




OFFICE OF
Educational Technology

Empowering Education Leaders:

*A Toolkit for Safe, Ethical, and Equitable
AI Integration*

October 2024



Empowering Education Leaders: A Toolkit for Safe, Ethical, and Equitable AI Integration

Miguel A. Cardona, Ed.D.

Secretary, U.S. Department of Education

Roberto J. Rodriguez

Assistant Secretary, Office of Planning, Evaluation and Policy Development

October 2024

Examples Are Not Endorsements

This document provides non-regulatory, education-specific guidance that is not intended to be comprehensive nor exhaustive. The examples and resource materials provided are intended for the user's convenience. The inclusion of any specific material is not intended to endorse any views expressed or products or services offered. These materials may contain the views and recommendations of various subject matter experts as well as contact addresses, websites, and hypertext links to information created and maintained by other public and private organizations. The opinions expressed in any of these materials do not necessarily reflect the positions or policies of the U.S. Department of Education. The U.S. Department of Education does not control or guarantee the accuracy, relevance, timeliness, or completeness of any information from other sources that are included in these materials. Other than statutory and regulatory requirements included in the document, the contents of this guide do not have the force and effect of law and are not meant to bind the public.

Contracts and Procurement

This document is not intended to provide legal advice or approval of any potential Federal contractor's business decision or strategy in relation to any current or future Federal procurement and/or contract. Further, this document is not an invitation for bid, request for proposal, or other solicitation.

Licensing and Availability

This guide is in the public domain and available on a U.S. Department of Education website at <https://tech.ed.gov>.

Requests for alternate format documents such as Braille or large print should be submitted to the Alternate Format Center by calling 1-202-260-0852 or by contacting the 504 coordinator via [email](mailto:om_eaos@ed.gov) at om_eaos@ed.gov.

Notice to Persons with Limited English Proficiency

If you have difficulty understanding English, you may request language assistance services for U.S. Department of Education information that is available to the public. These language assistance services are available free of charge. If you need more information about interpretation or translation services, please call 1-800-USA-LEARN (1-800-872-5327) (TTY: 1-800-437-0833); email us at Ed.Language.Assistance@ed.gov; or write to U.S. Department of Education, Information Resource Center, LBJ Education Building, 400 Maryland Ave. SW, Washington, DC 20202.

How to Cite

While permission to reprint this publication is not necessary, the suggested citation is as follows:

U.S. Department of Education, Office of Educational Technology, *Empowering Education Leaders: A Toolkit for Safe, Ethical, and Equitable AI Integration*, Washington, D.C., 2024.

Acknowledgements

Project Team

Empowering Educational Leaders: A Toolkit for Safe, Ethical, and Equitable AI Integration was developed under the leadership and guidance of **Roberto J. Rodríguez**, Assistant Secretary for the Office of Planning, Evaluation and Policy Development, **Anil Hurkadli**, Interim Deputy Director for the Office of Educational Technology, and **Bernadette Adams**, Senior Policy Advisor for the Office of Educational Technology.

This work was developed with support from Digital Promise under contract (91990019C0076) and led by **Pati Ruiz** and **Jeremy Roschelle** with: **Sana Karim**, **Babe Liberman**, **Gabrielle Lue**, **Eric Nentrup**, **Anthony Baker**, **Teresa Solorzano**, **Sarah Martin**, and **Carolina Belloc**. Graphics were developed by **Manuel Herrera** with additional figures developed by **Laura Cox**.

Contributing Members of the Educational Leader Community

- Erik Burmeister, Fremont Unified School District
- Dr. Kip Glazer, Mountain View-Los Altos Union School District
- Dr. Nneka McGee, San Benito Consolidated Independent School District
- Gina Fugate, Maryland School for the Blind
- Krystal Chatman, Jackson Public School District
- Melissa Moore, El Segundo Unified School District
- Dr. Mary Catherine Reljac, Fox Chapel Area School District
- Michael Hower, Fox Chapel Area School District
- Sallie Holloway, Gwinnett County Public Schools
- Andrew Fenstermaker, Iowa City Schools
- Dr. Patrick Gittisriboongul, Lynwood Unified
- Michael Nagler, Mineola Public Schools
- Mario Andrade, Nashua School District
- Dr. Mark Buckner, Oak Ridge School District
- Dr. Nathan L. Fisher, Roselle Public Schools
- John Malloy, San Ramon Valley Unified

Digital Promise Support Staff

- Amanda Armstrong
- Merijke Coenraad
- Diane Doersch
- Judi Fusco
- Marci Giang
- Keun-Woo Lee
- Kelly Mills
- Sierra Noakes
- Yenda Prado
- Heather Singmaster
- D'Andre Weaver
- Josh Weisgrau

Advisors

- James Basham
- Sarah Burriss
- Kristina Ishmael
- Nicole Hutchins
- Rohit Kataria
- Eleazar (Trey) Vasquez III

We also thank the many educators and nonprofit organizations that attended our listening sessions and contributed their ideas for translating the Department's recommendations for Artificial Intelligence in education into practical guidelines.

Table of Contents

Introduction	1
Mitigating Risk: Safeguarding Student Privacy, Security, and Non-Discrimination.....	5
Module 1: Opportunities and Risks.....	5
Opportunities	5
Risks.....	6
Discussion Points and Questions.....	11
Module 2: Privacy and Data Security.....	12
Monitoring Privacy and Security with AI.....	14
Discussion Points and Questions.....	15
Module 3: Civil Rights, Accessibility, and Digital Equity.....	17
Connecting AI with Educational Equity.....	20
Accessibility and Assistive Technology.....	20
Bias and Algorithmic Discrimination.....	20
Discussion Points and Questions.....	22
Building a Strategy for AI Integration in the Instructional Core.....	23
Module 4: Understanding the Evidence.....	23
Connecting Evidence to AI-Enabled Interventions.....	25
Discussion Points and Questions.....	26
Module 5: Considering the Instructional Core	27
Discussion Points and Questions.....	29
Module 6: Planning Your AI Strategy	30
Inform and Listen.....	30
Preparing for the Listening Session	31
During the Listening Session	32
After the Listening Session.....	32
Planning an Approach to Prioritization	33
Discussion Questions	33
Specifying Factors.....	34
Module 7: Guide and Support	35
Planning a Task Force.....	35
Ethics as a Cornerstone.....	36
Creating a Charter.....	37
Maximizing Opportunity: Guiding the Effective Use and Evaluation of AI	40
Module 8: Building AI Literacy for Educators.....	40
Defining AI Literacy.....	41
Activity 2: Evaluating AI.....	44
Activity 3: Integrating AI into Classrooms.....	46
Additional Considerations for AI Literacy	47
Module 9: Updating Policies and Advocating for Responsible Use.....	50

What Is a Responsible Use of Technology Policy?.....	50
How to Refine and Implement a RUP	50
Establish a baseline	50
Design or update the RUP	51
Implement the RUP and periodically improve it	52
Advocating for Responsible Use with Technology Partners	52
Module 10: Ambitious Action Plans	54
Why Build an Organization-Wide Plan for AI?	54
What Should be Included in an Organization-Wide Plan?	54
How can educational leaders define an organization-wide plan?	57
Appendix 1: Key Terms and Acronyms	59
Appendix 2: Building on Foundational Initiatives	61
Appendix 3: Reacting and Shaping Use Scenarios	62
Accessibility	62
Digital Tools Empowering Real-World Accessibility	62
Assessment Accommodations via Emerging Technologies	63
Cyberbullying	63
Deep Fake Images from GenAI	64
Threatening Messages from GenAI	64
Impersonation Campaigns.....	65
Informed Adoption.....	65
The "Wow" Factor	66
The New Kid on the Block	66
"...now with GenAI!"	66
Cost-Effectiveness	67
It's About Assessment	67
Monitoring and Surveillance	68
Keeping an Eye on Them	69
Digital Profiling.....	69
De-anonymizing Anonymous Posts	69
References	71

EMPOWERING EDUCATION LEADERS: A TOOLKIT FOR SAFE, ETHICAL, AND EQUITABLE AI INTEGRATION



Introduction

The U.S. Department of Education (Department) is committed to supporting innovative advances in educational technology (edtech) to improve teaching and learning across the nation's education systems and to support educators as they incorporate emerging technology into their learning communities. Artificial Intelligence (AI) in education is a complex and rapidly expanding topic because individuals have different knowledge, opinions, and perspectives on using AI, which has implications for all members of school communities.

Who are educators and educational leaders?

In this toolkit, we use "educators" to represent all educators. This includes teachers (regardless of certification area), teacher leaders, school principals, paraprofessionals, librarians, school-based mental health professionals, and other school-based educational personnel. Representation of all educators and teaching roles in a district best supports the needs of diverse student populations.

With this representation in mind, the toolkit is meant for a broad swath of educational leaders including superintendents and school administrators, curriculum and technology leaders, school principals, educators, parents and caregivers, and engaged community members. These educational leaders can be found in a variety of learning communities.

What is a learning community?

We use "learning communities" to refer broadly to schools, school districts, and other formal learning organizations, as well as informal learning organizations such as community centers and afterschool programs. The guidance in this document is relevant for multiple settings in which students learn.

Building on the Department's prior report, [Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations](#) ("AI Report"), this toolkit is designed to help educational leaders make critical decisions about incorporating AI applications into student learning and the instructional core. The AI Report provides definitions and illustrates broad classes of opportunities, along with explaining key tensions and risks of including AI. This document sets forth and expands upon the material in the AI Report and connects broad ideas about AI to the establishment of school and district use policies that will guide its effective implementation. This toolkit provides guidance for the effective use and integration of AI in teaching and learning and presents an overview of Federal laws and considerations that are essential to anchoring and ensuring the use of AI in a safe, secure, and non-discriminatory manner. Finally, the toolkit promotes the principles of transparency and awareness in the use of AI in schools, and emphasizes the importance of providing students, teachers, and parents opportunities to opt out of AI-enabled applications in school.

