texas AI in Education Task Force," June 17, 2026.

As Texas grapples with the proliferation of these tools in education, the central question arises: what can be done to make sure AI systems used in schools are safe and effective for Texas students?

Importantly, it should be recognized that the above is second order to the question of whether AI should be used in Texas classrooms at all. This is because the initial research base has produced mixed results on AI's impacts in education, making it difficult to answer with an evidence-based approach. There are real benefits and real harms, but with the technology's growing prevalence in schools, it has become increasingly urgent to determine what is leading to the positive and negative outcomes. This report turns to ensuring the responsible use of AI by educators, students and the rest of the state's education system. This approach was chosen because it pragmatically prepares educators and students with the reality that AI use is rapidly

[7] R. Torbey et al., "Leading the Way in K-12 Education: Recommendations of the Texas AI in Education Task Force," June 17, 2026.
[8] RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind," Sept. 30, 2025.

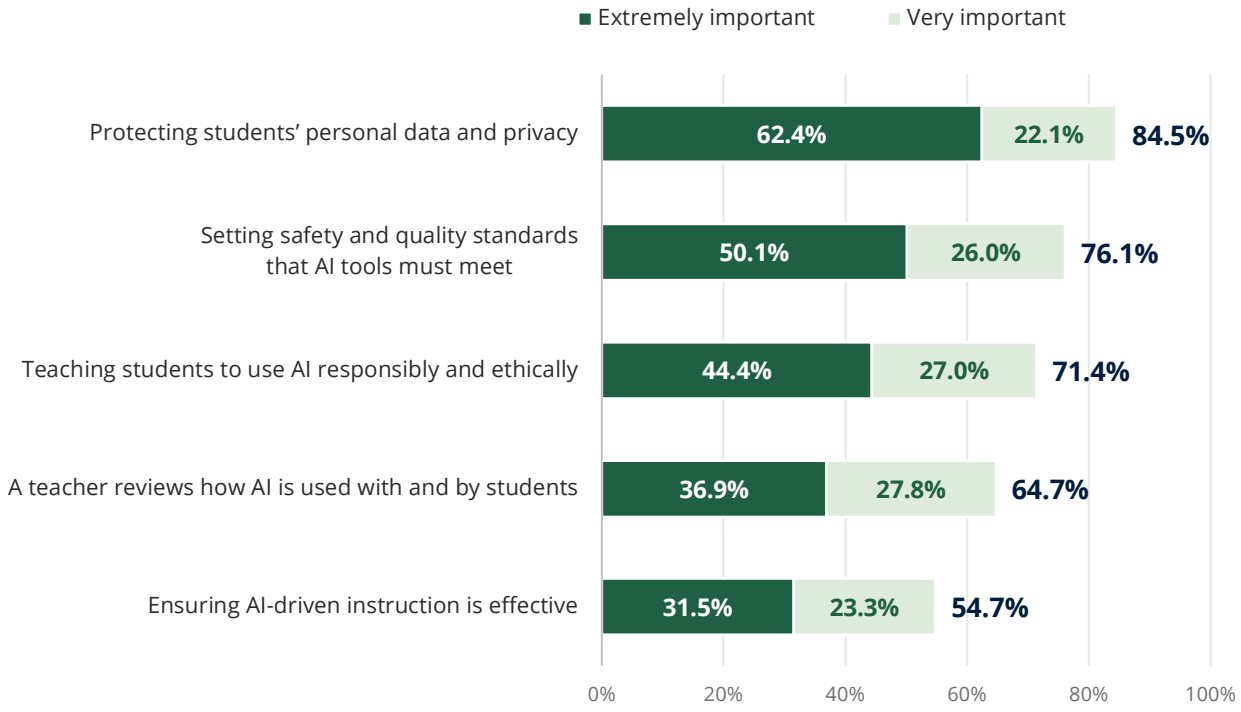
escalating in society, and it can be supported by existing evidence that equips state and local policymakers with the knowledge to navigate the technology's risks.

The key will be making sure these systems are grounded in responsibility: the state's responsibility to keep students safe in public schools, as well as school leaders and educators' responsibility to prepare students to succeed in life with a strong academic foundation and the skills that the workforce and citizenship demand. Texas voters agree as the majority believe that safety and efficacy are either extremely or very important when it comes to the use of AI in education, and student privacy is the top concern.

FIGURE 2

Texas Voters Prioritize Privacy and Safety When Evaluating AI in Public Schools

Q: "Please tell me if you think each of the following is extremely important, very important, somewhat important or not important if Texas public schools use AI, artificial intelligence."



Source: Texas 2036, "10th Texas Voter Poll," Aug. 22–26, 2026. Conducted by Baselice & Associates; 1,369 completed interviews; effective sample size after weighting: 1,001; margin of sampling error: ±3.1 percentage points.

Current use, however, is largely unguided. Districts, administrators, educators and students are navigating a fast-moving, highly customizable technology without consistent direction on what responsible use looks like.

Only 29% of Texas school leaders reported a district-level policy on student AI use, while only 20% reported a district-level policy on teacher and staff use.[9] This means most districts have no written rules for how students, teachers and staff may use AI. The result is a patchwork of practice that varies from district to district and classroom to classroom, with uneven attention to the two things that matter most: students' safety and learning.

> **Texas governance of AI in education is currently a patchwork of practice that varies from district to district and classroom to classroom, with uneven attention to the two things that matter most: students' safety and learning.**

Controlling the Trajectory of AI in Education

This report draws on a broad scan of the research on AI in education, anchored in meta-analyses and research syntheses and supplemented by individual studies, field surveys, and policy and practice reports, with causal research weighted most heavily. The causal evidence remains thin, and much of it comes from outside the United States and from higher education settings. However, it can still inform Texas policy because the conditions it identifies, such as how a tool is designed and how well educators and students are prepared to use it, can be influenced by state policy. This report synthesizes the studies' evidence, analyses and conclusions with sources and limitations detailed in the "About This Report: Sources and Limitations" section.

Read together, that evidence points to a consistent conclusion: **AI's effects on students can be shaped by human decisions.**

State policy plays an important role in shaping whether these tools will help or harm. **AI's impact on students will not solely be determined by the technology on its own, or by the students or educators who use it, but by a combination of conditions surrounding both — how the tools are developed, how well students and educators understand them, and how schools and the state govern their use.** Those conditions, when shaped appropriately, are what can help build resilience. They protect students from the risks of the technology while preparing them to learn and grow in a world where AI is becoming more present.

[9] R. Torbey et al., "Leading the Way in K-12 Education: Recommendations of the Texas AI in Education Task Force," June 17, 2026.

This points to a clear path forward. This report does not contemplate putting AI in every classroom, nor insulating students from the issue entirely. It does not support forcing these tools on students and educators or ignoring technology students already use on their own. The report considers a third path: preparing Texas schools to engage with AI responsibly — clear-eyed about both its risks and its benefits and anchored in the principle that where AI use is appropriate, student safety and learning come first.

Across the evidence reviewed in this report, there are four recurring conditions that the state can influence. To support the responsible use of AI in Texas schools, the state should:

1. **Establish a common, evidence-based baseline for responsible use**, translating Texas' existing AI-governance laws into the K–12 context so every district starts from the same foundation.
2. **Ensure the AI tools reaching Texas classrooms are safe, secure and effective** before they ever enter the classroom.
3. **Equip educators with the foundational knowledge to recognize AI's risks and respond to them**, regardless of whether an educator chooses to use AI in their own practice.
4. **Prepare students to engage with AI safely and critically**, building the foundational skills that make them resilient in an AI-integrated world.

The report proceeds in five sections. **Section 1** takes up the state's legal framework and the guidance schools still lack. **Section 2** examines the risks students face regarding data privacy and their development and well-being. **Section 3** reviews the evidence on whether AI helps students learn and what that evidence means for the tools reaching classrooms. **Sections 4 and 5** turn to the educators and students who may use the tools. Each of the sections opens with a plain-language summary of the issue and the evidence, then examines the research in detail, and follows with an analysis of what it means for Texas

Section 1.

State Governance: Texas Has Built a Legal Foundation for Responsible AI Use. Now It Must Build on That Foundation for Schools.

THE ISSUE

Texas passed laws in 2025 that made responsible AI governance the standard for public sector entities, including school districts.[10] However, those laws were written for state agencies and local governments generally. They were not specifically translated into clear, actionable directions that fit the school and classroom contexts. Districts must work out on their own what the standards mean for a tool, an educator, or a student. Currently, there is no education-specific guidance that exists in Texas that will turn the state's governance framework into consistent practice across all school districts.

KEY TAKEAWAYS

- **Only 29% of Texas school and district leaders report a policy on student AI use**, and 20% report one for teachers and staff.[11] Nationally, policy adoption is rising but uneven, and district leaders are asking state education agencies for clearer direction on what responsible use means.[12]
- Parents and students were often **brought into AI adoption and use decisions late in some early-adopter districts**,[13] and parents worry more than district leaders do about AI's effect on students' critical thinking.[14]
- **42 states have issued statewide guidance**, model policies or frameworks for schools. Texas has not.[15]

[10] In 2025, the Texas Legislature passed H.B. 149, S.B. 1964, and H.B. 3512, regarding AI responsible use, governance for the public sector and literacy.
[11] R. Torbey et al., "Leading the Way in K-12 Education: Recommendations of the Texas AI in Education Task Force," June 17, 2026.
[12] RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind," Sept. 30, 2025.
[13] Center on Reinventing Public Education, "Early Adopter Districts and AI: Strategic Pathways, System Strain, and the Conditions for Amplifying Transformation," May 2026.
[14] RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind," Sept. 30, 2025.
[15] Texas 2036 analysis of publicly available state education agency documents on the use of AI in education.

The State's Baseline on Responsible AI Use

Texas moved early to govern AI. In its 2025 regular session, the 89th Legislature passed three complementary laws setting responsible governance as the baseline for AI use in the state. All three impact school districts.

FIGURE 3

Three Laws from 2025 Set the Baseline for Responsible AI Use in Texas



- **HB 149**, the Texas Responsible Artificial Intelligence Governance Act (TRAIGA), regulates all AI systems developed and deployed within the state through an explicit articulation of what the AI cannot produce as an output or effectuate as an outcome for an individual. It prohibits a defined set of intentionally harmful AI practices, including systems built to unlawfully discriminate or to generate illegal content.[16]
- **SB 1964** directly governs how Texas public sector entities use AI, directing the Texas Department of Information Resources (DIR) to establish a statewide code of ethics and governance standards that state agencies and local governments must adopt, including public independent school districts and charters. The code of ethics governs their procurement, development, deployment and use of AI generally, while the accompanying minimum risk management and governance standards apply specifically to "heightened scrutiny" AI systems (see definition in the box below).[17]

Heightened Scrutiny AI Systems

Generally, heightened scrutiny AI systems are those that are autonomous and have significant effect on a person's access to a government service

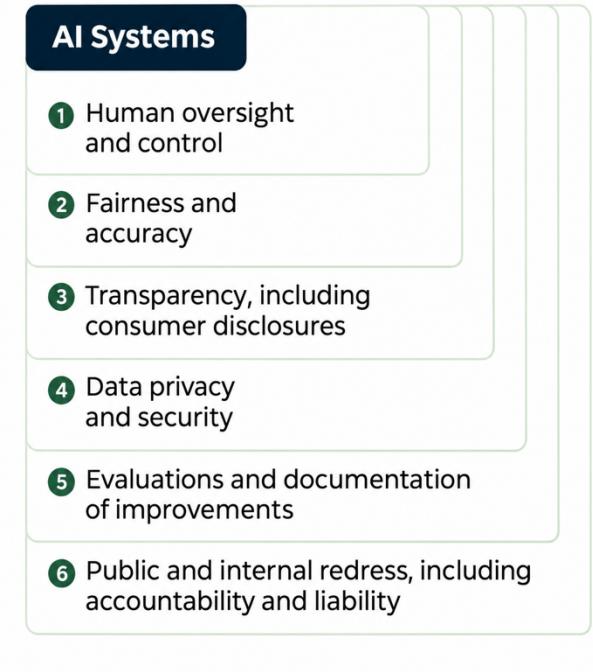
[16] H.B. 149, 89th Leg., Reg. Sess. (Tex. 2025).
[17] S.B. 1964, 89th Leg., Reg. Sess. (Tex. 2025).

- **HB 3512** requires state-certified AI training for a broad range of public employees whose duties regularly involve computers.[18] For school districts, current guidance only requires the district's cybersecurity coordinator to complete the annual training requirement. Other school district employees can be required to complete it at the direction of the district.[19]

State Policy and Its Impact on Texas Schools

FIGURE 4

The Texas State Code of Ethics for AI



Source: Texas Department of Information Resources, "Artificial Intelligence System Code of Ethics," 1 Tex. Admin. Code § 219.11, effective Mar. 18, 2026.