Defining AI

In this toolkit, the term "artificial intelligence" or "AI" has the meaning set forth in 15 U.S.C. 9401(3): a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments, abstract such perceptions into models through analysis in an automated manner, and use model inference to formulate options for information or action.

Defining Generative AI

We also use "Generative AI" (GenAI) in this toolkit. We define this term using the language from [The White House Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#):

"The term 'GenAI' means the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content."

Through listening sessions and attendance at many public events in the past year, the Department heard educators say that AI is here to stay, AI will keep changing, and safely integrating AI in educational settings will require informed leadership at multiple levels across the education system. In particular, this toolkit is informed by public listening sessions held from December 2023 to March 2024 that engaged 90 educators, who actively hold positions in PK-12 schools across the United States (e.g., classroom teachers, instructional coaches, or school

administrators), in broad discussions of the issues. Between December 2023 and March 2024, the Department's Office of Educational Technology also hosted 12 roundtable discussions on the use of AI in education with educators and educational leaders from across the nation to better understand their current perspectives and needs. Where this toolkit refers to listening sessions, it includes all these opportunities to hear from constituents.

This toolkit is organized into three distinct sections, each containing modules that can be accessed and revisited in any order depending on an educational leader's unique needs and priorities:

1. [Mitigating Risk: Safeguarding Student Privacy, Security, and Non-Discrimination \(Modules 1-3\)](#). Awareness of applicable Federal laws, rules, and regulations is an essential first step when planning for the use of AI in schools and classrooms. Educational leaders should know how existing Federal policies apply to the use of AI in their specific situations. This section invites leaders to learn about privacy and data security requirements; how civil rights, accessibility, and digital equity relate to AI; and a close consideration of the opportunities and risks associated with the use of AI. This section is relevant for an educational leader who wants to understand how proactively addressing student safety, privacy, and security can help shape their plans to use AI
2. [Building a Strategy for AI Integration in the Instructional Core \(Modules 4-7\)](#). New forms of AI have already permeated educational settings widely, and exploring AI firsthand is necessary to understanding it. Educators in our listening sessions strongly recommended that districts use the knowledge they have gained from past advances in edtech to build a clear and coherent strategy tied to the instructional core as a first step in planning for the use of AI, and then revising that strategy as they learn more about AI. That strategy should be informed by multiple sources of evidence on the use of AI. Leaders identified three additional steps for further informing their strategy for the effective use of AI-enabled tools in a manner that suits the needs of their students: (1) listen to and inform their communities, (2) establish priorities and pace for their community, and (3) guide and support implementation of a community's strategy via task force. This section provides resources to support educational leaders in considering the evidence supporting AI-enabled tools, and guiding leaders through each of these three essential steps. This path makes sense for an educational leader engaged in or beginning the strategic planning process around the use of AI.
3. [Maximizing Opportunity: Guiding the Effective Use and Evaluation of AI \(Modules 8-10\)](#). Although exploration and building coherent strategy are important early steps, the Department also urges educational leaders to be active in guiding the effective use of AI to enhance teaching and student learning, whether such tools are used for educator productivity or instruction. Educational leaders stressed three initial steps for shaping AI use: (1) developing AI literacy for educators, (2) revising responsible use policies, and (3) building a system-wide plan. This section is appropriate for an educational leader who has a clear strategy in place for the use of AI, and who is ready to focus on guiding, shaping, and continually evaluating the use of AI in their community.

Consider the metaphor of a mountain trek to represent the journey of incorporating AI in education. Like preparing for a challenging climb, achieving AI success requires careful planning, teamwork, and risk management. The trek-themed graphics in the toolkit highlight this proactive approach, reminding educational leaders of the importance of safety, ethics and equity no matter where they are on their AI journeys.

Regardless of which path an educational leader initially takes in this AI journey, we recommend navigating to the other modules in due course because the knowledge, questions, and actions in each of these three sections are designed to reinforce the others, together supporting the effective use of AI in education.¹

¹ This guide is responsive to President Biden's October 30, 2023, [Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#), which states:

To help ensure the responsible development and deployment of AI in the education sector, the Secretary of Education shall, within 365 days of the date of this order, develop resources, policies, and guidance regarding AI. These resources shall address safe, responsible, and non-discriminatory uses of AI in education, including the impact AI systems have on vulnerable and underserved communities, and shall be developed in consultation with stakeholders as appropriate.

MITIGATING RISK: SAFEGUARDING STUDENT PRIVACY, SECURITY, AND NON-DISCRIMINATION



Mitigating Risk: Safeguarding Student Privacy, Security, and Non-Discrimination

The modules in this section will help educational leaders weigh the opportunities and risks associated with AI-enabled tools in schools and enhance understanding of Federal laws, rules, and regulations that safeguard the rights and well-being of students as they use AI in a school setting. The modules end with discussion points and questions to facilitate conversations tailored to the local education agency. Later modules build on this content with an eye toward action. Modules 1-3 cover the following topics:

- [Module 1](#) highlights the range of opportunities and risks.
- [Module 2](#) reviews privacy and data security in existing Federal policy.
- [Module 3](#) covers student civil rights, accessibility, and digital equity.

An educational leader reading this section can ask: How might these safeguards inform and shape my strategy for using AI or AI-enabled tools in my educational setting?

Module 1: Opportunities and Risks

In its AI Report, the Department notes that educational decision-making similarly requires evaluation of both opportunities and risks.

Opportunities

The AI Report highlights numerous opportunities for the effective use of AI-enabled tools to enhance student assessment, address learner variability (i.e., that all learners vary along myriad factors across the whole learner, and that these strengths and challenges can vary by context), or enhance adaptivity of academic content (personalization, differentiation, or individualization). Specific opportunities that applications of AI might also address:

- Adapting instructional approaches and content based on students' strengths and not just their needs.
- Adapting to the needs of groups who learn together and supporting social learning.
- Adapting for students with disabilities, including support for neurodiverse learners.
- Adapting to support learners on active, open, and creative tasks (whereas prior products tended to adapt for short-answer or multiple-choice tasks).
- Adapting to support self-regulation, collaboration, and communication skills in addition to more discrete academic skills.

For educators, the AI Report emphasizes how AI-enabled tools can be used to reduce administrative burden, extend instruction in accordance with a teacher's plan, or make teacher professional learning more fruitful. The AI Report also explores the opportunity for AI-enabled tools to improve feedback loops that inform students and educators about how they can improve their knowledge and skills. Opportunities include using AI to capture better information about what students know and can do, to assist in analyzing that information, and to present the information to students and educators in more useful forms.

In its listening sessions, the Department heard about many opportunities educational leaders are considering to incorporate AI-enabled tools into a strategy for student and teacher supports, including the use of AI image generators to practice understanding concepts; AI-enhanced systems supports that simplify and facilitate strategic systems planning; and AI-enabled tools that assist with lesson planning or improve classroom materials, including language translations. Several states have produced their own policy guidance about the use of AI-enabled tools in education, as have many localities. State guidance reveals additional opportunities for the effective use of AI to translate materials to familiar languages for students and their families, as well as more detail on the administrative labor that could be reduced through use of GenAI (Rochelle et al., 2024). Although it is too early to determine which opportunities will be most fruitful, the AI Report advises educational leaders to focus on using AI as an educational tool to supplement existing structures rather than replacing the role of educators and other traditional educational systems.

The AI Report also advises districts to focus on aligning uses of AI to a shared vision for teaching and learning. The Department heard from many educational leaders who affirmed the importance of a human-centered, strategic vision for education.

Risks

Balancing the multiple potential benefits of AI in education, the AI Report also presented potential risks that are worthy of consideration in the use of AI related to safety, privacy, accuracy of information, fairness, and lack of evidence. Specific tensions were noted as well, such as:

- Between collecting information to personalize vs. enabling inappropriate levels or kinds of surveillance.

- Between a teacher exercising active judgment and decision-making in the use of an AI-enabled tool vs. technology automating educational or instructional decisions.
- Between sharing information for transparency vs. protecting student privacy.
- Between attending to how context varies across education settings vs. scalable solutions for all learners.

Certain uses of AI may pose a higher risk to the rights or safety of individuals when used in educational settings. Below is a non-exhaustive list of examples where the use of AI has posed an enhanced risk to the safety, privacy, or rights of students. In these and other instances, educational leaders should carefully consider whether AI is the right solution. Because AI-enabled systems may offer to automate processes for longer periods of time or with broader scope than in previous generations of edtech, the negative consequences of unmanaged risks with AI-enabled tools can be greater. It is essential that educational leaders employ robust AI risk management practices with any AI-enhanced tool, product, or strategy.

- Detecting student cheating or plagiarism
- Influencing admissions processes
- Monitoring students' physical movements or online activity
- Projecting student progress or outcomes
- Recommending disciplinary interventions
- Determining access to educational resources or programs
- Determining eligibility for student aid or federal education or facilitating surveillance (whether online or in person)
- Utilizing biometric identification, such as facial recognition, to gate access to accounts or assignments, to take attendance, or enhance security

Risks can occur at multiple levels in an educational setting. Some risks are harmful to individual students or educators and may also constitute a violation of laws governing privacy, security, or civil rights. An individual student experiencing discrimination due to a biased algorithm is an example. Other uses and applications of AI may pose risks to the educational system or to norms important in the school culture, such as undermining the role of educators in establishing relationships with students or machine-automated instructional decisions absent important teacher judgment and humans-in-the-loop. Additional risks could undermine trust in the broader learning community. For example, trust and confidence can be undermined if a school or district uses AI-enabled tools or systems that provide inaccurate or biased information to parents and caregivers.

Educational leaders should strengthen their approaches to risk management related to the use of technology in educational settings. One starting point for strengthening risk management comes from [The National Institute of Standards and Technology \(NIST\)](#). This starting point is not specific to education; rather, it addresses the use of AI across industries and applications. To guide discussions about risk management, NIST created the [NIST AI Risk Management Framework](#), which has four elements: Govern, Map, Measure, and Manage.