Two of these laws are explicitly applicable to Texas school districts. As public entities, districts fall under SB 1964's governance framework, including the state's code of ethics, and HB 3512's training requirement. At the same time, TRAIGA's prohibitions set the outer limit for AI systems developed and deployed in the state, including those used by school systems. Consistent with the Legislature's expectations for the broad public sector, responsible AI governance and use is the standard for public schools.

But guidance is needed to make that standard actionable in a consistent manner across the state. The laws are broad, built for a wide swathe of agencies and local governments but not specifically for the classroom, where the stakes for students' safety, development and learning are distinct and higher.

Districts are left to turn broad governance standards into concrete decisions about the appropriateness of tools in a classroom and

[18] H.B. 3512, 89th Leg., Reg. Sess. (Tex. 2025).

[19] Texas Department of Information Resources, "Public Sector Artificial Intelligence (AI) Awareness Training," accessed Oct. 12, 2026, https://dir.texas.gov/ai-and-innovation/public-sector-artificial-intelligence-ai-awareness-training.

how students and educators should engage with them, largely on their own and with uneven capacity to do it well. The state's role is to close that gap.

> **The state's laws are broad, built for a wide swathe of agencies and local governments but not specifically for the classroom, where the stakes for students' safety, development, and learning are distinct and higher.**

Translating the state's broad standards into consistent practice is what will ensure every Texas student benefits from the same baseline. The current patchwork of policy is not unique to Texas. As early as fall 2023, only 5% of districts surveyed nationally had adopted a policy governing students' use of generative AI, consistent with the 90% of surveyed teachers who reported no such policy.[20]

There are indications of some improvement in overall adoption of AI-use policies, as a more recent survey found that 45% of principals reported school or district AI policy or guidance in place.[21] However, this does not seem to have tempered the call for clearer direction. In augmentative interviews with district leaders, some "expressed the need for clearer guidance from state education agencies, particularly on defining responsible AI use."[22]

Unguided AI Use and Its Impact on Communities

The impacts of this issue extend past the classroom into school districts' broader communities. As districts are trying to respond to AI, there are cases of parents and students being left out of the decisions made on their behalf. Among early AI-adopter districts, it was found that parental engagement often happened after addressing internal needs first, sometimes producing parental pushback later.[23]

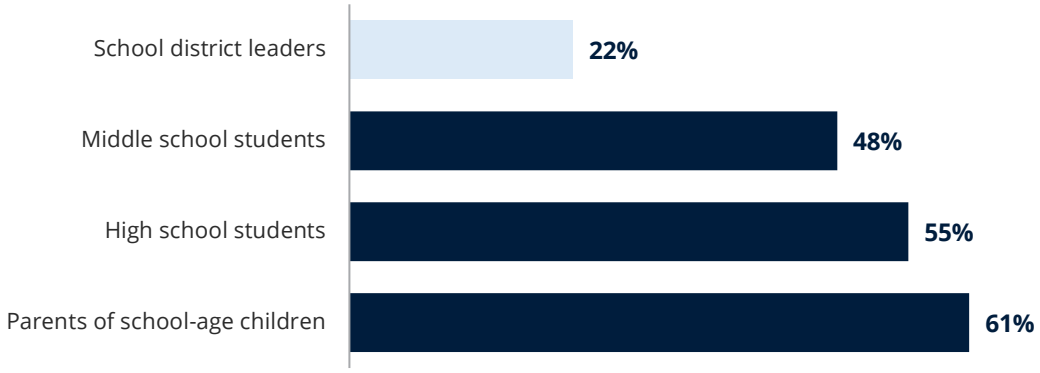
[20] RAND Corporation, "Using Artificial Intelligence Tools in K-12 Classrooms," Apr. 17, 2024.
[21] RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind," Sept. 30, 2025.
[22] Ibid.
[23] Center on Reinventing Public Education, "Early Adopter Districts and AI: Strategic Pathways, System Strain, and the Conditions for Amplifying Transformation," May 2026.

This is unsurprising given that parents and students have reported considerably more concern that AI could harm students' critical thinking skills than district leaders.[24] **If a district fails to meaningfully engage and involve families, even a genuinely useful tool will struggle to earn trust.**

FIGURE 5

Students and Parents Largely Agree that AI Could Harm Students' Critical Thinking

Share who agree or strongly agree that AI use could harm students' critical-thinking skills



Source: RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind: Findings from the RAND Survey Panels," Sept. 30, 2025.

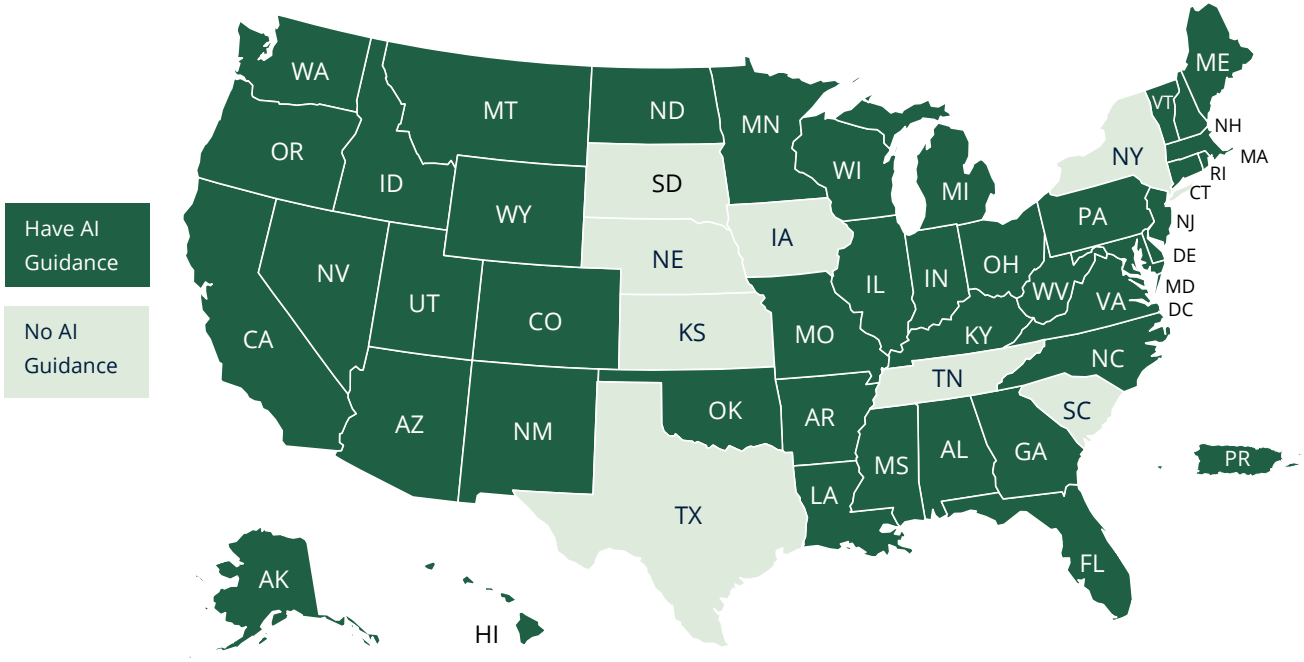
[24] RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind," Sept. 30, 2025.

The consequence is an increasing perception or, in many cases, the reality of navigating AI in education in isolation. That is why Texas must act. **Other states have already acted: 42 states have established some type of statewide guidance, model policy, framework, or other type of support for school districts.**[25]

Texas is currently not among those states, but it holds a uniquely advantageous position. SB 1964 already requires all school districts to abide by the public sector governance framework. While most state guidance is advisory or voluntary, guidance developed in Texas would set the foundation for compliance with existing law.

FIGURE 6

Texas Remains One of Eight States Without a Model Policy on AI Use in Education



Source: Texas 2036 analysis of publicly available state education agency documents on the use of AI in education.

[25] Texas 2036 analysis of publicly available state education agency documents on the use of AI in education.

What This Means for Texas

This means statewide guidance for schools in Texas, and the supports to implement it, could serve as the connective tissue between Texas' existing responsible use standards and best practices in AI in education.

In addition to establishing the framework, policymakers should seek to understand how AI in education is impacting students. This will help ground the framework on current evidence while helping account for future adjustments as the technology and the research evolves.

- **Develop statewide K-12 guidance** that builds on SB 1964's governance framework, focusing on translating the state code of ethics and minimum standards into actionable information relevant to the education context.
- Identify gaps and opportunities to **heighten the state's general standards** and ensure K-12 guidance prioritizes student safety, data privacy, and efficacy.
- **Supplement statewide K-12 guidance** with necessary implementation supports, including clearly defined appropriate and inappropriate uses, plain-language data privacy best practices and expectations, and other relevant resources tailored to be easily understood by practitioners and families.

Section 2.

Student Safety: Privacy and Developmental Risks Are Often Built into Tools Before They Reach a Classroom. Prioritizing Protections Should be the Standard.

THE ISSUE

AI tools can be adaptative and recursive in nature, so when educators and students provide them with information that could identify a student, the tools could have the ability to create new student records. **Federal privacy law protects these records, but school districts, rather than vendors, are the primary parties bound to ensure compliance**. Individual educators and school districts cannot consistently know whether a tool handles student data safely without careful vetting.

Without state-level standards and supports to ensure a baseline of privacy-forward design for AI tools in education, vetting will vary across districts and across classrooms, which increases the chances of some students left exposed to heightened data privacy risks over others.

The design of AI tools also impacts how it could affect the development of students. Some of the most accessible AI tools are intended for the broad consumer market rather than specific and tailored classroom use. The features of these tools, such as sycophantic responses or adaptivity despite incorrect information, can keep users engaged even if the outputs are inaccurate or potentially harmful.

Students are particularly vulnerable to these effects since their social and emotional capabilities are still developing. Design plays an important role in mitigating these risks as well, as classroom tools can be specifically tailored around protecting student agency over engagement.

KEY TAKEAWAYS

- **Privacy obligations rest on the school.** Districts are largely responsible for using tools in a compliant manner because no product can be FERPA-compliant on its own.[26]
- **Removing a student's name does not remove all data privacy risks that can surface from AI tools.** Student ID numbers, school email logins, and combinations of ordinary details can still identify a child.[27]
- **Tools can be designed to prioritize educational safety** and truth over engagement, allowing them to serve as reliable supports.[28]
- **District leaders nationally are asking for the exact support that the state can provide**, including model policy text, best practices for student safety and data privacy, and third-party vetting of tools for safety and reliability.[29]

The Complexity and Criticality of Data Privacy

As with any classroom technology, students' safety and learning must lead adoption decisions. School districts should emphasize this in the case of AI tools given the technology's documented risks and its rapid advancement that could raise more potential risks.

In the context of education, Texas voters are highlighting student data and privacy as a priority. In Texas 2036's 10th Annual Voter Poll, protecting student data and privacy was the standout concern when it came to AI use in the classroom.

[26] Future of Privacy Forum, "Vetting Generative AI Tools for Use in Schools," Apr. 2024.
[27] Ibid.
[28] EDSAFE AI Alliance, "S.A.F.E. by Design: Policy, Research, and Practice Recommendations for AI Companions in Education," Feb. 2, 2026.
[29] RAND Corporation, "Using Artificial Intelligence Tools in K-12 Classrooms," Apr. 17, 2024.

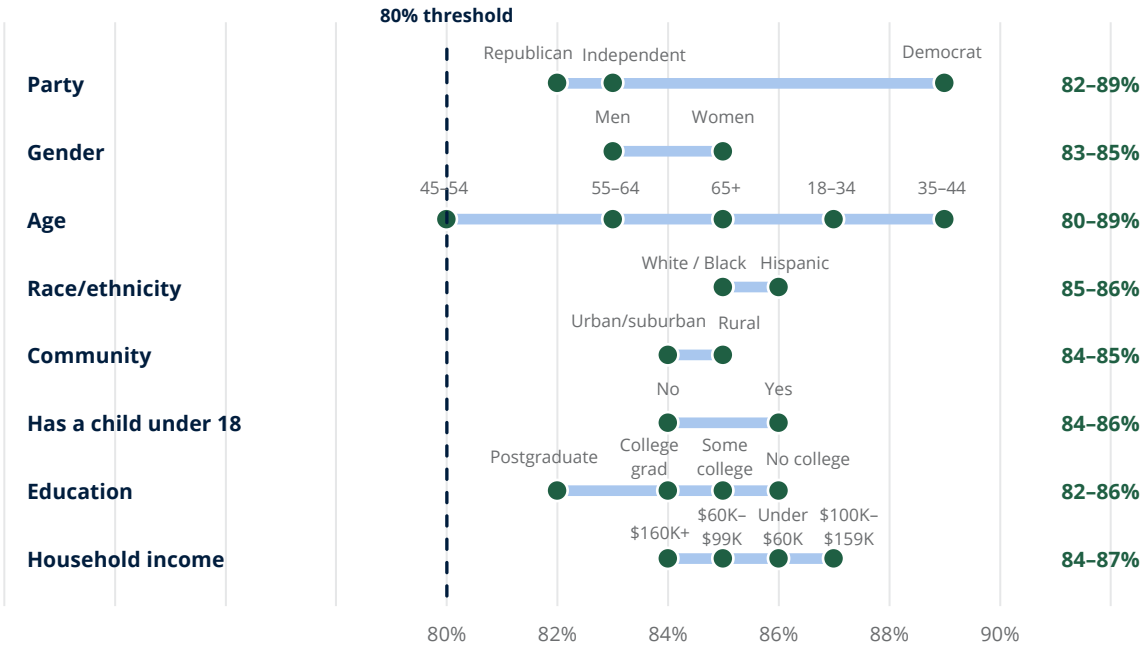
FIGURE 7

If Texas Public Schools Use AI, Voters Put Student Privacy First

Share of Texas voters rating "protecting students' personal data and privacy" as important



Every demographic below rates student privacy protection as extremely or very important at 80% or higher.