Govern: Develop a district culture that involves everyone in mitigating risks, enabling use of AI to advance district strategies

Map: Identify both how AI can serve important needs and lead to specific risks by key factors such as grade-level, subject matter, and knowledge of the students and the setting.

Measure: Prioritize risks, and work with AI product vendors and with locally-gathered evidence to measure both the positive impacts and degree of risk mitigation.

Manage: Strengthen capacity to manage implementation of AI tools to realize desired benefits while managing both anticipated and emergent risks.

[Module 10](#) of this toolkit suggests how to use these risk management elements in an educational setting.

The following non-exhaustive list provides real-world examples of various risks associated with the use of AI in schools and educational settings. As AI technologies continue to evolve and as their use expands, additional and unanticipated risks may also emerge.

Transparency About AI Risks: In some cases, educational leaders have not been able to obtain transparent information about how their AI-enabled tools or systems function and collect data including essential information about the tool or system’s performance, evaluation practices, and processes to mitigate bias. This lack of information presents a significant challenge in providing transparency to educators, parents, students, and their wider educational communities.

Bias and Fairness: Some test proctoring systems that rely on facial recognition algorithms have unfairly and disproportionately flagged non-white students and students with disabilities for cheating. AI in education has shown potential to reinforce existing inequities, biases, and unfairness due to biased training data, algorithmic discrimination, and insufficient diversity in AI development. Product developers have more recently undertaken efforts designed to address these issues, including improving data diversity, developing fairness-aware algorithms, and promoting diversity in AI development.

Harmful Content: Students often use AI to create images for their school projects. However, GenAI has been found to perpetuate negative and harmful stereotypes based on race, sex, and disability when given prompts to construct images representing different student cohorts.

Malicious Use: GenAI has been used to perpetuate cyberbullying by students, including creating false and negative narratives about others, or generating “deep fake” fabricated images depicting peers and later posting these images on social media. Some have also used AI to impersonate officials in communicating malicious, false, or harmful messages to the broader school community.

Hallucination Risk and Wrong Information: Despite ongoing research, GenAI has still been found to produce inaccurate or factually incorrect outputs at times, such as narratives describing historical figures who never existed or wrong answers to math problems.

Overreliance: Educators who rely on AI-enabled tools and systems to identify students who may be "at risk" and in need of extra support might sometimes overlook valuable insights regarding context, background, or alternate risk factors not recognized or accounted for in these systems.

Urgency to Adopt: GenAI chatbots became widely available to schools before adequate guidance was provided to help educators and students use them responsibly. This distracted educators from managing the core functions of teaching and learning. Educational leaders have faced pressure to allow various AI-enabled technologies into schools due to rapid interest and the fear of falling behind.

Managing Risks

Additionally, we provide a non-exhaustive list of practices to manage the risks of AI use to protect rights and safety in the use of AI. These practices mirror the minimum risk management practices that Federal agencies must implement in their own use of rights- or safety-impacting AI. For example, the Office of Management and Budget recently put out Memorandum OMB M-24-10, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," which directs agencies to manage AI risks and ensure that AI applications are safely deployed, transparent, and aligned with ethical standards. We recommend similar considerations for educational leaders to ensure responsible AI use in education and to safeguard student rights and privacy. As AI technologies evolve, so too should effective practices for mitigating risk in the use of AI.

1. **Complete an AI impact assessment.** Educational leaders should document the intended purpose for the AI and its expected benefit; the potential risks of using AI; and the quality and purpose of the data used to develop, train, test, and operate the AI.
 - a. Especially for predictive AI models, the decision about which intervention to use based on the AI model is a major factor in ensuring effective and equitable outcomes. Models based on historical data are insufficiently reliable and carry greater risk for impacting rights if used for punitive interventions (such as disciplinary action) or significant opportunities (such as admissions decisions) as opposed to supportive interventions (such as tutoring services).
2. **Test the AI for performance in a real-world context.** Educational leaders should conduct tests to demonstrate that the AI will achieve its expected benefits and that associated risks will be sufficiently mitigated. Testing conditions should mirror as closely as possible the conditions in which the AI will be deployed.
 - a. Educational leaders are also encouraged to leverage pilots and limited releases with strong monitoring, evaluation, and safeguards in place to carry out the final stages of testing before a wider release.

3. **Independently evaluate the AI.** Educational leaders should obtain independent assessments that the system works appropriately as intended and that its expected benefits outweigh the potential risks.
4. **Conduct ongoing monitoring.** In addition to pre-deployment testing, educational leaders should institute ongoing procedures to monitor degradation of the AI's functionality and to detect changes in the AI's impact on rights or safety.
5. **Regularly evaluate risks from the use of AI.** The ongoing monitoring process should include periodic human reviews to determine whether the deployment context, risks, benefits, and agency needs have evolved. Human review is especially helpful after significant modifications to the AI or to the conditions or context in which the AI is used.
6. **Mitigate emerging risks to rights and safety.** Educational leaders should take steps to mitigate risks identified through ongoing monitoring or periodic review. Steps may include updating the AI to reduce its risks or implementing procedural or manual mitigations, such as more stringent human intervention requirements.
7. **Ensure adequate human training and assessment.** Educational leaders should ensure there is sufficient training, assessment, and oversight for operators of the AI to interpret and act on the AI's output, combat any automation bias, and ensure human-based components of the system effectively manage risks from the use of AI.
8. **Provide additional human oversight, intervention, and accountability as part of decisions or actions that could result in a significant impact on rights or safety.** Educational leaders should determine whether there are decisions or actions in which the AI is not permitted to act without additional oversight and implement appropriate human oversight, intervention, and accountability accordingly.
9. **Provide public notice and plain-language documentation.** Educational leaders should provide reasonable and timely notice to affected populations when an AI system is being used. Educational leaders should also provide information to members of the educational community on how the AI was trained and tested, and on how risk mitigation practices were implemented.
10. **Identify and assess AI's impact on equity and fairness and mitigate algorithmic discrimination when it is present.** Educational leaders should determine whether the AI model results in significant disparities in the model's performance across demographic groups. Educational leaders should mitigate disparities that lead to, or perpetuate, unlawful discrimination or harmful bias.
11. **Consult and incorporate feedback from affected communities and the public.** Educational leaders should consult affected communities, including underserved communities, and solicit public feedback in the design, development, and use of AI.
12. **Conduct ongoing monitoring and mitigation for AI-enabled discrimination.** Educational leaders should monitor for AI-enabled discrimination against protected classes during ongoing monitoring.
13. **Notify negatively impacted individuals.** Educational leaders should notify individuals in a clear and accessible manner when the use of AI results in an adverse decision or action that specifically concerns them, such as placement in a particular class or identifying cheating within a student's response to an assignment.

14. **Maintain human consideration and remedy processes.** Educational leaders should provide a fallback and escalation system in the event that an impacted individual would like to appeal or contest the AI's negative impacts on them. These remedy processes should not place unnecessary burden on the impacted individual.
15. **Maintain options to opt out for AI-enabled decisions.** Educational leaders should provide and maintain a mechanism for individuals to conveniently opt out from AI functionality in favor of a human alternative. Opt-out mechanisms should be prominent, readily available, and accessible.

Discussion Points and Questions

These discussion points and questions are intended to catalyze conversations within learning communities regarding Opportunities and Risks.

Discussion Points:

- Discuss why it is important to frame the conversation about using AI in terms of both opportunities and risks.
- Invite members of your learning community to share opportunities they have explored and are excited about.
- Involve members of your learning community in elaborating their concerns about risk. Use examples to illustrate that risks can extend beyond privacy and data security and include significant negative impacts to individual people and the integrity and trust of the learning community.
- Explore how your learning community presently manages risks due to use of technologies and how this approach could be upgraded to address the breadth and depth of risks that may emerge as AI is adopted.

Discussion Questions:

- Which opportunities to use AI best align to values and vision for the educational system that our educational leadership team serves?
- How can our educational leaders obtain information about risks related to each educational opportunity, for example, from students, educators, community members, vendors, researchers, or through local risk study teams?
- How can the recommended minimum risk management practices and the NIST AI Risk Management Framework inform the processes our leadership team will establish locally to manage risk?
- How can our leadership team manage the tensions between desire among students, parents, or staff to use (or not use) AI and incomplete or worrisome information about risks?
- How can we make sure that vulnerable populations are not disproportionately affected by the risks of AI?

Module 2: Privacy and Data Security

In order to ensure the effective use of AI in education, educational leaders must be well versed and current in their knowledge regarding relevant laws, rules, and regulations related to privacy and data security. In order to maintain compliance and safeguard their communities of students, families, educators, and school staff, the Department regularly provides information about privacy and data security legal requirements and best practices through its [Student Privacy Policy Office \(SPPO\)](#). SPPO offers extensive information on protecting student privacy, including protecting students' education records and the personally identifiable information (PII) contained therein.

Visit the SPPO [website](#) for resources offered through SPPO's Privacy Technical Assistance Center (PTAC) that address the [Family Educational Rights and Privacy Act \(FERPA\) and Protection of Pupil Rights Amendment \(PPRA\)](#).

Visit the Department's [Office of Educational Technology](#) website for content and events to support edtech developers in safeguarding student data.