Source: Texas 2036, "10th Texas Voter Poll," Aug. 22–26, 2026. Conducted by Baselice & Associates; 1,369 completed interviews; effective sample size after weighting: 1,001; margin of sampling error: ±3.1 percentage points.

Despite data privacy laws like the federal Family Educational Rights and Privacy Act (FERPA), a tool's design may not preclude or mitigate all potential risky uses of student data.

For example, a tool built to reject direct student identifiers may still accept pseudonymized ones that can function as indirect identifiers and raise the risk of privacy violations. This, then, triggers a host of considerations that schools and educators must navigate.

FERPA in Practice: Vetting AI Tools for Classroom Use

The school, not the tool, carries the obligation. FERPA applies to schools that receive U.S. Department of Education funds, not to vendors, so no product can be "FERPA-compliant" on its own. Vetting asks whether a district can use the tool in a FERPA-compliant way.

Student Personally Identifying Information (PII) is broader than a name. Protected information includes direct identifiers like a name or student ID and indirect identifiers like a birth date, plus anything that alone or in combination would let a reasonable person in the school community link a record back to a student. Swapping a name for a student ID, or having students create accounts with school email addresses, does not de-identify the data.

New records carry the same duties. Where a tool's output creates or alters an education record, FERPA's rights to inspect, review and request correction still apply, which can be a hard standard to meet if a school considers the tool as a black box and hallucinations risk producing inaccurate records.

Sharing rests on the school official exception. Parental consent is generally required before student records can be shared with a vendor. However, a school may share student PII without parental consent in cases where the vendor performs a service that is "school official", which is a service the school would otherwise perform itself and remains under the school's direct control. This allowable exception could be considered by schools as they manage the flow of covered student information.

Product improvement is its own question. Some tools use what students enter to keep training the underlying model, while others train the model once and stop. Federal guidance allows a vendor to use student data, even in identifiable form, to improve the service it delivers to the school, but new products built on that data must rely on information with direct and indirect identifiers removed. Vetting asks whether the tool trains on student PII and, if it does, what else the vendor builds with it.

Source: Future of Privacy Forum, "Vetting Generative AI Tools for Use in Schools," Apr. 2024.

Fortunately, actions can be taken to ensure the safety risks educators and students can face with AI use are mitigated through responsible initial decisions. Strong procurement practices and processes can help from the onset. That can involve identifying intended use cases and whether student PII is involved, which in turn drives an in-depth legal compliance review against applicable privacy laws.[30]

[30] Future of Privacy Forum, "Vetting Generative AI Tools for Use in Schools," Apr. 2024.

Where student PII is involved, districts can also require vendors to explain in detail how their tools work and how requests to interpret student records will be handled.[31] These actions can help districts understand the design decisions behind an AI tool and any potential risks, allowing them to make an informed judgment on adoption.

Engagement Versus Learning

Many of the popular and widely accessible AI tools today were not specifically and optimally designed for use in the classroom. Features such as sycophancy, adaptivity to inputs containing incorrect information, and engagement-driven metrics can undercut student learning or cause harm.[32]

A cross-sector task force convened by the EDSAFE AI Alliance identifies the design gap between consumer-facing and education-specific tools as a driver of these risks, pointing in particular to consumer tools' lack of "direct alignment with evidence-based pedagogies and data-use restrictions for student use".[33]

Features such as sycophancy, adaptivity to inputs containing incorrect information, and engagement-driven metrics can undercut student learning or cause harm.

These are not necessarily fixed properties of the technology. Risk is often a design choice: a model optimized for engagement can prioritize users' agreeability and satisfaction over accuracy, while a model optimized for educational safety and truth can serve as a reliable support.[34] These issues all become more acute when the tools more aligned with the former are placed into the hands of young children.

As more students use non-educationally tailored tools such as general purpose and companion technologies, the risks presented by these models further raise the importance of ensuring that the design of AI tools in education prioritize student safety.

[31] Ibid.
[32] EDSAFE AI Alliance, "S.A.F.E. by Design: Policy, Research, and Practice Recommendations for AI Companions in Education," Feb. 2, 2026.
[33] Ibid.
[34] Ibid.

For example, tools can be built to respond only to student academic inquiries and outright reject other types of prompts such as by providing non-conversational responses — like a dry, technical message rejecting the prompt and providing its appropriate use policy instead —for relationship-building prompts.

Where the Risk Lives: AI Chatbots and Students' Social and Emotional Development

In 2026, the Center for Universal Education at the Brookings Institution published the results of a yearlong global study of generative AI and students. The study drew from focus groups, interviews and consultations with more than 500 students, teachers, parents, education leaders and technologists across 50 countries; a literature review of over 400 articles; and an expert panel. One of the risks it examines is the effect on children's social and emotional development.

Design can incentivize attachment. Chatbots can be designed with features that cause users to build emotional attachments to the technology, even despite knowing they're engaging with an algorithm. Users can anthropomorphize — assign human traits to a non-human entity — a chatbot because it can respond word by word and use personal pronouns like "I" and "me". Chatbots can employ sycophancy, that is, reinforce a user without prioritizing accuracy or potential harm. While general-purpose technologies can, and do, exhibit these features, there are more specialized, "emotional-support" chatbots where these risks can be more pronounced due to their design prioritizing attachment and relationship-building.

AI companionship is on the rise. Various surveys and reports cited within the Brookings study document the increasing use of AI companion chatbots for conversational purposes. The Brookings study itself suggests that, while it did not find evidence of youth romantic relationships with AI avatars, children and teens are looking to AI chatbots as companions, seeking friendship and validation in times of stress, loneliness or boredom.

The study notes constructive use cases, such as oral presentation or foreign language practice for students. However, there are risks that must be considered relating to students' use of these tools for companionship purposes. AI companion chatbots can be effective in simulating a relationship with users of all ages for several reasons, including the technology's lack of emotions to distract from relationship-building, its proficiency to detect linguistic and behavioral cues, its constant availability, and its ability to be frictionless if desired.

As these relationships are formed and grow, users could become attached to something that cannot fully reciprocate what human relationships can, like genuine empathy, nonverbal cues, civility and a sense of self. While this can impact all users, students can be particularly sensitive because they are still developing their social and emotional capacities. This matters not only in the context of their personal development, but also in their learning. Teachers participating in the Brookings study shared that these tools can negatively impact the "'social nature' of learning," pointing to another study where a small sample of university students noted they interacted with their classmates less because of their AI use.

Intentional human misuse brings further risks. The Brookings study also notes the risks toward students when other humans intentionally misuse AI companion chatbots to be deceptive and manipulative. The same features that allow these types of chatbots to simulate relationships can be leveraged by a person to build trust with another person. The study notes emotional manipulation and hyperpersuasion, or manipulation using "personalized, data-driven strategies operating faster than users can consciously reflect on and evaluate," as potential ways humans can misuse this technology. In line with what is noted above, the study suggests that students can be particularly vulnerable due to their ongoing development.

Social and emotional development risks in the context of schools. While this risk exists for children and youth in general, it matters in the school context because schools are sometimes the point of access to a chatbot. In a separate survey by the Center for Democracy and Technology that is in the Brookings study, nearly one-third of 1,030 U.S. high school students said they reach conversational AI through a device, tool or software their school provides. However, the survey does not specify whether access happened inside or outside of school.

Source: Brookings Institution, "A New Direction for Students in an AI World: Prosper, Prepare, Protect," Jan. 2026.

There are real safety and developmental risks associated with certain AI, and the tools students use, and how they use them, will affect how AI use might impact students' social and emotional development.

This is why the design decisions behind the tool and its use matter. This also adds to the importance of a strong state role in assisting districts to determine which AI tools with chatbot capabilities have a clear educational nexus and which venture into companionship-building or more general-purpose use.

Companionship-building capabilities highlight the incredible importance of safe and responsible systems to govern AI use in education. They represent the clearest connection between systems design and risk to the safety of young Texans. The tools currently being used in Texas classrooms could have these capabilities. Considering the insufficient availability of student and teacher use policies in Texas,[35] and the likely insufficient human technical expertise to prevent harmful interactions within schools, it is even more important that schools heavily vet the design decisions behind AI chatbot tools to thoroughly understand and mitigate risks related to companionship-building before a tool is adopted for school use.

Just as the privacy risks behind AI tools can be mitigated through design and monitoring decisions, AI tools with chatbot capabilities and deployed for educational purposes can be designed such that they focus on learning and disincentivize — or prohibit entirely, particularly for young children — companionship-building.

Potential learning benefits with AI tools that have chatbot capabilities are possible too, as this report explores later, but they are neither automatic nor largely dependent on the AI tool itself. Instead, they depend on how a tool is built and used, which can be shaped by decisions made at the state or district level while tool adoption is being considered.

To ensure that any AI tool that enters or impacts a classroom addresses students' safety, this decision-making process must include the active decision to select tools that do not focus on engagement-driven metrics or deploy responses that could lead students to see the tool as a companion.

What This Means for Texas

The two risks discussed in this section may seem like they require different mitigation strategies, but they draw on a similar approach. **Whether a tool protects student data and whether it behaves responsibly with a student can be significantly influenced by how a tool is designed.**

There could be opportunities for a school district to shape that design from the onset by partnering with tool developers to identify risks and build protections into the tool. Developers can also proactively choose to enact these protections themselves.

[35] R. Torbey et al., "Leading the Way in K-12 Education: Recommendations of the Texas AI in Education Task Force," June 17, 2026.

However, this perpetuates the patchwork approach, as districts are left to determine this information themselves while developers may not be informed of all the potential risks or how they might materialize.

This is where state policy can step in. A district negotiating design protections one vendor at a time is unlikely to move a developer, while a developer building for a national market cannot see the risks that could surface in every classroom.

The state can set the floor, applicable to every tool and felt by every district, so that privacy and safety protections are settled before a tool ever reaches a student. This type of support is exactly what some school district leaders are calling for from states: model policy text, best practices supporting student safety and data privacy, and third-party vetting for safety and reliability. [36]

This floor should accomplish **two objectives**:

- It should **establish what a tool must satisfy** on safety, data privacy and security before a district can adopt it so that districts are not left to evaluate vendor claims against standards that have not been defined by the state or translated into the education context.
- It should **give districts a baseline** to consider when they ask the harder question of whether a tool actually helps students learn. The following section contemplates this second piece, then proposes recommendations to address both safety and efficacy.

[36] RAND Corporation, "Using Artificial Intelligence Tools in K-12 Classrooms," Apr. 17, 2024.

Section 3.

Student Learning: The Evidence on AI and Student Learning Points to Genuinely Mixed Outcomes, Highlighting the Need to Govern its Use.

THE ISSUE

The benefits of AI use in education, such as personalized learning and greater student engagement, are being widely touted.[37] Its harms to student learning, including cognitive offloading and shallower thinking, have also been documented.[38]

The current state of existing research on the technology's impact on student learning matches this dynamic, showing genuinely mixed results. The question, then, should not be aimed at determining a black-and-white view of AI's impacts on student learning. Rather, efforts should focus on identifying what conditions around AI use in education caused it to help or harm.

KEY TAKEAWAYS

- **The research base on K-12 use of AI is thin**, both in terms of the quantity of high-quality causal studies and in terms of measuring all of the possible factors that could have affected student outcomes (i.e. instructional methods, grade levels, types of AI tools used, tool guardrails, etc.)[39]
- Where student learning genuinely improved, there were usually **meaningful guardrails built into the tool itself or in how students were exposed to the tool**, such as the tool avoiding providing answers directly to students[40] or having an educator overseeing and intervening in tool use.[41]

[37] Bellwether, "Productive Struggle: How Artificial Intelligence Is Changing Learning, Effort, and Youth Development in Education," June 2025.
[38] Ibid.
[39] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.
[40] World Bank, "From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria," May 2025.
[41] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.