In addition to updating their knowledge and understanding of state and local privacy and data security laws, educational leaders can begin their work on privacy and data security when using AI by reviewing and updating their understanding of the following Federal laws and policies:

- [Children's Online Privacy Protection Act \(COPPA\)](#): Enforced by the Federal Trade Commission (FTC), COPPA imposes requirements on operators of websites or online services directed to children under 13 years of age, and on operators of other websites or online services that knowingly collect personal information online from a child under 13 years of age. Presently, FTC is engaged in new rulemaking which considers how industry developers incorporate children's data in AI technology creation and in the deployment of AI systems, who should be authorized to access and use that data, and how these systems interact with children.
- [Family Educational Rights and Privacy Act \(FERPA\)](#): Educators must consider whether the information about students shared with or stored in an AI-enabled system (i.e., information about student background or interests) is subject to Federal privacy laws such as FERPA. Enforced by the U.S. Department of Education, FERPA governs the privacy of students' education records and the PII contained therein, as maintained by educational agencies and institutions (such as school districts, colleges, and universities) or by parties acting for such agencies or institutions. FERPA applies to agencies and institutions that receive funds under any program administered by the U.S. Department of Education. Among other things, FERPA generally requires educational agencies and institutions to obtain prior written parent or "eligible student" consent before disclosing such records and PII, with certain defined exceptions. Certain edtech products may utilize AI or other technologies to collect lengthy back-and-forth dialogues with

students; absent the proper safeguards, these exchanges and data may contain more sensitive information collected without student or parent consent.

- [Individuals with Disabilities Education Act \(IDEA\)](#): Implemented by the U.S. Department of Education, this law enacted in 1975 governs special education and related services provided to children with disabilities, as well as early intervention services provided to infants and toddlers with disabilities and their families. IDEA contains confidentiality requirements to protect the privacy interests of children with disabilities from birth until age 21 who are referred for services. Like FERPA, IDEA requires that a parent provide prior written consent before PII is disclosed to a third party and that educators obtain informed consent from parents (with some specific exceptions). For example, “alt text” can support understanding information from images, and with GenAI, it will become possible to synthesize or customize alt text to the needs of specific learners; however, issues related to the unauthorized release of PII through use of alt text, such as data being shared with third-party vendors, will need to be addressed. The toolkit discusses IDEA further in the “[Civil Rights, Accessibility, and Digital Equity](#)” module.
- [Children’s Internet Protection Act \(CIPA\)](#): The Federal Communications Commission provides oversight of CIPA, a law that imposes requirements on schools or libraries that receive E-Rate discounts for internet access. CIPA requires schools receiving discounts offered by the E-Rate program to adopt an internet safety policy addressing unauthorized disclosure, use, and dissemination of personal information regarding minors. Such safety policies will likely need updating as the value of data about students increases and cybersecurity threats correspondingly increase. The Federal Communications Commission’s [CIPA Guide](#) offers a more in-depth understanding of CIPA requirements.
- [Protection of Pupil Rights Amendment \(PPRA\)](#): Enforced by the U.S. Department of Education, PPRA affords parents of students certain rights regarding, among other things, participation in surveys and the collection and use of information for marketing purposes. PPRA requires local educational agencies receiving Federal funds from the Department to develop and adopt policies, in consultation with parents, that afford parents the right to inspect any instructional materials (excluding academic tests or academic assessments) used by the local educational agency as part of the educational curriculum, upon request.

As they implement a strategy to guide the use of AI in schools, educational leaders will likely need to seek further guidance on data privacy and security in order to ensure compliance with these and other statutes. Many reputable organizations, including non-governmental and nonprofit organizations, provide additional information that could be helpful.

Monitoring Privacy and Security with AI

Educational leaders have effectively addressed student privacy and security in the wake of prior waves of technological advancement and use of edtech tools in schools. However, with the advent of AI-enabled technologies, issues related to student privacy and data security can become significantly more complicated and involved. For example:

- AI-enhanced products and services may inappropriately or even inadvertently collect or generate student data that contains PII.
- AI-enhanced products and services may not adequately answer the questions parents, students, educators, and others may have about data collection and use (or where necessary, obtain consent from the relevant user).
- AI-enhanced products and services may not be transparent in disclosing how and why they collect student data or PII, and educators, students, or parents may not be provided an affirmative option to opt out of data collection.
- AI-enhanced products may not adequately protect student data, resulting in wrongful use or disclosure of student information.
- AI-enhanced products could result in students under the age of 13 accessing inappropriate or harmful content.

Further, within emerging applications that incorporate AI, these issues may intensify, and new threats may appear. Drawing on information relayed in listening sessions with educators and in public conferences, the Department notes the following enhanced concerns, separate and apart from potential violation of applicable privacy laws:

- AI-enhanced products may collect new and highly sensitive forms of data, such as students' physical characteristics or their whereabouts, and even the questions that students informally ask. This is in contrast to more traditional edtech tools and products, which collected more limited or discrete input, such as keyboard touch and mouse input, in response to structured educational prompts. Disclosure of a student's voice or face could result in harm if, for example, synthesized content shows them falsely expressing something they did not say through methods such as "deep fakes."
- AI-enhanced products may collect extensive data about students' activities and locations, enabling new access into students' lives and broader activities. This data enhances the potential for individual surveillance and, if data are leaked, could result in bad actors' use of these details to harass, bully, or otherwise harm students.
- AI-enhanced products collect high volumes of student data, and applying algorithms to this volume of data may enable making broad inferences about students' future behavior, such as what kinds of products they might later buy. Further, even data without obvious PII may disclose a student's identity because identity can be inferred from a large volume of data points. Experts have already reported seeing an increase in scale and frequency of cyberattacks in schools due to the increasing value of student data for non-educational uses (United States Government Accountability Office, 2022).

- AI-enhanced products may collect data through communications and interfaces that mimic human social interaction, including chats and other human-like forms of interaction. This may result in educational participants being less careful about what they input. For example, a well-intended teacher could inadvertently put sensitive and private information about a student's disability into an unprotected or unencrypted chat interface in order to obtain a customized lesson plan or an Individualized Education Program (IEP), without being aware of how the data could be used by the developer or by others granted access to the data.
- In AI-enhanced products, it may be difficult to inspect how an algorithm or AI system collects and makes use of data, since data is not being collected via more traditional or familiar means such as fixed surveys, fields within the curricular pages, or other pre-determined formats. The data in these systems may be a "black box." For example, an AI-based product might synthesize a new dialogue that obtains PII from students, even if the dialogues previously observed in a product demonstration or inspection did not do so.

Discussion Points and Questions

These discussion points and questions are intended to catalyze conversations about Privacy and Data Security.

Discussion Points

- Identify relevant Federal, state, and local laws and their key provisions. Describe steps your institution has already taken with regard to Federal, state or local laws with AI or prior waves of edtech, and what you plan to do now.
- Discuss community concerns about privacy from the various roles and perspectives within the school district. This should include representation from the district's various demographic groups, educators and staff across different teaching assignments and years of tenure, and community members including students' family or guardians.
- Review past actions to address privacy and data security, such as actions taken in response to data breaches, or previous plans to address potential breaches.
- Elaborate ways in which risks are intensifying in an age of AI. Consider why bad actors may have stronger incentives to breach data privacy, the new kinds of private information that AI-based systems may store, and the quantity of detailed information about students' activities that may be available.

Discussion Questions

- With AI use increasing, what new kinds of privacy incidents might occur in our educational settings? How might guidance from trusted learning organizations be leveraged to design responses when an incident does occur? For example, some educational leaders leverage resources developed by [The Consortium for School Networking \(CoSN\)](#), including their [tools and resources](#).

- For existing products widely used in our schools, how is AI being incorporated in each respective vendor's product roadmaps, and how transparently can they explain any changes to privacy and data security risks? For example, consider the [sample vendor letter](#) shared by [Advanced Learning Partners](#).
- What key privacy and data security risks do we observe when students, educators, and community members use products that include AI? For example, how are risks, such as those in the [Center for Democracy & Technology report](#), being experienced and addressed?
- How might we use existing communication channels to inform our community about privacy and data security risks in a way that respects their concerns and provides information and resources that build confidence about privacy and data security?
- How might we support educators and other adults to protect student privacy and security? For example, implementing resources developed by CoSN including the [Student Data Privacy Toolkit](#).

Module 3: Civil Rights, Accessibility, and Digital Equity

Fairness is an important American value, upheld by laws and educational policies that promote a national culture of equitable systems. Educational leaders should understand interplay between at least three key areas of policy to practice fairness in their schools or districts as it relates to AI: Civil Rights, Accessibility, and Digital Equity. This module highlights Federal laws, rules, and regulations as they relate to Civil Rights, Accessibility, and Digital Equity; there is potential overlap across the three topics.

Civil Rights

The Department's [Office for Civil Rights](#) (OCR) enforces civil rights laws that protect students against discrimination based on race, color, national origin, sex, disability, and age. These laws prohibit discrimination inside and outside the classroom. OCR handles a wide variety of cases and responds to potential violations of civil rights law, including those related to discriminatory discipline, racial harassment, unequal access to educational resources, denial of language services to English learners, and other denials of equal educational opportunities based upon race, color, national origin, sex, and disability. Across all of these civil rights laws, public schools and universities receiving federal funds, as well as private schools and universities receiving federal funds, are required to address discrimination that unfairly blocks students from participating in or learning from educational activities. For example: Title VI of the Civil Rights Act of 1964 prohibits discrimination on the basis of race, color, or national origin; Title IX of the Education Amendments of 1972 prohibits discrimination based on sex; and the Age Discrimination Act of 1975 prohibits discrimination based on age in programs receiving federal funding. Together, these statutes emphasize the need for AI in education to be inclusive, free from bias, and accessible to all across various demographics.

Accessibility for Individuals with Disabilities

On April 24, 2024, the Department of Justice issued its [final rule revising the Title II regulations](#) for implementing the Americans with Disabilities Act (ADA). The rule establishes specific requirements, including the adoption of technical standards to make services, programs, and activities offered through web and mobile applications by state and local government entities accessible to people with disabilities. This includes content provided by public schools and universities. This regulation adopts an internationally recognized accessibility standard for web access, the [Web Content Accessibility Guidelines \("WCAG"\) 2.1](#).

The U.S. Department of Education's Office of Special Education and Rehabilitative Services administers programs and services supporting millions of children, youth and adults with disabilities, including programs authorized under the IDEA. The IDEA assists states and public agencies in the delivery of early intervention, special education, and related services to more than 7.6 million (as of school year 2022–23) eligible infants, toddlers, children, and youth with disabilities. In addition to early intervention services to eligible children, ages birth through 3, the law makes available a free appropriate public education to eligible children with disabilities throughout the nation and ensures special education and related services to those children. The

IDEA also provides substantive protections on various aspects of a student's educational placement and experience, including the development of an IEP, the delivery of appropriate services, review of potential disciplinary actions, and more that may be impacted by the adoption of AI tools in the educational setting.

The IDEA highlights the importance of educational resources that leverage students' strengths and resources to address their needs. The Department provides guidance and technical assistance on improving results and outcomes for children and youth with disabilities served under IDEA and Section 504. This includes resources such as [the Center for Innovation, Design, and Digital Learning](#) (CIDDL), which has provided guidance on the application of AI in education, specifically in special education in their [Inclusive Intelligence: The Impact of AI on Education for All Learners](#) report. Additionally, [Section 504 of the Rehabilitation Act of 1973](#) prohibits discrimination on the basis of disability by recipients of federal financial assistance, and Title II of the ADA prohibits discrimination on the basis of disability by public entities whether or not they receive such financial assistance. In the context of AI, IDEA and Section 504 underscore the need to design tools that are inclusive, adaptable, and compliant with accessibility standards, ensuring fair and effective support for students with disabilities. Learn more about these statutes on the [U.S. Department of Education website](#).

Digital Equity

Digital equity ensures that everyone has fair access to technology, internet connectivity, and digital literacy resources, regardless of socioeconomic status, geography, or other factors. Recent legislation, including the Bipartisan Infrastructure Law and the Digital Equity Act, aims to close the digital divide by funding broadband access and digital skills programs across underserved communities. Achieving digital equity is essential to maximizing the benefits of AI in education, ensuring all students and educators can access and effectively use digital tools to enhance learning.

The [National Educational Technology Plan \(NETP\)](#), along with other related Department publications, includes guidance on addressing digital equity. NETP identifies three distinct but interconnected equity “divides” to better understand the relevant issues, all of which are discussed in more detail below.

- The **Digital Access Divide** addresses unequal access to hardware, connectivity, content, and experiences using technology in education.
- The **Digital Use Divide** addresses differential access to high-quality uses of technology that provide students with meaningful opportunities to explore, create, and engage in critical analysis of academic content and knowledge.
- The **Digital Design Divide** addresses allocations of professional development and resources to involve educators in designing appropriate uses of technology for their students.

The three digital divides outlined in the NETP express themselves in new ways with AI, as follows:

- **Digital Access and AI.** Building and running the underlying models that power GenAI are presently very expensive, and the business models for recovering these costs are not fully established. Additionally, subscription models for premium services are often affordable only to some schools and/or families and not others. Likewise, GenAI-based tutoring models that require high bandwidth access to a Large Language Model (LLM) may generate high per-student costs, which can result in unaffordable pricing for most students. Although GenAI may appear to be free, it is not, and the costs may be distributed in ways that present a barrier to equitable access.
- **Digital Use and AI.** Uses of AI may differ depending on multiple factors, including existing digital access or access to financial resources. In listening sessions and at national conferences, the Department has heard concerns from educators that students in schools serving higher-poverty communities with fewer resources will experience AI as a cost-cutting measure and substitute for human educators and mentors, while those attending schools serving wealthier communities will experience AI as augmenting what human educators provide. Alternatively, AI may be used exclusively in remedial roles with some students, while in more expansive and developmental roles with others. In addition, due to bias in AI-systems, educators are concerned that students with disabilities or students of color may have a higher chance of being accused of using AI inappropriately, such as overusing AI. Educational leaders should seek evidence to ensure that any negative consequences of using or misusing AI do not disproportionately impact students from protected groups. Leaders should seek evidence that AI use is fair to all students.
- **Digital Design and AI.** In its AI Report, the Department recommends centering educators in the design of AI-enabled tools and systems and in how AI is used in education. AI Literacy (see [Module 8](#)) is a prerequisite for building the capacity of educators to contribute in a meaningful and effective manner in the design of educational experiences with technology. Additional professional learning relevant to the specific educational content, pedagogical approach, and technology should also be provided to educators. Engaging educators in working groups or task forces is one way to accomplish this goal. In general, educators need time, access, knowledge, and dedicated resources to participate in design.

These examples point to the general fact that AI is entering learning communities that have existing digital divides impacting student opportunities and learning. For more than a decade, schools have used AI-related systems to predict student test scores, reading levels, and other achievement characteristics. School leaders have seen both the benefits and the potential bias in such systems. Local educational leaders will need to elaborate how the features of their specific system will come into play as AI becomes prominent in their local educational practices—and what they can do throughout their system to advance digital equity.

Connecting AI with Educational Equity

Accessibility and Assistive Technology

[Universal Design for Learning \(UDL\)](#) is a set of widely recognized evidence-based learning principles and a “framework to improve and optimize teaching and learning for all people based on scientific insights into how humans learn” (CAST, 2024). UDL-informed AI tools can support learning for all students by providing multiple means of representation, engagement, action, and expression. Federal laws such as IDEA, the Every Student Succeeds Act (ESSA), and the 21st Century Assistive Technology (AT) Act call for equitable access to assistive technologies for students with disabilities and learning differences. The Department’s Office of Special Education Programs (OSEP) provides resources such as training and consultations to encourage assistive technology use that support student access, learning, and assessment. See the [OSEP website](#) to learn more.

AI can accelerate the integration and use of UDL and AT into teaching and learning to reduce learning barriers, benefitting students with disabilities and neurodiverse learners. Importantly, GenAI is evolving to support multimodality—the combination of various communication modes to create meaning and applies in many contexts. With AI, multimodality presents itself in technology interfaces that can listen to speech, provide captioning, translate languages, and even recognize American Sign Language. Further, AI may assist students with visual impairments by producing reliable “alt text” that describes images. It can render increasingly higher quality speech-to-text and can carry out dialogues with students through conversation and accelerate [UDL](#). In turn, multimodality could power important new assistive technologies.

However, to date, accountability for accessibility and WCAG standards for educational resources is uneven. Therefore, educational leaders may need to urge developers using AI to build on their strengths and meet the needs of their learners.

Bias and Algorithmic Discrimination

With the growing use of AI in schools and its ability to operate on a mass scale, schools may create or contribute to discrimination.

Algorithmic discrimination happens when automated systems treat people unfairly, or negatively impact them, because of their race, ethnicity, gender, religion, disability, or other legally protected characteristics. AI models in education have the potential to exhibit algorithmic discrimination when used for tracking, for example, student preparedness for graduation or student discipline issues, in any predictive way based on historical data. This is because the AI systems were trained on historical data that could proliferate historic and systemic biases that have existed in our education systems and learning communities.

As deployers of automated systems, educational leaders should take proactive and continuous measures to protect students, educators, and learning communities from algorithmic discrimination, and to use systems in an equitable way.

Algorithmic discrimination highlights how biases embedded in technology can lead to unfair treatment, but these biases are not limited to machines. Educators themselves may also exhibit societal and cultural biases and discrimination when evaluating, and disciplining, students on their AI use. While concerns about AI tools, such as ChatGPT, facilitating cheating are widespread among educators, research has so far demonstrated that students do not use AI to as large an extent as teachers assume (Center for Democracy and Technology, 2023). This research has also shown that a minority of teachers have received guidance on detecting student use of tools like ChatGPT for cheating. Together with human bias, this insufficient guidance could result in disproportionate disciplinary actions against historically marginalized students and erode trust between students, families, and educators.

Below are examples of concerns related to bias and discrimination in the use of AI in education, raised by educators and expressed in the Department's listening sessions and public conferences. Similar situations could arise if educational leaders give decision-making autonomy and power to AI systems and tools without monitoring bias and algorithm discrimination.

- AI applications may monitor students and teachers based on their online activity, screen time, and physical activity when using the school internet and devices, and through surveillance cameras with facial recognition and other AI-enabled capabilities.
- An AI application that automates course selection or career recommendations may systematically deter students from pursuing courses or careers that they could succeed in based on historical biases that suggest that "certain students didn't succeed in the past."
- An AI application that detects student behavior, for example, when taking a test or in a school hallway through computer-based vision, may unfairly single out students of color for disciplinary action. An application that employs the AI subfield of computer vision to interpret images or video from a device's camera or a school's security camera may have been trained on material that is biased, for example, against students of color and students experiencing poverty, and has not been reviewed by humans in our current context. Unchecked, these biases may show up in reality, for example, by disproportionately identifying students of color as requiring disciplinary action.
- An AI application that sequences online mathematics modules for a student may direct English learner students to remedial modules more often. The English learner students may know the math but not perform well with the reading load that surrounds the math.
- An AI application could block educational progress by discriminating against students from a particular background by repeatedly failing to recognize their speech in English and asking them over and over to reiterate their input (whereas a human hearing the same speech may understand it the first time).
- An AI image generator might provide images drawn from racist, sexist, and/or ableist stereotypes when students are using the tool for creative assignments.

Discussion Points and Questions

These discussion points and questions are intended to catalyze conversations among your learning communities with regards to Civil Rights, Accessibility, and Digital Equity.

Discussion Points:

- Introduce and explain key concepts such as bias, algorithmic discrimination, multimodality, and digital divides, and how these concepts can clarify the relationship between equity and AI.
- Explicitly consider the biases and inequities that AI tools—specifically those used for monitoring and surveilling student digital, online, and in-person activity—can have on historically marginalized students.
- Provide guidance for AI educators on how to build trust with their students and detect and respond to responsible and irresponsible student use of AI tools. This approach is essential for preventing bias, avoiding the disproportionate disciplining of students from historically marginalized backgrounds, and reducing the risk of distrust between students, their caregivers, and the educator.
- Understand relevant areas of Federal law and policy, e.g., civil rights, accessibility, and the digital divides, and discuss how they currently show up in your local educational setting as you adopt AI and other emerging technologies.
- Ask special education educators to offer examples about how multimodal features could be helpful to specific populations and also what risks should be mitigated for the populations they serve.
- Explore examples of the [three digital divides](#) and discuss how your institution addresses them currently and when emerging technologies are adopted.