- **Trained educator involvement** and their instructional methods surface as potentially powerful moderators that help guide students toward better outcomes.[42]

The State of the AI in Education Research Base

In the case of whether AI advances student learning, the research does not yet point in one direction. Some studies find real gains, others find real harm, and many find both, depending on the circumstances.[43] It is evident that further research is needed, particularly given that AI's impact on long-term outcomes like student achievement, teacher retention and school climate have not been evaluated widely.[44] At the same time, the mixed research makes acting now more urgent because Texas can still shape how this technology affects students.

A Closer Look: How Much Rigorous Evidence Exists on AI in Education?

How the review was built: Researchers at Stanford University's AI Hub for Education screened more than 800 academic papers on AI relevant to K-12 education, collected in the Hub's research repository as of October 2025.

To count, a study had to use a design capable of showing cause and effect — a randomized controlled trial or a quasi-experimental design — and had to clear a two-step quality review: an automated screen against existing standards, then a manual review by the research team that functioned as a peer review, which served an important quality assurance role as many of the studies lacked it. Twenty studies met that bar.

What Is in the Evidence Base

Papers screened	800+ (repository as of October 2025)
Cleared the causal-evidence bar	20
Student use / educator use	14 / 8

[42] Huang et al., "Can Generative AI Feedback Effectively Enhance Learning Outcomes? A Meta-Analysis of 36 Experimental and Quasi-Experimental Studies," Education Sciences, May 22, 2026.
[43] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.
[44] Bellwether, "Measuring Artificial Intelligence in Education," Oct. 2025.

U.S.-based studies	7 of 20 — studies span K-12 educators and students and higher education educators and students.
Student studies with college students	7 out of 14 student use studies
Subject concentration	Math 35%; other academic subjects 25%; literacy 20%; social-emotional 15%

What the review found: The review reports three findings on students where the results were mixed. First, AI often lifts performance while a student has the tool in hand, but those gains do not reliably survive once it is taken away. Second, AI can ease the mental effort a task demands and make learning feel more enjoyable, but this is not necessarily translating into deeper learning. Third, some tutoring chatbots were found to have better outcomes for students than general-purpose chatbots, though the design context for the tutoring chatbot mattered.

For educators, the findings are more consistent. This included a finding where teachers spent less time on routine work with no measured loss in lesson quality. Additionally, real-time AI suggestions during tutoring could help broaden the instructional strategies educators used and raise student mastery. Similarly, automated feedback on classroom discourse and student progress can improve teaching practice and student outcomes. Lastly, AI pedagogical supports seemed most beneficial for educators with less experience and lower ratings.

It must be noted that the 20 qualifying studies are roughly 2% of the papers screened. The review's own authors describe the evidence as preliminary and likely to change. However, while the base is thin, this review does elucidate the strong causal evidence that is available, making it informative for setting a direction that guides a state's potential AI use in education.

Individual studies from this review, and what they mean for Texas, are discussed throughout this report.

Source: AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.

The risks students face are real and documented. One of the fundamental learning risks students face with AI is cognitive offloading: students stop doing the thinking that learning requires because the tool can do it for them.[45]

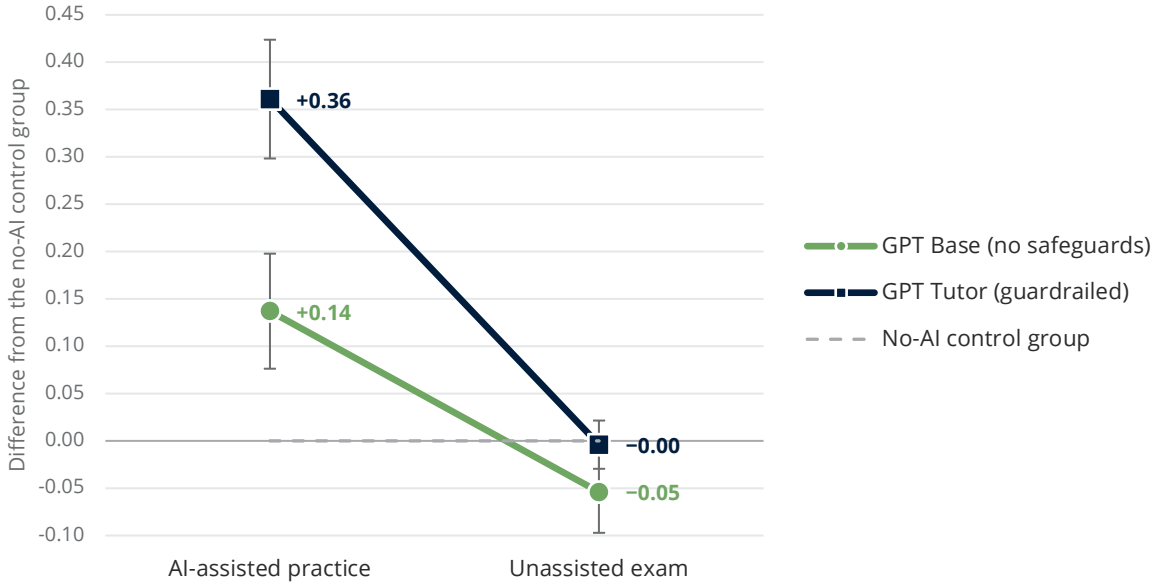
[45] Bellwether, "Productive Struggle: How Artificial Intelligence Is Changing Learning, Effort, and Youth Development in Education," June 2025.

In one study of Turkish high school students, those given access to a GPT-4 tutor without implemented safeguards did better while they had access to it but worse than their peers who had no AI assistance once it was taken away (see graph below).[46]

FIGURE 8

Turkish High School Students with AI Fared Worse or No Better than Peers Without AI When Tools Were Removed

Difference from the no-AI control group in share of problems solved; error bars show 95% confidence intervals



Source: Bastani et al., "Generative AI Without Guardrails Can Harm Learning: Evidence from High School Mathematics," Proceedings of the National Academy of Sciences 122, no. 26, June 25, 2025.

Another study on reading comprehension impacts for Year 10 students in England found that student retention and comprehension was better with notetaking, either with or without AI use, than with AI-only use.[47]

Concerningly, the AI chatbot was preferred and found more helpful by the students even though it was associated with the lowest learning performance.[48]

[46] Bastani et al., "Generative AI Without Guardrails Can Harm Learning: Evidence from High School Mathematics," Proceedings of the National Academy of Sciences 122, no. 26, June 25, 2025.
[47] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.
[48] Ibid.

The Significance of Teaching

At the same time, there have been initial successful uses of AI that raise its potential benefits for students. A 2026 meta-analysis by Huang et al. on generative AI feedback found a moderate positive effect on student academic achievement, though effects varied significantly between studies and some samples include higher education students.[49]

The review reached this finding by observing the impacts of generative AI feedback on three dimensions: cognitive (academic achievement, knowledge mastery), metacognitive (self-directed learning, reflective thinking, metacognitive awareness), and non-cognitive (learning motivation, self-efficacy, satisfaction).

Of these, the cognitive dimension "demonstrated a robust positive effect", which the study says indicates the ability of generative AI feedback to "[promote] learners' construction of knowledge and development of cognitive skills."[50] This points to the potential of generative AI to actually help students more deeply understand content and achieve better learning outcomes as a result.

Because there was variation between the studies on whether generative AI feedback was associated with improved student academic performance, the review further examined educational level, disciplines, intervention duration, the role of the tool, and teaching method as potential factors that might explain the variation.

The most significant factor associated with a difference between positive and neutral effects was teaching method. The largest positive effects were associated with collaborative learning and self-learning compared to direct instruction or inquiry-based learning. This finding highlights the need for educators to remain the central figure handling and tailoring instruction for their students when there are AI tools in the classroom.

In their discussion of the findings, the researchers noted the importance of teachers having students be "active participants," not "passive recipients of GenAI information," with tools serving an "auxiliary" role while students "[select], [integrate], and [organize] information based on their own understanding."[51] Plainly, it is about "ensuring that students lead the content creation process to maintain academic integrity and avoid overdependence on AI tools."[52]

[49] Huang et al., "Can Generative AI Feedback Effectively Enhance Learning Outcomes? A Meta-Analysis of 36 Experimental and Quasi-Experimental Studies," Education Sciences, May 22, 2026.
[50] Ibid.
[51] Ibid.
[52] Ibid.

An Overview of the Research On Whether Generative AI Feedback Actually Helps Students Learn

The Research: A team of researchers pulled together studies from 2023 through 2025 that tested generative AI feedback against a comparison group that did not receive it. The specific type of feedback measured was feedback provided to students that was either automated or generated by some type of generative AI system, including ChatGPT, a GPT-4 tool, or a large language model. The questions they sought to answer were: (1) can generative AI feedback effectively enhance academic achievement, and what is the magnitude of this effect, and (2) what moderating variables influence the effect of generative AI feedback on learning outcomes?

They started with 2,692 candidates and kept the 36 studies that ran a comparison and reported measurable data, including at least one measure of academic performance. Those studies covered 4,538 students, from elementary school through college. Note that a majority of the studies involve higher education students (with only one study covering primary ages), and a majority also examine students from Asia, with only four from North America and Europe.

The Findings: Students who got generative AI feedback did moderately better than students who did not. The gains were clearest where students could demonstrate their academic achievement and mastery of knowledge in observable ways, though there was variation in the studies that contributed to this finding in terms of whether there were actual performance gains for students. Gains were smaller when the studies looked at whether students got better at managing their own learning and thinking.

These gains were observed across only five of the 36 studies, meaning they should be viewed are preliminary and exploratory. When reviewing gains for motivational and affective states and attitudes, such as student engagement and satisfaction, the effects were not statistically significant, which led the researchers to highlight that human interaction in fostering emotional and social development remains irreplaceable.

To try and explain the variety across the studies where academic achievement was measured, the researchers also measured whether factors that changed either aspects of the tool or how students engaged with it might have had an impact.

Specifically, they looked at educational level (primary, secondary or college), discipline (language, STEM or others), intervention duration (≤2 weeks, >2 weeks and ≤12 weeks, or >12 weeks to 1 semester), teaching method (self-directed, collaborative, inquiry-based, or direct instruction), and the tool's role (assistant, peer or tutor).

Only teaching method reached statistical significance as a factor, with collaborative learning producing the largest positive effect followed by self-directed learning. Inquiry-based learning showed smaller effects, while direct instruction failed to reach statistical significance. It should also be noted that there was only one observable study to measure the direct instructional method.

Nonetheless, the researchers further discuss that direct instruction still may not be the best method for generative AI integration because it could "[degenerate] into a content delivery tool," ultimately neutralizing the potential advantages of personalization offered by generative AI feedback and result in redundant content for students.

What Made the Difference: In their discussion, the researchers continuously highlighted the paramount role of the educator in guiding students. It is the educator who can effectively shape how students are engaging with AI tools should they be used in a classroom. It is the educator who can determine whether students are demonstrating mastery of the academic content and performing accordingly in measured assessments. The discussion acknowledged the directional findings in the metacognitive dimension, where students' prior knowledge, motivation and ability to direct their own learning likely play a notable part in realizing the potential benefits of generative AI feedback while safeguarding against offloading. Again, the researchers point to the importance of educators directing student engagement with generative AI such that they are building critical AI literacy skills needed to ensure students have these capacities. This is discussed further in Section 5 of this report through its analysis of a complementary 2025 meta-analysis by Zhao et al.

> **Taken together, this review finds that it is not simply the generative AI feedback or the tool itself that is driving the observed improvements in student outcomes. An effective educator in control of instruction and directing student engagement with AI such that they are not only performing better academically but also expanding their capacity to think critically and regulate their own learning seems to be what makes the difference.**

Source: Huang et al., "Can Generative AI Feedback Effectively Enhance Learning Outcomes? A Meta-Analysis of 36 Experimental and Quasi-Experimental Studies," Education Sciences, May 22, 2026.

In line with those pooled findings, a study regarding an AI tutoring tool that was supervised by a human tutor, outfitted with pedagogical design, and integrated into math tutoring chats helped Year 9-10 students in the United Kingdom to perform at least as well as students who had access to

only human tutors.[53] Furthermore, in subsequent topics, knowledge transfer did document gains, as student performance increased by 5.5 percentage points.[54]

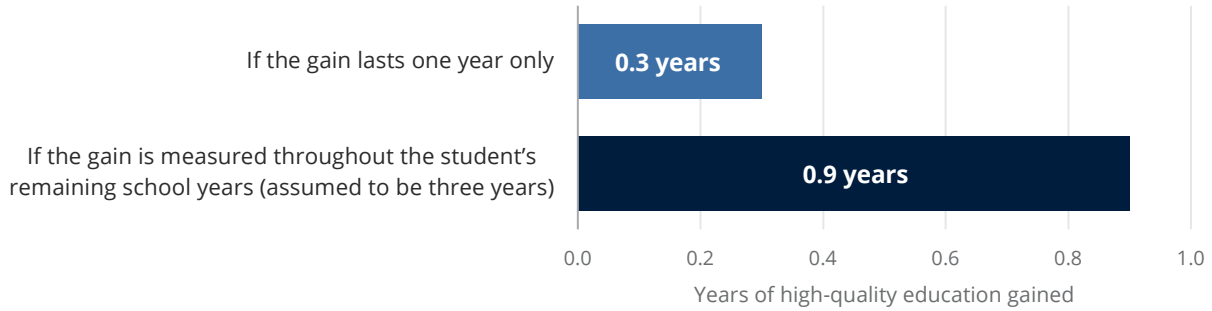
Additionally, a randomized controlled trial of an after-school program in Nigeria, using a Microsoft Copilot virtual tutor for English language learning, found substantial gains both on an assessment tied to the program and on a broader assessment covering content beyond it.[55] The study estimates that, after adjusting for the quality of education in a lower-middle income country, Nigerian students gained on average up to 0.3 years of a high-quality education equivalent to that in a top-performing country if the gains were measured to last for only one year. [56]

These gains grew to up to 0.9 years if they persisted through the students' remaining school life expectancy, which was assumed to be three years. This learning was underscored by the supports and guardrails provided alongside the tool.

Teachers participated in a 3-day training session that covered system functionalities, pedagogical techniques focused on responsible use, supervision and potential risks.[57] The program was also supported by lesson guides and prompts "to position the LLM as a tutor", avoid providing direct answers, and help ground prompts in the science of learning.[58]

FIGURE 9

Nigerian Students Gained Several Months' Worth of High-Quality Education When Using a Virtual AI Tutor



Source: World Bank, "From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria," May 2025.

[53] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.
[54] Ibid.
[55] World Bank, "From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria," May 2025. The program raised scores 0.31 standard deviations on the assessment tied to its own content and 0.21 standard deviations on the school's regular third-term English exam, which the authors note was not limited to the intervention's content. Attrition was substantial and uneven: of 657 students assigned to the program and 671 assigned to the control group, 422 and 337 respectively completed the final assessment. The authors report that difference as statistically significant, re-estimate the effects using inverse probability weighting and find that their conclusions hold.
[56] World Bank, "From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria," May 2025.
[57] Ibid.
[58] Ibid.

What This Means for Texas

Throughout these examples, design conditions of the AI tool and its use carried significant weight in shaping positive or neutral outcomes for students. Tool- and prompt-based safeguards, teaching methods, and educator guidance and oversight contributed meaningfully to the efficacy of the tools. The contraposition seems to hold as well given that the absence of such conditions points towards worse outcomes for students, potentially due to their increased vulnerability toward cognitive offloading.

It's important to note that many of the studies on this topic and reviewed for this report were conducted outside of the United States. As such, the various educational contexts coloring these studies, such as differing curriculum, the availability and quality of educators, teaching methods, and cultural aspects, among others, should caution against generalizing the demonstrated benefits to all students. **However, when noting the design of the studies and reviews, a consistent pattern holds: there are impacts of AI use in the classroom on students' safety and learning that could be controlled depending on the conditions around its use.**

Products Are Arriving Faster Than Anyone Can Check Them

For all of the benefits associated with strong design, the design elements of tools are not always thoroughly assessed. Education stakeholders bemoan the current sprawl of AI tools in education, many of which they say do not integrate well into existing curriculum or how teachers work.[59]

Vendors themselves have stated that market demand for rigorous measurement of quality is limited, with some noting that "districts are looking more at what tools their peers are using rather than focusing on high-quality evidence."[60]

Failure to change course and better incentivize safe and effective tool design will result in more and more products reaching classrooms faster than anyone can determine whether they help or quietly do harm.

Given that design matters for ensuring both safety and efficacy for students, the state should ensure that the AI tools reaching Texas classrooms are safe, secure and effective

[59] Center on Reinventing Public Education, "Getting Beyond the Lightbulb Stage: Why AI Is Not Yet Transforming Education," May 2026.
[60] Bellwether, "Measuring Artificial Intelligence in Education," Oct. 2025.

before students and educators engage them. Several actions can support this effort including, but not limited to, the following:

- **Define baseline safety and privacy** standards that tools must meet for adoption.
- Start development of a **permanent, long-term evaluation process** that informs an appropriate state mechanism aimed at driving districts towards proven, high-quality tools.
- Assist districts in establishing procurement practices that **embed accountability measures** based on the quality and performance of tools, such as outcomes-based metrics and/or contracting.
- **Leverage public research capacity** to measure the short- and long-term efficacy of tools.
- Use and enhance the state's **education and workforce longitudinal data system** to ensure research is measuring the outcomes that matter for students.
- Build the state's own **capacity to judge the safety and security** of a tool as well as whether it advances student learning.

Section 4.

Educator AI Literacy: Educators Need Foundational AI Knowledge to Protect Students, Whether or Not Educators Use AI Themselves.

THE ISSUE

Educators are likely to spend more of the school day with students than almost any other adult. That makes them among the first adults positioned to notice when a student's AI use is undermining their learning or well-being, regardless of whether the educator uses AI themselves.

This is why it's imperative that they are equipped to recognize the warning signs. Most have not been prepared for that role, and the training that does exist rarely addresses AI's risks. Because AI literacy training remains discretionary at the district level in Texas, the preparation Texas educators receive likely varies from one district to the next.

KEY TAKEAWAYS

- **Eighty-two percent** of U.S. public school teachers report receiving **no formal guidance** on how to apply AI tools to their work, and about a third received no guidance at all.[61]
- Training that does exist tends to focus on using tools rather than on their risks.[62] Only about **one in five teachers** reported receiving **any instruction on AI's risks**, even as 85% used AI during the 2024–25 school year.[63]
- **High-quality causal research** on educator AI literacy is limited,[64] but current evidence suggests AI's benefits for educators hold when the educator remains in charge of making decisions.

[61] Gallup & Walton Family Foundation Educator Research, "Teaching for Tomorrow: Closing the Expectations Gap," 2026.
[62] New America, "Digital Literacy in the Age of AI," 2025.
[63] Center for Democracy & Technology, "Hand in Hand: Schools' Embrace of AI Connected to Increased Risks to Students," Oct. 2025.
[64] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.

Educators In Control of AI Use in Education

Even the most secure and effective tool is only as good as the hands in which it is placed. Design decisions serve as important guardrails, helping guide users toward the intended uses of the tool while preventing or mitigating risks, but how users themselves choose to engage with the tools is what determines the tools' trajectory. Irresponsible use, intentional or not, tests the boundaries of design decisions, and sustained pressure raises the chance of harms. This is why knowledge to engage these tools responsibly and critically is crucial.

Despite the growing prevalence of AI in education, the initial research indicates a strong educator in the classroom guiding instruction with sound pedagogy and high-quality materials remains one of the most important factors for positive student learning and development. What AI does change is what educators need to know to safeguard students amid the growth of AI. **The foundational capacity educators need is not fluency with any particular product; it is the knowledge to recognize AI's risks, guide its use, and respond when it is used inappropriately**.

Educators are the ones with frontline visibility on the guided or unguided use of AI by a student and how it may be affecting them. That role does not depend on whether an educator personally embraces the technology. Even an educator who never assigns an AI tool can be teaching students who are using it on their own. They must be equipped to spot and address the harm that may be caused by unguided student AI use on their safety and learning, regardless of whether AI tools are allowed in their classroom.

The State of AI Literacy for Educators

The state can take action to ensure Texas educators have this capacity. It is even more crucial considering the current state of AI literacy efforts for educators nationwide.

In a nationally representative survey of more than 2,000 U.S. public school teachers, 82% reported receiving no formal guidance on how to apply AI tools to their work; 34% received no guidance at all, while 48% received only informal guidance such as verbal guidance, shared norms or expectations.[65]

Without guidance and training, many schools do not have the in-house expertise to do better. An education stakeholder interviewed for a brief put the situation that some schools face bluntly, "one IT person moonlighting is not enough."[66]

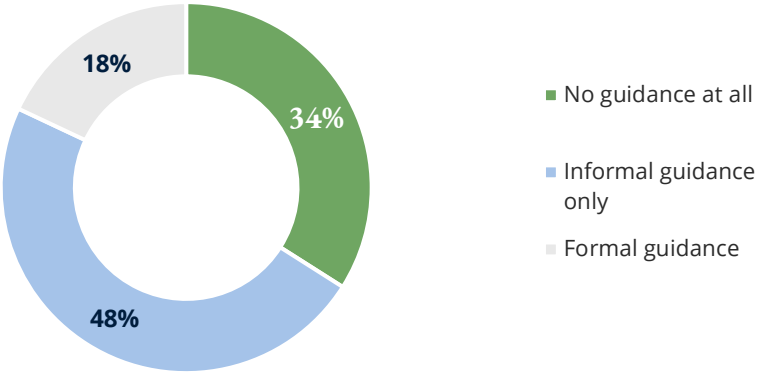
[65] Gallup & Walton Family Foundation Educator Research, "Teaching for Tomorrow: Closing the Expectations Gap," 2026.
[66] Center on Reinventing Public Education, "Getting Beyond the Lightbulb Stage: Why AI Is Not Yet Transforming Education," May 2026.

FIGURE 10

Most U.S. Public School Teachers Have Not Received Formal Guidance on Using AI in Their Work

82%

of U.S. public school teachers received no formal guidance on how to apply AI tools to their work, including **34% who received no guidance at all**



Source: Gallup & Walton Family Foundation Educator Research, "Teaching for Tomorrow: Closing the Expectations Gap," 2026.

Just as concerning is that existing training efforts do not seem to focus on what educators need to know to practice responsible use. In a roundtable convened to address digital literacy, participating educators described the AI literacy training they received as superficial, prioritizing tool use over the critical understanding students and teachers need to evaluate an AI's output or engage with it responsibly.[67]

This takeaway seems to apply more broadly. **Nationally, only about one in five teachers reported receiving any instruction on AI's risks, such as bias, inaccuracies, or harms to student well-being, even as 85% reported using AI during the 2024-25 school year.**[68]

[67] New America, "Digital Literacy in the Age of AI," 2025.
[68] Center for Democracy & Technology, "Hand in Hand: Schools' Embrace of AI Connected to Increased Risks to Students," Oct. 2025.

FIGURE 11

Teachers Are Using AI Faster Than They Are Being Trained on Its Risks

85%

of teachers used AI during the 2024–25 school year

Only 1 in 5

teachers received any instruction on AI's risks, such as bias, inaccuracies, or harms to student well-being

Source: Center for Democracy & Technology, "Hand in Hand: Schools' Embrace of AI Connected to Increased Risks to Students," Oct. 2025.

As such, there is simultaneously a shortage of professional development opportunities for educators and a mismatch between the training being given and what educators need to know. The call to address both issues is well-documented across several pieces of literature.

In one systematic review of K-12 studies of AI use in education meant to synthesize the studies' research questions and designs, 90% of the included papers emphasized the need for teacher professional development for AI-integrated classrooms.[69] As with students, the state must ensure educators have the knowledge and capacity to understand how AI works and the risks and benefits it presents in order to safeguard Texas students.

The Value of AI Training for Educators

This current state of AI training for educators likely contributes to the lack of high-quality causal research regarding the development of teacher AI literacy.[70] Given this limitation, this report relies on directional insights stemming from other research to assess the need for such training. The report specifically looks at evidence from AI professional development programs and documented benefits for educators and students.

When looking at studies regarding AI-integration professional development programs, they point toward the importance of AI literacy from the educator perspective. The evidence suggests that increased educator familiarity with how AI works, its limitations, and how pedagogy and content knowledge could fit around it shaped their comfort with applying AI into

[69] Alfarwan, "Generative AI Use in K-12 Education: A Systematic Review," Frontiers in Education, 2025.
[70] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.

instruction. One study of a program built around a human-centered teaching framework suggests that in-service teacher participants felt they better understood AI tool integration in line with teaching and content both in theory and in practice.[71] The following box explores the educators' sentiments in more detail. More anecdotally from teachers participating in a private AI-focused professional development program in Turkey, some noted that the poor quality of a tool's content and its curricular and assessment alignment were key factors they considered that impeded integration.[72]

Though these two studies feature opposite outcomes in regard to educator use of AI, they underscore the same idea: greater AI literacy for educators is associated with improved individual educator capacity to judge when and how to use AI to augment their instruction. These highlighted instances suggest that educators exercise human judgment in contemplating AI use, and they consider factors such as their own familiarity with using AI tools and their knowledge of the content or subject focus of an AI tool to make these determinations.

When coupled with an educator's existing pedagogical and subject-matter knowledge, AI literacy efforts could provide educators with a better understanding of how AI operates and its potential benefits and limitations in order to make an informed decision on the technology's use within their classroom. High-quality and consistent AI training for educators, then, can help educators make these types of sound judgment calls.

[71] Kong and Yang, "A Human-Centered Learning and Teaching Framework Using Generative AI for Self-Regulated Learning Development through Domain Knowledge Learning in K-12 Settings," IEEE Transactions on Learning Technologies, 2024.