Discussion Questions:

- What questions should district leadership ask AI developers regarding their assessment and correction of bias and algorithmic discrimination in their AI products?
- Where are algorithmic biases and discrimination risks most pronounced and potentially consequential as a district's leadership team considers using AI to automate decisions in school systems, including teaching and learning decisions as well as student support and logistics decisions, and monitoring online activity?
- How can district leadership and educators use AI in a way that is not surveilling students but prioritizes their safety?
- How can district leadership provide guidance to educators to effectively evaluate student use of AI and effectively respond to student misuse so that educators are not disproportionately disciplining their students?
- What can a district's leadership team do to address the three divides discussed in the NETP, as they apply in our educational setting?

BUILDING A STRATEGY FOR AI INTEGRATION IN THE INSTRUCTIONAL CORE



Building a Strategy for AI Integration in the Instructional Core

The modules in this section are designed to support educational leaders in establishing a strategy that supports the effective use of AI in advancing the instructional core, providing room to address both current and future AI technologies.

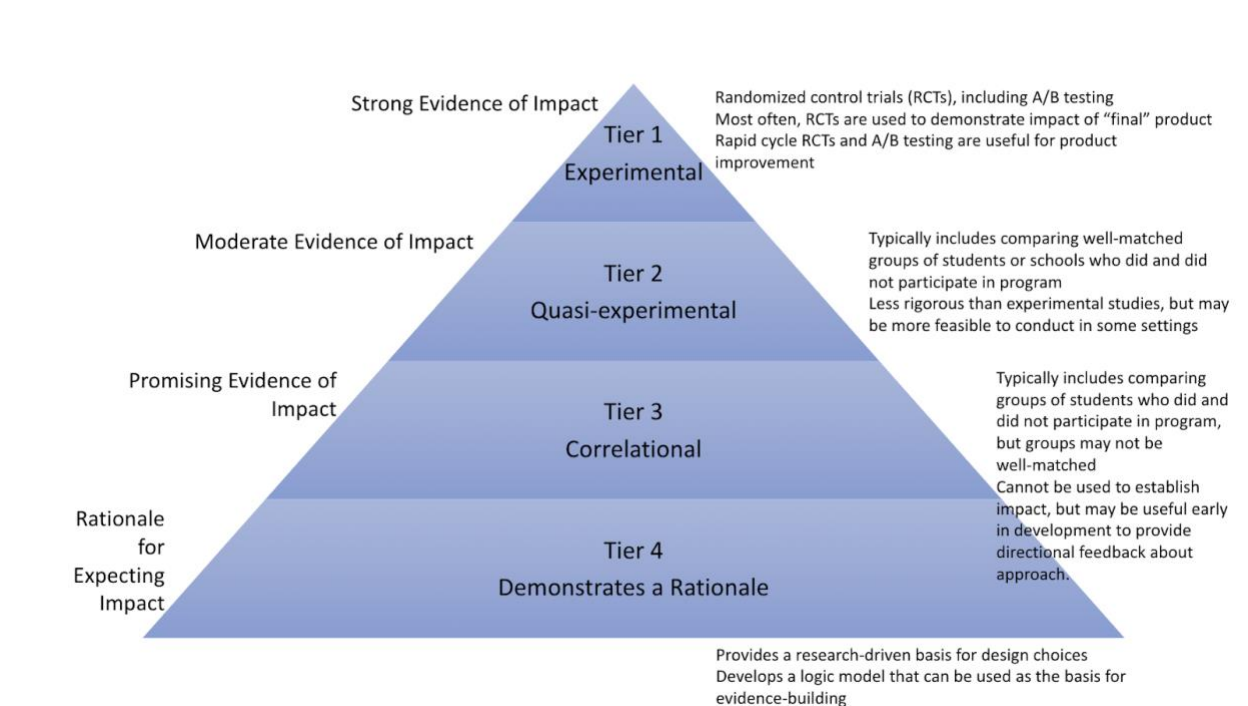
- [Module 4](#) emphasizes the importance of **evidence**.
- [Module 5](#) considers the **instructional core**.
- [Module 6](#) recommends that educational leaders begin their strategic planning by developing awareness and knowledge, through **informing and listening**. Educators can counter the perceived rush to apply AI in diffuse ways, with unknown consequences, by **setting priorities and pacing** for their community.
- [Module 7](#) urges educational leaders to establish a task force to provide responsive **guidance and support** for educators, students, parents and caregivers, and the larger learning community who are using AI in educational settings.

An educational leader who is beginning to read this section can ask: How might I leverage a range of data and perspectives to build a strategy for responsible use of AI in my educational community? How might I deepen my understanding of the evidence and the connection between AI and the instructional core? How might I set a pace for implementation and clear goals for success?

Module 4: Understanding the Evidence

In 2015, the [Every Student Succeeds Act \(ESSA\)](#) reauthorized the Elementary and Secondary Education Act (ESEA) of 1965. ESEA both encourages and requires school and district decision makers to choose evidence-based educational products and services (“interventions”) that have been shown to improve student outcomes through high-quality research and evaluation, including research that demonstrates a rationale for the use of technology. ESEA defines four

tiers of evidence. The highest quality of evidence is referred to as “Strong Evidence,” and the weakest evidence is referred to as that which “Demonstrates a Rationale.” Types of evidence at each tier are described in the graphic below. In considering the topic of evidence, educational leaders should review both the kinds of evidence discussed in Federal policy (discussed below) and also evidence from community members’ lived experience ([Module 6](#)).



Educational leaders should understand the characteristics of sound evidence in each tier and are encouraged to use the evidence tiers as they evaluate adoption and responsible use policies of both AI and non-AI products, tools, and services in their learning communities.

Additional federal resources provide information regarding use of evidence. The [What Works Clearinghouse](#) (WWC) reviews research on educational resources, determines which education research studies meet rigorous standards, and summarizes education research findings. For some applications of AI, such as intelligent tutoring, the WWC has existing reviews detailing key findings within educational research. Newer products, however, may not yet be specifically reviewed. [WWC Practice Guides](#) summarize research-based principles with broad applicability, and these can be very helpful for understanding principles that could be incorporated into a product’s design. In addition, the Office of Educational Technology offers an [Edtech Evidence Toolkit](#) with resources including one-pagers, case studies, and examples to support educational leaders in making evidence-informed edtech adoption decisions within schools. More generally, Federal agencies including the Department, the [Institute of Education Sciences](#), the [National Science Foundation](#), and the [National Institutes of Health](#) fund research to evaluate and publish technological approaches to improving education.

For educational leaders who may find it difficult to comb through the extensive literature on their own, forming a research-practice partnership with a researcher may be helpful. Research-practice partnerships are collaborative relationships between researchers and practitioners that aim to address real-world challenges in education through the practical and collaborative application of research in educational settings (Coburn & Penuel, 2016; Farrell, et al., 2022). The Department's Institute of Education Sciences, the National Science Foundation, and certain philanthropic organizations (such as The Lumina Foundation, The New School Venture Fund and The Bill & Melinda Gates Foundation) all fund research-practice partnerships. Research-practice partnerships between educational leaders and researchers are typically designed to a) promote evidence-based practices in schools, b) bridge the gap between academic research and practical application in educational settings, and c) facilitate the translation of research findings into actionable strategies to positively impact learner outcomes.

Connecting Evidence to AI-Enabled Interventions

Research on AI in education has a 50-year history, providing sound evidence for some applications of AI in teaching and learning. For example, many in the field have identified AI-based tutoring as a promising application, and Intelligent Tutoring Systems (ITS) have been studied for a long time (Steenbergen-Hu & Cooper, 2013). Indeed, meta-analyses (statistical summaries of many independent studies) exist for AI-based mathematics tutoring (Kulik & Fletcher, 2016; Xu et al., 2019).

Successful interventions described in the literature might include elements that newer GenAI-based tutors lack. For example, GenAI-based tutors may not have information about how to sequence instructional content, how to diagnose common student misconceptions, or how to identify strengths in student knowledge that can enable the student to learn a target idea. In another instance, the literature reveals (VanLehn, 2011) that tutors are especially beneficial when they can intervene on a specific problem-solving step that a student got wrong, rather than intervening only when the overall answer to a problem is wrong. How to use AI and machine learning to intervene appropriately with different students is an active research issue (Gao et al., 2024) as is the more general issue of how to incorporate evidence-based pedagogies into the response strategies of GenAI-based tutors (Weiss, 2024).

Because newer GenAI tutors may not intervene at the detail level of a problem-solving step (Chine et al., 2022), prior evidence may be helpful in probing areas where the newest products may fall short of what has been shown to work.

Many applications of advanced technology are initially exciting, but when evidence is collected to evaluate the application, no or minimal benefits are detected. Educational leaders should carefully look at the quality of evidence for novel applications of AI in education. Given the rapidly changing landscape of AI-based products, evidence at any of the four tiers could be useful to support educational leaders' decision making. We provide examples below:

- **Tier 4 (Demonstrates a Rationale):** Educational practices, interventions, program components, and tools at the Tier 4 level have a well-defined logic model, are supported

by prior research with positive findings (such as the WWC Practice Guides) and have efforts underway to determine their effectiveness. Methods incorporating literature reviews, development of logic models, pilot testing, and surveying/user feedback are all examples of developing a Tier 4 evidence base.