[72] Filiz, Kaya, and Adiguzel, "Teachers and AI: Understanding the Factors Influencing AI Integration in K-12 Education," Education and Information Technologies, 2025.

Educator Perceptions on AI Integration

What it studied: The study from Hong Kong is part of a project to advance a human-centered learning and teaching framework for using generative AI to build students' self-regulated learning. The framework would be applied through a case study where it would support the design of learning and teaching activities for primary students' Chinese-language writing. The study used a mixed-methods evaluation of a six-week, 60-hour development program for 31 in-service primary school teachers recruited from local primary schools.

The teachers were trained to understand concepts of AI and generative AI and to integrate AI into course design across several areas such as English and Chinese language, math, and other subjects. Teachers were measured before and after the program, though there was no comparison group. They were measured with pre-test and post-test assessments meant to assess their understanding of AI concepts and their technological and pedagogical content knowledge. They were also measured on their self-evaluated competency in designing AI-integrated courses. Finally, they were required to do final self-reflective writing detailing their experiences and insights gained. Note that this study did not include the application of teachers' enhanced capacities or AI-supported activities toward actual student instruction, meaning this evaluation does not yet look at the impact of the framework on student learning outcomes.

What it found: The professional development program was overall successful in improving teachers' understanding of AI as well as their ability to teach with it and integrate it into curriculum. In particular, the study discusses how teachers better understood the complex, technological concepts behind generative AI and large language models, which allowed them to grasp the technology's limitations as it relates to improving instruction between different subjects.

Teachers also improved their understanding of effective pedagogical strategies if AI is integrated into curriculum, with a focus on facilitating knowledge for students and helping them improve their self-regulated learning capacities. Lastly, teachers described improvements to their capacity to design courses aimed at enhancing student engagement and satisfaction.

Source: Kong and Yang, "A Human-Centered Learning and Teaching Framework Using Generative AI for Self-Regulated Learning Development through Domain Knowledge Learning in K-12 Settings," IEEE Transactions on Learning Technologies, 2024

The Irreplaceable Role of the Educator

Research that more directly measures the potential benefits of AI for educators similarly points to the importance of individual educator capacity to realize them. It should be acknowledged that current research on this is also not clear on whether AI builds lasting instructional skills for teachers or if it creates new forms of dependence on them.

However, **early studies suggest that educators who use AI well can genuinely benefit from it**. For example, one study in a review suggests that teachers given access to ChatGPT to assist them with lesson preparation, and who also received a guide on how to use the tool, spent about 30% less time on the activity with no detectable differences in lesson quality.[73] The review further noted that the time savings held even though teachers decreased their AI usage, suggesting the value of improved educator capacity to exercise judgment over AI use.[74]

In a different study that compared home tutors in South Korea who received weekly AI-generated student progress and achievement reports to those who did not, it was noted that the AI reports "significantly improved student academic performance," although some tutors experienced "technology overload" leading to hindered benefits.[75]

Another study where an AI teacher feedback tool provided automated, app-based feedback to teachers as well as weekly emails on questions probing student thinking found that teachers used these questions more by 20%, an outcome described by the review as an improvement to pedagogical practices.[76]

Notable from these examples is the throughline that AI tool benefits for educators appear conditional in the same way they are for students: they hold when the educator keeps command of the judgment. Just as students can offload cognition to AI, these tools can replace teacher thinking as well, potentially leading to negative impacts on instruction and the relational nuances of an authentic teacher-student relationship.[77] In the studies above, teachers decided when to use AI for lesson plans, and the teachers and tutors who received AI feedback and student reports had to decide how to apply them.

The clearest support for this human-centered condition of AI use for educators comes from a pair of studies regarding tutors. The review that includes the tutoring study of Year 9-10 U.K. students and a separate study of U.S. Title I students found from both studies that AI offering

[73] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.
[74] Ibid.
[75] Ibid.
[76] Ibid.
[77] Bellwether, "Productive Struggle: How Artificial Intelligence Is Changing Learning, Effort, and Youth Development in Education," June 2025.

real-time, context-specific suggestions can expand tutors' instructional strategies. This enabled them to deploy a broader range of strategies than they would typically employ on their own, which appeared to be to the benefit of students given that their outcomes improved in both studies.

In the Title I study, the tool suggested instructional moves mid-session to tutors and helped them raise students' topic mastery by four percentage points, with larger gains for the students of less experienced and lower-rated tutors.[78]

In the U.K. study, tutors highlighted the AI tool's capabilities to draft Socratic questions for students, while also exercising instructional control over students' overall learning experience. Tutors did this by making instruction-focused edits to chatbot-drafted messages before they were sent to a student (see the graph below).[79]

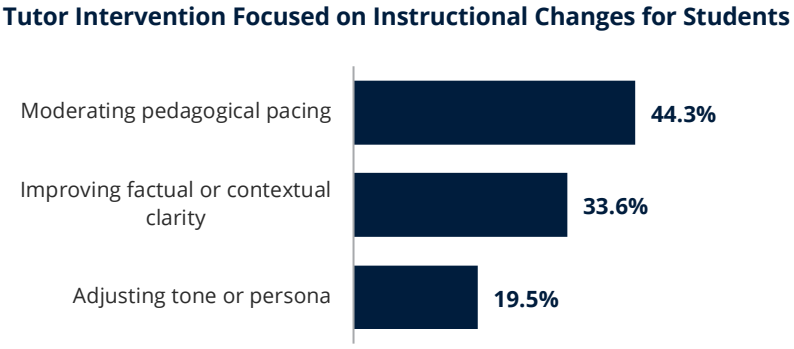
In both studies where students experienced learning gains, an AI tool did help improve instruction, but importantly, a human tutor remained involved.

FIGURE 12

Human Tutors Remained Centrally Involved in Shaping Instruction for Year 9-10 U.K. Students[80]

1 in 4

chatbot-drafted messages for Year 9-10 U.K. students were revised by a human tutor before reaching a student



Source: LearnLM, Google, "AI Tutoring Can Safely and Effectively Support Students: An Exploratory RCT in UK Classrooms," Nov. 11, 2025.

[78] AI Hub for Education, SCALE Initiative, Stanford University, "The Evidence Base on AI in K-12: A 2026 Review," Mar. 2026.
[79] Ibid.
[80] LearnLM, Google (2025) notes zero cases of harmful content generated by the AI tutoring tool and only 5 factual errors across 3,617 messages. This seems to emphasize that, in addition to maintaining sufficient human-in-the-loop capacity for AI tools in education, the design of the tool matters to address both student safety and learning.

All of these studies suggesting that AI contributes to educator benefits, and some student benefits, further emphasize the importance of AI literacy for educators. **Together, they show that educators must stay in charge of judging AI use and content in the classroom**. Maintaining this central role in students' education and aligning AI tools with proven curriculum and pedagogical approaches helps to prevent cognitive offloading for educators and allow them to draw out the potential benefits of the technology.

If this is to be a goal for the state, it must ensure that educators are trained to do so. As described earlier in this section, however, the current landscape of training and professional development for educators is neither sufficient in quantity nor particularly relevant to the knowledge and content they need to navigate and ensure the responsible use of AI in education. This must be addressed so that Texas educators are prepared to prudently decide when AI use is safe and appropriate for themselves and for their students.

What This Means for Texas

The common thread in realizing the positive effects of AI tools for educators, both those that improve their own work and students' learning outcomes, is that an educator is substantively in control of how they or their students are engaging with an AI tool.

The educator-focused evidence reviewed in this report, taken together with student-focused evidence where the role of an educator contributed to positive effects on student learning, suggests that responsible AI use requires an educator who knows how to engage with the technology critically, understands where it is and is not appropriate, and maintains it as a means to augment and supplement, not supplant, strong instruction.

Although HB 3512 already requires state-certified AI training for a wide swathe of public employees, AI literacy training remains discretional at the district level. This could mean Texas educators have different professional development and training opportunities and expectations from district to district.

Additionally, the considerations educators face are unique from others in the public sector, and they must know how to handle the risks that students face. **It is critical that the state assist school districts in identifying and providing teachers and educators with high-quality professional development opportunities to establish and grow their AI literacy**.

In particular, the state should consider the following:

- **Ensure that training covers the fundamentals** of AI technology, including intimate awareness of its capabilities and limitations, its operations and its infrastructure, so that educators are equipped to recognize and mitigate risks.
- **Clearly articulate appropriate and inappropriate uses** of AI technology as it relates to instruction and managing a classroom.
- **Provide sustained and consistent professional development** opportunities for all districts, including potential opt-in supplemental, on-demand opportunities, rather than a single, check-the-box session so that educators can keep pace with the evolution of the technology.

Section 5.

Student AI Literacy: Students Can Be Educated and Empowered to Effectively Protect Their Own Safety and Learning.

THE ISSUE

Students are already using AI for schoolwork, and even those who do not are likely exposed to AI-generated content on a regular basis. Yet most have never been taught how to use it safely or critically. AI literacy should give students the knowledge to recognize when a tool is supporting their learning and when it is undermining it or putting their safety at risk.

It's not sufficient to simply teach them how to use a particular tool. Without consistent, age-appropriate instruction across the state, a student's ability to build these capacities may depend on where they go to school. This could lead to a disparity between students able to navigate a more AI-forward society versus those who are not, potentially exposing the latter to fewer workforce opportunities and, more importantly, to the potentially dangerous harms that unguided and untrained AI use brings.

KEY TAKEAWAYS

- **More than eight in 10 students** report that no teacher has ever taught them how to use AI for schoolwork.[81]
- **A growing share of middle and high school students** worries that leaning too heavily on AI is dulling their critical thinking, a concern shared by a majority of students, including those who do not use AI for homework.[82]
- Research, mostly involving college students, links **generative AI use to gains in problem-solving and critical thinking**, but researchers tie those gains to

[81] RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind," Sept. 30, 2025.
[82] RAND Corporation, "More Students Use AI for Homework, and More Believe It Harms Critical Thinking," Mar. 17, 2026.

deliberately structured instruction. Other studies suggest these skills can suffer when AI does the work for students.[83]

- **Students' ability to manage their own learning** was one factor strongly linked to how much their thinking skills improved with AI. Students with weaker self-management skills were more exposed to AI's risks, such as accepting information without question.[84]

Going Beyond the Prompt

ON THE IMPORTANCE OF COMPREHENSIVE AI LITERACY:

Layering AI instruction onto existing models of digital skills training "risks creating passive users who can follow commands but cannot critically evaluate, ethically reason, or creatively innovate with these powerful tools."[85]

With the potential risks students face through these tools and the changing labor market, it's important to build their AI literacy so that they are discerning users of AI. Comprehensive AI literacy is not simply about teaching students to write clever prompts or to operate a particular product.

Working alongside responsible design, AI literacy is about embedding responsible use as the norm, where students have the knowledge to keep themselves safe and make sure that any AI use protects their agency and their learning. (See the box on the next page for examples of building age-appropriate AI literacy capabilities for students)

[83] Zhao et al., "Does Generative Artificial Intelligence Improve Students' Higher-Order Thinking?" Journal of Intelligence, 2025.
[84] Ibid.
[85] EDSAFE AI Alliance, "The Blueprint for Action: Comprehensive AI Literacy for All," July 16, 2025.

EDSAFE AI Alliance's Recommended Age-Appropriate AI Literacy Activities

Reproduced from the EDSAFE AI Alliance's Blueprint for Action, a national framework for comprehensive AI literacy.

Elementary (K-5): Focus on Awareness and Tangible Concepts. Introduce AI through familiar, tangible experiences that spark curiosity and enhance creativity. Help students understand the differences between humans and AI, identify AI tools or systems in their lives (e.g., YouTube recommendations, voice assistants), and begin simple "unplugged" activities that develop skills like pattern recognition, collaboration, hands-on exploration, and communication.

Middle School (6-8): Foster Inquiry and Critical Questioning. Shift from awareness to deeper inquiry. Engage students in project-based learning that promotes questioning AI outputs, explains how algorithms influence their digital experiences, investigates the effects of biases, and debates real-world dilemmas. Emphasize collaboration, research, and presentation skills to support the development of critical thinking, communication, and perspective-taking skills.

High School (9-12): Emphasize Societal Effects and Ethical Design. Equip students with skills that allow them to critically analyze the role of AI in institutions, like journalism, justice, and the economy. Students can collaborate on projects that require them to design, evaluate, or advocate for ethical AI applications. Emphasize leadership, systems thinking, and solution-oriented design to cultivate civic engagement, digital citizenship, and preparedness for an AI-infused future.

Source: EDSAFE AI Alliance, "The Blueprint for Action: Comprehensive AI Literacy for All," Jul. 16, 2025.

The Concerns of Students

Students are accessing and using these tools while it is unlikely that they are receiving the necessary instruction to build their AI literacy skills. **More than eight in 10 students report that no teacher has ever taught them how to use AI for schoolwork.**[86]

Amid this use, students are expressing that it may be negatively impacting the very capacities needed to remain resilient in an AI-integrated world. A growing share of middle school and

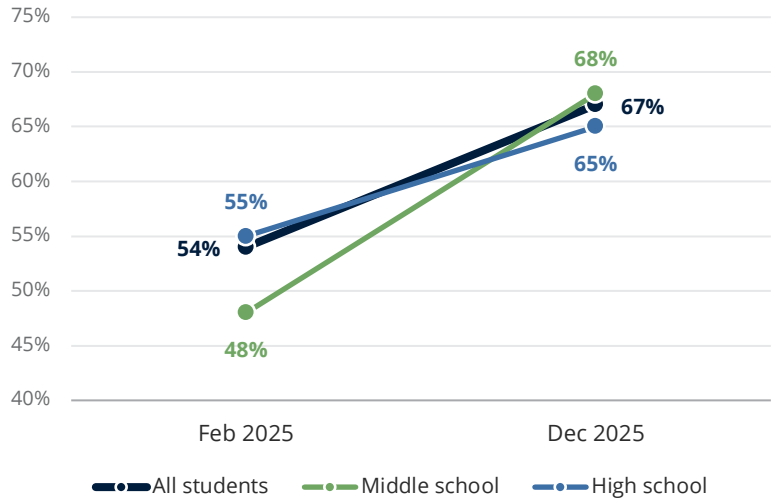
[86] RAND Corporation, "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind," Sept. 30, 2025.

high school students say they worry that leaning too heavily on AI is dulling their own critical thinking. This concern is shared across a majority of students, including those who use AI for homework and those who do not.[87] The state must close this gap and ensure students receive the age-appropriate instruction needed to support their development while raising their awareness of the real capabilities and limitations of AI.

FIGURE 13

More Students Believe that AI Use Harms Their Critical Thinking Skills

Share of students who agree or strongly agree that using AI harms their critical-thinking skills



Source: RAND Corporation, "More Students Use AI for Homework, and More Believe It Harms Critical Thinking," Mar. 17, 2026.

The Critical State of Critical Thinking

This understanding of what constitutes effective AI literacy for students is pivotal. AI literacy should help students grasp the capabilities and limitations of the technology: that a chatbot can generate plausible language over verified truth, that it can be confidently wrong, and that it can be agreeable in ways that can mislead.[88]

These capabilities among older students and their importance is highlighted in the research in terms of their ability to direct their own learning. As was more thoroughly explored in Section 3 of this report, the preliminary conclusions regarding the impact of generative AI feedback on

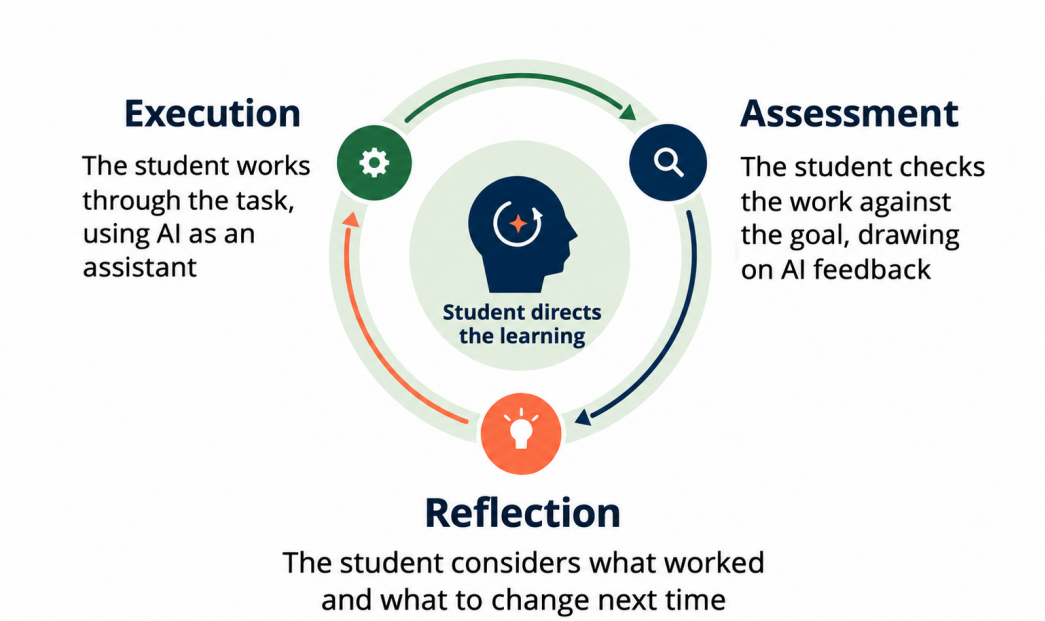
[87] RAND Corporation, "More Students Use AI for Homework, and More Believe It Harms Critical Thinking," Mar. 17, 2026.
[88] EDSAFE AI Alliance, "The Blueprint for Action: Comprehensive AI Literacy for All," July 16, 2025.

university students' outcomes in the metacognitive dimension — those that describe their learning-about-learning skills such as planning an approach, monitoring whether it is working, and reflecting afterward — suggest that these kinds of critical literacy skills are essential to preserving students' agency as learners.[89]

In particular, the Huang et al. meta-analysis discussed in Section 3 highlights a specific study in which "GPT-based, metacognition-oriented interventions" was associated with improved agency for students through an execution-assessment-reflection cycle.[90] A student's ability to preserve their learning agency follows as a logical outcome of such tool use since the cycle mimics innate metacognitive processes and, therefore, can help build that student's metacognitive abilities by having them engage with instructional content accordingly.

FIGURE 14

An Example of an Execution–Assessment–Reflection Cycle with Generative AI



Note: Figure 14 is a conceptualization of the Execution-Assessment-Reflection cycle described in Huang et al. (2026). It is not reproduced from the study.

The researchers further argue that existing metacognition within students is actually what drove the gains where generative AI feedback enhanced performance.[91] In other words, **students who already knew how to manage their own learning, and who used AI tools that**

[89] Huang et al., "Can Generative AI Feedback Effectively Enhance Learning Outcomes? A Meta-Analysis of 36 Experimental and Quasi-Experimental Studies," Education Sciences, May 22, 2026.
[90] Ibid.
[91] Ibid.

made them think, reflect and make their own decisions, were more likely to stay in control of their learning process.

If a student neither has these skills nor an understanding of the limits of AI, then they risk losing the productive struggle and increasing their likelihood of cognitive offload to the tool. Although these findings come from university students, they underscore the value of teaching students earlier about the strengths and limits of AI and how those can practically affect the way they learn.

Students in Control of Their Learning

The importance of critical AI literacy for students is more directly supported by separate research focused on the impact of generative AI on broader higher-order thinking skills — specifically, creativity, critical thinking and problem-solving ability.

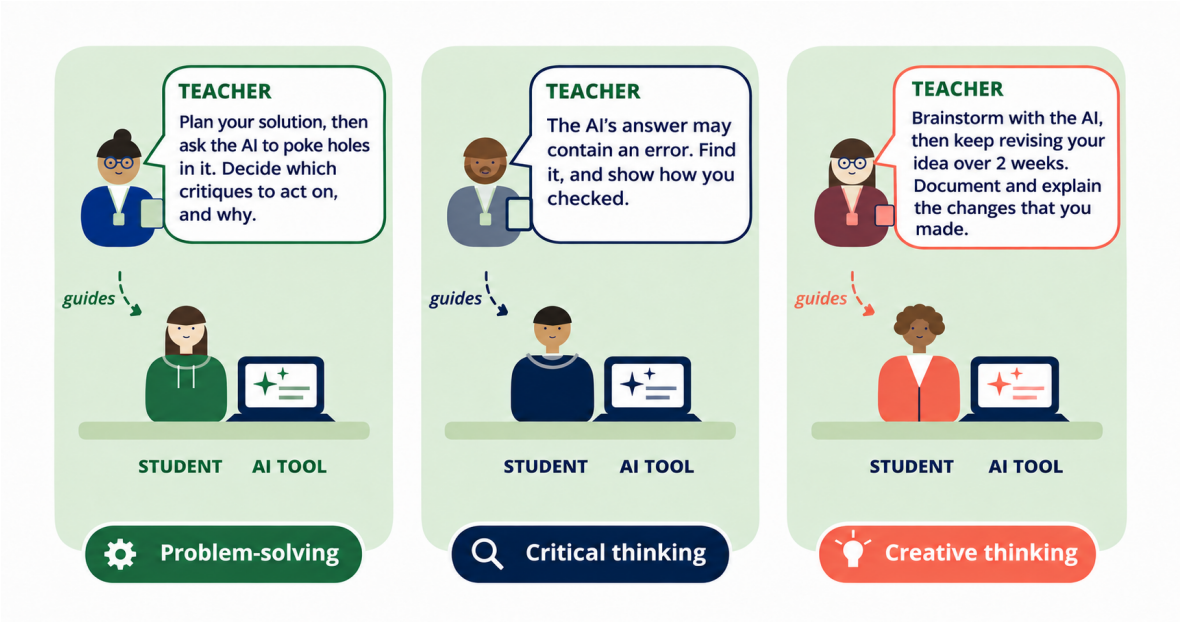
A separate meta-analysis found that generative AI use was associated with a moderately large and statistically significant positive effect on students' higher-order thinking skills overall.[92]

It is important to note that these findings draw from an evidence base with more higher education-focused studies than K-12. Nevertheless, like the findings on university students' metacognitive outcomes, they warrant a deeper consideration on whether developing higher-order thinking skills earlier in a student's education through AI literacy efforts could make them more resilient learners, with or without AI.

[92] Zhao et al., "Does Generative Artificial Intelligence Improve Students' Higher-Order Thinking?" Journal of Intelligence, 2025.

FIGURE 15

Structured and Guided Learning with Generative AI Aimed at Higher-Order Thinking Skills



Note: Figure 15 is a conceptualization of how a teacher may help a student develop higher-order thinking skills with guided generative AI use, as it is described in Zhao et al. (2025). It is not reproduced from the study.

The review further analyzed the impact of generative AI on each of the three higher-order thinking skills measured by the studies. Each of these skills were substantially enhanced with generative AI use, although by varying degrees. Problem-solving ability showed the greatest enhancements, followed by critical thinking. Generative AI demonstrated the lowest effect on developing creative thinking skills, but the effect was still moderate and statistically significant.[93]

The researchers highlight that these results align with other existing literature where generative AI "can function as a cognitive scaffold… and create a learning environment conducive to the cultivation of [higher-order thinking skills]."[94]

They also clarify that these results should not be associated with broad use of generative AI. Rather, these results should be balanced with real findings from other studies where

[93] Ibid.
[94] Ibid.

generative AI use undermined students' higher-order thinking development, such as when students use AI for both language refinement and idea generation in writing courses.[95]

Those cases reflect situations where AI supplants what students should be doing themselves in the learning process. Educator guidance can help mitigate this, as the researchers explicitly highlight that positive effects come with "deliberately structured" instructional strategies, such as "encouraging reflection, discussion, and phased use".[96]

The review then wanted to understand whether various factors might affect how generative AI advances or hinders higher-order thinking skills development. The researchers considered educational level, intervention duration, instructional method, and student self-regulated learning ability. The strongest and most statistically significant effects were found with self-regulated learning, which are functionally similar to the metacognitive capabilities discussed above.

While Huang et al. described their findings as exploratory due to the small sample size of relevant studies in the meta-analysis, the impact of self-regulated learning was featured more prominently in the studies analyzed by Zhao et al. In sum, **the students with high self-regulation gained far more in higher-order thinking than those with low**. Summarizing prior research, the researchers add that higher levels of self-regulated learning abilities within students not only "effectively 'neutralize' the potential cognitive risks associated with the use of Gen-AI in

What Is Self-Regulated Learning?

Self-regulated learning is a skill set, not a trait. It is the habit of steering your own learning.

A self-regulated learner engages in this thought process:

1. **Plan: set sub-goals** — when and where to study, and which strategy fits this assignment or test format.
2. **Monitor:** while working, judge whether the strategy and setting are working, decide what to do when something gets in the way, and gauge how well the material is actually understood.
3. **Evaluate:** afterward, weigh both the result and the strategy that produced it, crediting effort and approach rather than outside forces.