- **Tier 3 (Correlational):** Correlational studies can be useful to look for “no harm” and for suggestive evidence of benefit. For example, educational leaders can use correlational evidence to examine whether the intensive use of a new AI-based product is associated with negative or positive outcomes, both for students overall and within specific student groups. Likewise, correlational studies can reveal whether using a product for a recommended or longer period of time is more strongly associated with benefits than using the product for a minimal amount of time. Positive correlation between the recommended usage and desired outcomes suggests benefit.
- **Tiers 1 and 2 (Quasi-Experimental and Experimental):** Studies at these levels compare a new approach (e.g., using AI) to an existing approach, either by randomizing students to the two approaches (Tier 1) or comparing well-matched students who use the alternative approaches (Tier 2). Although these studies have a reputation for being costly and difficult to conduct, faster and less-expensive options are increasingly common. For example, a district may have adopted a literacy solution in its non-AI form, and the literacy product may have a new version with an AI-based chatbot. It may be feasible to conduct an A/B comparison with and without the chatbot at low cost and modest complexity. Educational leaders may prefer to wait for rigorous Tiers 1 or 2 evidence to become available over time before adopting AI-enabled products.

Discussion Points and Questions

These discussion points and questions are intended to catalyze conversations among your learning community regarding Evidence.

Discussion Points:

- Discuss why using evidence is important and how your educational organization presently uses evidence in decision making.
- Share ways in which your leadership team can learn more about how to use evidence to support edtech adoption.
- Share evidence for AI-enabled products that are being considered for adoption or purchase and evaluate the quality of this evidence, identifying what information is missing/desirable for making a sound decision.
- Consider how you might collaborate with researchers (e.g., from a local university or nonprofit) to evaluate evidence for AI-enabled products.

Discussion Questions:

- What are your educational leadership team's requirements for evidence for different ways of using AI (and all edtech) in your setting?
- What low-cost and timely evidence gathering activities can be integrated into existing procurement processes to support evidence-based edtech procurement decisions?
- What steps can be taken to move toward building a stronger evidence base for use of AI in schools, including not only efficacy on average but also avoiding potential harm to any student?


Module 5: Considering the Instructional Core

Defining the Instructional Core

Richard Elmore (2008) defines the instructional core as the connection between students, teachers, and content, grounded in high-quality instructional tasks. To improve student learning at scale, it is essential to address the instructional core by focusing on three key principles:

1. raising the level of content taught to **students**
2. enhancing **teachers'** skills and knowledge related to that content
3. increasing students' active engagement with the **material**

Together, these efforts form a strong triangle that can drive meaningful improvements in education outcomes.



The concept of the “instructional core” pinpoints the difference between innovations that deeply advance student learning and those that merely scratch the surface of educational improvement. In another foundational article, Cohen and Ball (1999) write: “Since World War II, efforts to improve schools have numbered in the thousands.... Unfortunately, three decades of research has found that only a few interventions have had detectable effects on instruction and that, when such effects are detected, they rarely are sustained over time” (p.1). They propose

that meaningful improvements require reconfiguring instruction, which they define as the interactions among teachers, students and curricular resources. For students, this means they are central to the learning process, actively engaging with challenging content and meaningfully engaging with skilled, knowledgeable educators. For teachers, this means enhancing their expertise, pedagogical strategies, and instructional practices in order to effectively teach and engage students and help them learn and master rigorous content. For educational leaders, this means providing professional learning opportunities, tools, and targeted mentoring that specifically advance teachers' instructional skill and capacity and impact the level of student work in classrooms.

Multiple states and school districts have centered the instructional core in their work. For example, Nebraska's approach observes as a first principle: "Increases in student learning occur only as a consequence of improvements in the level of content, teachers' knowledge and skill, and student engagement." And Kentucky provides [Characteristics of Highly Effective Teaching and Learning](#), a shared framework for discussing best practices for creating ideal learning environments.

The Instructional Core and AI

The Department sees the potential for effective use of edtech to enhance and strengthen the instructional core. There are, of course, multiple examples of the use of AI-enhanced tools or applications that fall short of meeting this potential. For example, some applications of AI generate relevant or attractive content for lesson planning but might not be accurate or improve the quality of the content. Some applications of AI can save teachers' time but fall short of providing them new knowledge, skills, or strategies that will enhance their interactions with students in scaffolding robust content. And as every teacher knows, technology can sometimes be used in a manner that increases students' engagement superficially without promoting the skills and growth mindset that students need to exercise critical thinking and reasoning and meaningfully exert effort and persistence around an instructional task.

To apply AI with purpose, educational leaders must be attentive to avoid superficial or inconsistent applications of AI and instead maintain a focus on how AI can be utilized as a tool to improve the quality of teaching and learning, student-teacher interactions, and instructional tasks at the heart of the instructional core. For example, in teaching reading, the National Center on Improving Literacy (2022) argues that by integrating comprehension with decoding while **building on students' prior knowledge**, educators can accelerate reading growth. An effective use of GenAI might **revitalize the content students are asked to read** while also enhancing **how students and teachers interact** while they read. In mathematics, students need practice using concepts to explain mathematical ideas, but too often many mathematics assignments remain focused on practicing low-level procedures (TNTP, n.d.). An effective use of GenAI might **provide students more opportunities to practice explaining** core mathematical concepts, integrating both student feedback and supplemental resources for teachers to help them develop expertise in connecting student reasoning with mathematical concepts. Next Generation Science Standards (NGSS Lead States, 2013) set expectations that classrooms will

increase attention to “cross-cutting concepts,” which connect individual units of study in the core subject of science. An effective use of GenAI might connect the units of study to broader scientific concepts and other academic subjects, like social studies or math.

Centering the instructional core requires interconnected changes, some of which reach beyond the classroom triangle in which teachers, students, and content interact. Assessment is a huge factor in what changes or stays the same related to instruction. The effective use of AI can transform how teachers and students make use of assessment results for instructional planning, decision-making, and the implementation of effective instructional interventions. In our earlier AI Report, the Department observed the opportunities for AI to enable measuring what matters while also observing that the assessment community has strengths in addressing fairness and mitigating bias that will need to be brought to bear on AI-enabled assessment. Educational leaders play an essential role in supporting classroom **educators to make the right connections** between assessment systems and improvements to the instructional core.

The Department’s AI Report notes that many educator preparation programs attend to technology in superficial or inconsistent ways. Of course, this also pertains to in-service teacher professional development. For educators to integrate AI into the instructional core in their classrooms, they will need more time and higher-quality professional learning designed to develop and enhance their expertise with using edtech. Educational leaders can influence change by asking hard questions about whether and how AI literacy for teachers can help advance the instructional core.

Finally, the AI report emphasizes the potential of AI to reduce the administrative burden on teachers, giving them more space to focus on their students and on instruction. The potential to make teachers’ lives better can be an initial motivator to explore the use of AI; however, teachers’ aspirations will quickly seek to expand the potential for AI to support effective instruction, beyond enhancing time management and productivity. Educational leaders have a role in keeping a focus on advancing the instructional core as additional time or capacity becomes available.

Discussion Points and Questions

These discussion points and questions are intended to catalyze conversations that connect the opportunities surrounding AI to local efforts to improve the instructional core.

Discussion Points:

- Discuss how AI-enhanced products complement the instructional core as a catalyst.
- What are “higher level” or better-quality tasks that students, with the use of AI-enabled tools, can successfully do?
- If AI supports students in some interactions as they work on these tasks, what are the most valuable ways teachers should interact with students?
- How can teachers learn new knowledge, skills, and instructional strategies using AI?
- Share ways in which your leadership team can connect current or potential uses of emerging technologies with its priority on the instructional core.

- Create guidance for selecting AI-enhanced products that are not pure supplements to instruction, but rather support desired instructional changes among teachers, students, and resources.
- Explore the additional factors in your setting – beyond students, teachers, and resources (e.g., data from edtech tools and digital platforms and the interpretation and analysis of those data to answer administrative questions) – that most strongly influence the instructional core and how AI might enhance these factors or pose new risks.
- How might you increase reflection and action to create the conditions necessary to harness the potential of AI?

Discussion Questions:

- Imagine an observation in a classroom that is using AI. What does success look like?
- How are students engaged in tasks, what tasks are they doing, what are they doing and saying?
- How are students supporting each other?
- What are teachers doing and how are the knowledge and the level of skills they are bringing to the instructional process improving?
- What are their families saying and noticing?
- What guidance and examples ([Module 9](#)) are your educational leadership team providing to ensure coherence among the different ways of using AI (and all edtech) as it relates to the instructional core?
- How can staff development and collaborative spaces prioritize interconnectedness between instructional core and AI, rather than seeing these as separate issues?
- When one element of the instructional core is changed, how are you balancing the others to ensure consistency?
- What cost-effective and timely professional learning experiences can support your team as they prepare to leverage AI-enhanced tools in their practice?
- How are educational leaders examining the comprehensiveness of teachers' unique roles and opportunities to learn?
- How are supervising teachers receiving additional support to coach and assist student teachers on their use of AI in relationship to the instructional core?
- How can teams in your educational setting use research to understand the differences between promising uses of AI-enhanced tools and those which are likely to fail to meaningfully change the instructional core? Further, how can research and staff expertise enable consideration of equity divides (e.g., in the NETP) throughout this process?

Module 6: Planning Your AI Strategy

Inform and Listen

The Department, through listening sessions, learned that educational leaders are experiencing wide variation in their communities' knowledge about AI, their excitement and concerns, and their views on important next steps.

As a first step, educational leaders can host a listening session in their own communities to engage a variety of voices and involve invested community members in developing guidance and recommendations. As discussed in the Evidence module ([Module 4](#)), the experiences of an educational community are also evidence that should be weighed while developing a strategy for responsible use of AI. [Module 5](#) provides insights into the process of hosting a **virtual** listening session. A virtual listening session commonly lasts 60-90 minutes, and it is advisable to hold multiple sessions for various audiences within your educational community.

Preparing for the Listening Session

The Department anticipates that individual educational leaders have their own best practices for planning meetings, including, for example, what format to use (virtual, physical, or hybrid), who to invite and how to invite them, what accommodations to provide for those with special needs, and how to capture notes. As many educational leaders already know, an event agenda can be a useful organizing document for everyone who will participate in conducting the event.