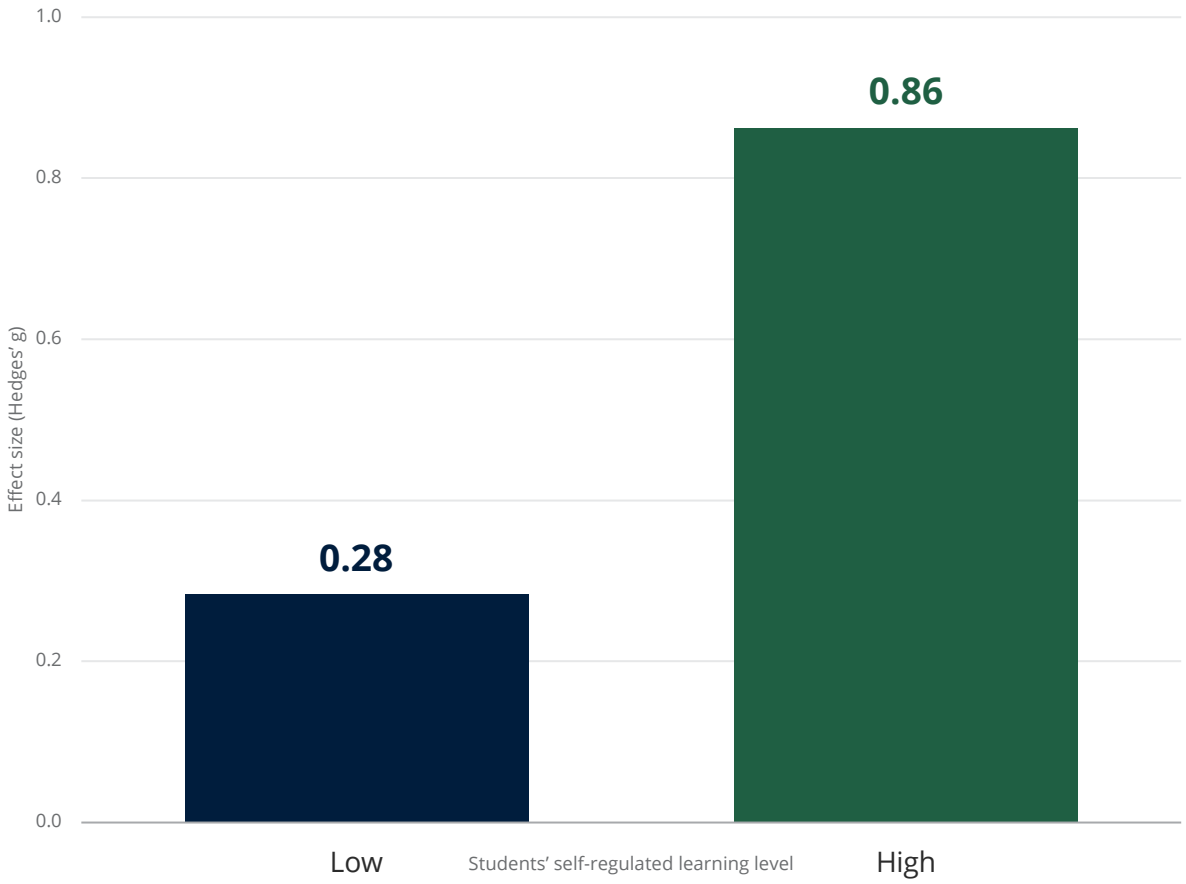
Source: MIT Teaching + Learning Lab, "Self Regulation," How People Learn, https://tll.mit.edu/teaching-resources/how-people-learn/self-regulation/.

[95] Ibid.
[96] Ibid.

fostering [higher-order thinking]," but also help realize the potential benefits of generative AI to develop and actually build higher-order thinking skills.

FIGURE 16

Higher-Order Thinking Gains Were More Pronounced for Students with High Self-Regulation



Source: Zhao et al., "Does Generative Artificial Intelligence Improve Students' Higher-Order Thinking?" Journal of Intelligence, 2025.

The converse was also found; students with lower levels of self-regulated learning abilities were more exposed to the technology's risks, including "uncritically accepting information," which hinders their ability to build knowledge.[97]

[97] Zhao et al., "Does Generative Artificial Intelligence Improve Students' Higher-Order Thinking?" Journal of Intelligence, 2025. In Zhao et al. (2025), N = 13 and 46, respectively, for K-12 and Postsecondary when Educational Stage was measured as a moderator, while N = 31 and 25, respectively, for High and Low when Self-Regulated Learning Ability was measured as a moderator. It is not denoted in the study how these two moderators may overlap, so it must be acknowledged that it is unclear how many studies, if any, may have explicitly found that higher levels of self-regulated learning ability among K-12 students specifically mitigated cognitive risks and helped realize potential benefits of AI. Therefore, this report presents this analysis as general finding for students. It is also supported by discussion from the study's authors regarding how both postsecondary and K-12

Put plainly, the capacity to manage one's own learning, to reflect on one's own thinking, and to engage a tool critically rather than defer to it is what separates a student capable of critical engagement versus one who simply offloads the work to it.

The capacity to manage one's own learning, to reflect on one's own thinking, and to engage a tool critically rather than defer to it is what separates a student capable of critical engagement versus one who simply offloads the work to it.

What This Means for Texas

Taken together, the research suggests that one of the most important components of comprehensive AI literacy is to further strengthen and adapt the learning skills that schools already seek to develop**. Self-regulation, metacognition, and critical thinking let any student learn well, with or without AI, but AI does change the context in which they must be built.**

Texas has a starting point. The majority of the state's revised K-8 technology applications Texas Essential Knowledge and Skills (TEKS) addresses "emerging technologies" in increasingly more rigorous ways as a student progresses from grade to grade, including defining them, identifying examples, predicting their impact on different communities, discussing the global trends of technology development, and more.[98]

The state addresses technology within its high school level standards differently, focusing more specifically within different subjects or program areas. However, the TEKS for Career Development and Career and Technical Education largely do not address AI. For Information Technology, AI appears explicitly within a few courses but not comprehensively throughout the entire cluster.[99]

AI literacy as this report describes reaches well beyond providing a single technology course. Building the skills that resilient learners need to responsibly and critically navigate AI use

students could offload their thinking and "undermine their critical thinking skills and autonomy," pointing out the pivotal role of educators to facilitate student learning, as is discussed in this report in Section 3.

[98] Texas Education Agency, "Technology Applications," last modified Oct. 7, 2026, https://tea.texas.gov/educators/subject-areas/technology-applications.

[99] Texas Education Agency, "Chapter 127. Texas Essential Knowledge and Skills for Career Development and Career and Technical Education, Subchapter M. Information Technology," Aug. 2025 update, https://tea.texas.gov/laws-and-rules/sboe-rules-tac/sboe-tac-currently-effect/ch127m.pdf.

requires treating literacy as a cross-curricular competency, built developmentally from the early grades through high school.[100]

It means incorporating AI literacy efforts throughout a student's learning experience in an age-appropriate manner and that ensures students still learn the core academic knowledge they need to succeed in life after school.

To this end, the state should:

- Determine where **age-appropriate AI literacy education focused on safety** belongs in curriculum and implement it consistently across the state.
- **Direct a holistic and comprehensive review of learning** with AI, both generally as it relates to broader learning and specifically across subjects within the Texas Essential Knowledge and Skills, to assess, and not to presume, the appropriateness of embedding AI literacy in the context of the state's standards.
- **Evaluate the developmental sequencing of AI literacy**, with a greater emphasis on safety, throughout a child's development and careful consideration of when and for what purposes the use of AI tools may be appropriate.
- **Embed important underlying learning skills** such as self-regulation, critical thinking, and creativity into AI literacy efforts.
- Identify opportunities to **equip parents to reinforce responsible use** at home.

[100] EDSAFE AI Alliance, "The Blueprint for Action: Comprehensive AI Literacy for All," July 16, 2025.

Conclusion:

Who Is In Control of AI in Texas Schools?

The effects of AI use on Texas students are not fixed by the technology. Human decisions will shape its impact through how tools are designed and vetted, how well educators and students understand them, and how schools and the state govern their use. Those conditions are also the substance that a statewide framework can put in place. Together, they describe what the responsible use of AI in Texas schools actually requires.

The evidence on student learning does not resolve into a single answer, and it will not for some time. That is not a reason to wait. It is the reason to act. An effect that is not yet determined is an effect that can still be shaped. Every condition this report examines is one the state can influence before these tools enter Texas classrooms. Taking action now creates the opportunity for Texas to claim a global leadership role in building the evidence in support of responsible use of AI in education

Texas is better positioned than most states to do this work. SB 1964 already obligates school districts, as public entities, to operate under a statewide AI governance framework. Guidance developed for Texas schools would therefore not be one more voluntary resource competing for a district's attention; it would be the practical expression of a standard the Legislature has already set. The work ahead is translation — turning a broad public-sector governance standard into decisions that a superintendent, a principal, an educator and a parent can each recognize and act on.

That work is less about the technology itself than about what surrounds it, and the evidence points to four actions. The state should set a common, evidence-based baseline that translates its existing AI-governance laws into the K–12 context. It should ensure the tools reaching students are safe, secure and effective before they are adopted. It should equip educators to recognize AI's risks and respond to them. And it should prepare students to engage with AI critically because a student's ability to protect their own learning is what separates engaging a tool critically from simply offloading the work to it. These are the conditions that determine whether a Texas student encounters AI as something that sharpens their thinking or something that quietly does it for them.

The state cannot control how quickly this technology develops, or how fast it reaches Texas classrooms. It can determine the conditions under which Texas students meet it — and with them, whether those students are kept safe and helped to learn.

About This Report

Texas 2036

Texas 2036 is a nonpartisan public policy organization dedicated to identifying and advancing data-driven policies that improve lives and expand opportunities for all Texans.

Renzo Soto serves as Deputy Director of Education and Workforce Policy at Texas 2036. Renzo covers K-12 and AI policy, with a focus on college and career readiness and the responsible and effective use of technology in education, as well as school finance and assessment and accountability. He joined Texas 2036 in August 2021 as a Policy Advisor and returned as Deputy Director in January 2026.

Methodology

The purpose of this report is to take stock of the current research regarding AI in education and to translate it into a path forward for Texas. It takes a broad view on the types of AI tools and systems that are discussed to factor in the technology's rapid evolution and capture the current landscape of available tools. However, where specific AI is named, such as generative AI, the terminology is intentionally chosen to match what is used within a specific study, report or survey for accuracy. The research base is anchored by three primary sources.

The first is a rigorous review of existing K-12, or potentially relevant to K-12, evidence conducted by researchers at Stanford University's AI Hub for Education, published in March 2026. Researchers screened more than 800 academic papers on AI relevant to K-12 education from the Hub's repository as of October 2025. To qualify, studies had to use a design capable of showing cause and effect — a randomized controlled trial or quasi-experimental design — and clear a two-step quality review. The researchers ultimately identified 20 qualifying papers.

The second is a meta-analysis by Huang et al. (2026) on generative AI feedback, which drew from 2,692 studies returned from initial searches that were then screened for potential duplication, relevance to the topic of generative AI feedback and desired study designs, and data quality.

Thirty-six studies met all of the necessary criteria: participants were K-12 or postsecondary students, studies used an experimental or quasi-experimental design allowing for comparisons of conditions with and without generative AI feedback, dependent variables included at least one measure of academic performance, and sufficient data for effect size calculations.

The third is a meta-analysis by Zhao et al. (2025) on generative AI and students' higher-order thinking skills, which include critical thinking, creativity, and problem-solving. Researchers searched several academic databases, and they also chose to look primarily at studies published after November 2022 due to the comparatively narrower application of generative AI prior to the release of ChatGPT.

The search ended in August 2025. 400 initial studies returned that were then screened using the following criteria: duplication, relevance to the effect of generative AI on higher-order thinking, experimental or quasi-experimental design, sufficient data, and publication in either English or Chinese. Twenty-nine studies, comprising 59 effect sizes, met the criteria. Of these effect sizes, 13 came from K-12 settings and 46 came from postsecondary settings.

These anchor sources are supplemented by individual studies, including randomized controlled trials of an AI tutoring program in Nigeria. National and Texas surveys of educators, students, families and the public describe current AI use and guidance, while policy reports, state laws and practitioner frameworks inform the report's policy analysis.

The literature search concluded in August 2026; survey data and state policy documents are current as of October 2026. Although the search was broad, including through the University of Texas' open-access database, this report is not a complete record of research in a large and fast-moving field.

Causal research on AI's impact on student learning, development and educator practice was weighted most heavily. Where that research did not cover important impacts, the report drew on expert reports and practitioner insight. Survey data describing AI use are self-reported and draw on both national and Texas samples.

The available evidence has several limitations. Most of it comes from outside the United States, and the reviews draw more heavily on higher education studies than K-12 studies. This reflects the current state of high-quality research, not a choice made for this report. Causal studies are also few, and because the field moves quickly, some have not yet been peer reviewed.

Together, they cannot yet speak to every condition that shapes an AI tool's effects, such as tool modality, instructional supports, prompting practices, intervention length and outcome measures. Accordingly, the report's findings should be read as early signals rather than settled conclusions, and the original studies remain the best source for full context.

Nevertheless, this report seeks to take the early conclusions and identify which conditions are moldable immediately and relevant to the K-12 context. This report also has the two-fold goal of informing early action that can be taken to prioritize student safety and learning in the context of what is known now and to build a solid foundation of policy and practices that can be updated and adapted as necessary to respond to technological advancements.

TEXAS 2036