Some specific considerations for a listening session about AI may include:

- **Local experts:** Consider featuring leaders, educators, students, or community members that represent the various demographics of the school ecosystem who have knowledge, expertise, and the ability to talk about AI in terms that are likely to be clear and accessible to those who will be participating.
- **Concrete examples:** Consider inviting educators or students who can provide examples of how they have used AI in responsible ways, or conversely, establishing clear limits on when AI should not be used.
- **Balanced panelists:** Consider whether a panel might be able to set a responsible tone for discussing the range of opportunities and risks with AI, and thus model productive discourse about AI. This panel should include a range of people that best represent the school ecosystem. This might include, for example, those who can share a family perspective, those who identify as having a disability, and out-of-school-time programs who work closely with the district.

In addition, educational leaders may wish to consider:

- Who is the **audience** for this listening session?
- Should the listening session be **invitation only** or **held as a town hall**?
- What **information** about AI, edtech, and education will the session leader share and receive during this session?
- Who can serve as **trustworthy** and **knowledgeable** facilitators, panelists, or speakers to engage this audience? What is their connection to the audience?
- What **questions or concerns** should the organizers collect before the listening session so that leadership can be better prepared?
- How can the organizers facilitate **interaction** while managing challenges that could arise during a session? What are the various modes of engagement that the organizers should consider?

During the Listening Session

The listening session should prioritize hearing from district and community members. Educational leaders can set expectations for the participatory aspects of a listening session via an opening presentation. In planning the opening presentation, consider how the presentation addresses the following considerations:

- **Purpose:** State why you are holding the listening session; one example might be to engage the community in building a strategy to guide both current and future uses of AI in schools. Describe what the district has already identified as key principles, values, guidelines, or commitments. Explain how insights from the listening session will be used.
- **What is AI?** Establish some common ground in the audience's experiences of AI (for example, automated maps or phone-based voice assistants), accessible definitions (for example, "automating based on associations" from the AI Report), and clear examples. Share what kinds of uses of AI are already occurring in the school community.
- **Where is the district in the development of a strategic process?** For example, what policies, guidelines, and guardrails already exist (for AI, edtech, and non-technology issues)? On which principles or values is there agreement? What are some of the new challenges?

Educational leaders should keep the following in mind when hearing from the audience:

- **Responses:** How are community members responding to the integration of AI? It can be helpful, for example, to acknowledge the spectrum of concern-to-enthusiasm that participants may bring to the listening session.
- **Moderated interaction:** Consider whether accepting questions and comments via notecards, an electronic poll, or other means might be preferable to allowing anyone to speak on a live microphone.
- **Call to action:** Help participants understand how they can further contribute to solutions. Explain how their feedback will be applied and how they can continue to provide guidance.

Here is an example of a [planning document](#) that you can use and share with all speakers and members of the planning team.

Outline materials in an [event agenda](#) and share the document with all speakers for visibility.

After the Listening Session

Audience members should not leave a listening session without follow-up, especially after sharing their comments and experiences. Educational leaders might consider the following to acknowledge that the audience's time and input is valued and will be used for future development:

- **Send a thank you message** to attendees.
- **Summarize** the session, potentially using an AI tool to help identify themes that are then reviewed by humans.
- **Share** a brief recap of key insights and lessons learned from the listening session with attendees. Include opportunities for continued engagement and other next steps.

Pace is a community agreement often dictated by consensus and governed by the needs of a team’s most vulnerable members. Both in its listening sessions and at public conferences, the Department learned of educational leaders’ concerns with the pace of AI advances.

Furthermore, educational leaders, parents, educators, and students all expressed reservations about how past technology breakthroughs have impacted their communities. In light of past experiences with previous edtech integration, educational leaders in listening sessions expressed desires to set priorities and to regulate the pace of adoption of newly emerging technologies. Rather than “letting a thousand flowers bloom” by allowing students and educators to introduce whatever applications seems attractive, leaders seek a more focused and strategic approach to integrating AI into their settings.

The Department encourages educational leaders to intentionally set priorities according to a shared vision for education (as discussed in the second recommendation in the [AI Report](#)). Further, the Department encourages educational leaders to establish a pace and support their community in saying “no” or “go slow” as conditions warrant.

Planning an Approach to Prioritization

GenAI is being used in every type of edtech and for wide-ranging educational applications (e.g., teaching and learning, assessment, counseling, school safety, logistics, operations, and more), and the market is moving at a blistering pace. Because of this rapid movement, priorities and pace should not come from the technology itself.

In its AI Report, the Department recommended starting from a shared educational vision, which is likely stated in an educational institution’s mission, strategy, portrait of a graduate, and other documents. In addition, many states, including [Washington](#) and [Delaware](#), are now producing guidance specific to AI. The nonprofit sector is providing additional input. Reviewing the [three introductory modules](#) is also a good starting point to ground planning in a strong foundation that mitigates the risks of AI in education described in Section 1 of this toolkit.

Each educational institution will likely need to develop its own approach to prioritization based on how it expresses its mission, vision, strategy, and other key input.

Discussion Questions

Educational leaders can ask:

- What are our most important unmet educational goals that AI can help resolve?
- Where are we starting in terms of responsible use of edtech?

- What has worked previously in terms of a pace for responsible adoption and effective technology implementation?
- How do we ensure attention to important problems that do not have immediate technological solutions (e.g., truancy) with attention to problems that are a fit for AI and other emerging technologies?
- What is our human resource capacity and readiness to implement AI in the best possible ways?

Specifying Factors

The Department recommends that educational leaders list key topics on which their teams will gather information before making decisions about priority and pace. A partial list of considerations could include:

- **Alignment to Educational Strategy:** How tightly aligned is this opportunity to use AI to the educational institution's vision, mission, strategy, and goals?
- **Risk Analysis:** Considering the areas of concern described in Modules [1](#), [2](#) and [4](#), what are the risks with this use of AI in our educational setting, and to what extent are those risks managed?
- **Content Moderation:** To what extent will the AI tool impact student learning due to the data used to train the underlying model?
- **Evidence:** What evidence (see discussion in [Module 4](#)) supports prioritizing this opportunity to use AI?
- **Capacity and Readiness:** To what extent can we engage the necessary leadership, resources, and talents in our educational community to support success?

Educational leaders may wish to consider four or more levels of adoption such as:	
Encouraged Use	Some proposed uses of AI emerge as high priority given the factors that were determined to be important in the educational setting, and may be encouraged.
Allowed Use	Some proposed uses of AI may be permitted, but not especially encouraged, and leaders may wish to monitor how much use occurs and what the impacts are.
Limited Use	Some proposed uses of AI may be worth exploring in a bounded and closely monitored way.
No Use	Some proposed uses of AI will be unrelated to important educational goals, too risky, lacking evidence, or too hard to implement well.

Educational leaders should set expectations on pacing for how soon and how much use of AI is expected, for example, compared to existing resources and approaches. **Likewise, allowing enough time to evaluate, reflect, and adjust plans for AI use is important.** Pace may also be reflected in how many AI opportunities are allowed or encouraged.

Module 7: Guide and Support

In implementing a strategy to integrate AI in your school, there will inevitably be situations where the challenges faced exceed present capabilities. In such cases, seeking external assistance may become necessary. This paradigm shift into the GenAI era is a similar venture that may require assistance to get started, to maintain, and to complete successfully. As is widely reported in educational news media, educators, students, parents and caregivers, community members, and others in learning communities are asking for guidance about responsible AI use. States, districts, nonprofits, and other organizations are responding. Building on their recommendations, the Department urges educational leaders to establish a task force in their institution (or another similarly chartered council, committee, or team) to respond on an ongoing basis. A task force can provide more specific guidance and support to educators and others about responsible uses of AI that are well aligned to the institution's priorities.

Planning a Task Force

The Department expects the needs for support and guidance to vary across settings, and thus there is no single charter for a task force that will work in all settings. Questions to ask include:

- **Who will be included in the task force?** The process of choosing task force members will depend on the resources available in individual districts. No matter how they are chosen, it is important to include diverse experiences and voices. This can either be accomplished through an open call for volunteers, applications from interested community members, or strategic selections to ensure diversity. If the latter, the selection process must be transparent, consistent, and fair. Members of the task force must have an understanding of the particular risks that might arise in the use of AI, as identified in previous modules (see the [Privacy](#) and [Civil Rights](#) modules).
- **Can the task force be integrated into existing district structures?** Some districts may not have the resources for a separate task force; also, interested members in a task force may not have the time and bandwidth to join an additional group. Districts could explore existing spaces where the task force can be added, such as district academic support teams.
- **How will task force members be compensated?** It is not easy for educators, students, and community members to add more time to busy schedules to attend task force meetings and accomplish task force duties. If districts consider compensation in some form, it could, for example, be monetary, include extra credit for students, or release educators from other responsibilities.
- **Who will the task force support and guide?** A task force might be chartered mainly to support educators, or to help building-level decision makers, or to work with students and families, or some combination of these groups.
- **What will be the task force's charge?** A task force might help educators implement an acceptable use policy, given emerging applications of AI, or to address equity divides (see [Module 3](#)), or to vet possible applications of AI for broader adoption and implementation.

- **What knowledge will the task force use?** A task force will need input on the educational institution's desired pace and priorities, as well as information garnered from listening sessions, and guidance from leadership.
- **What resources will be available to the task force?** How much time will they have to meet and work? Will they be able to conduct surveys? Will they be able to go to conferences or attend professional development trainings to increase their own knowledge about AI? Who will help the task force gather evidence from research?
- **How will the task force provide support and guidance?** Will members of the task force each work with members of the educational community in their specific building or team? Will the task force provide professional learning to educators? Will they conduct awareness or other campaigns? Or will they provide recommendations to others who have responsibility for support and guidance?

Ethics as a Cornerstone

To build public trust and ensure that AI tools are integrated in ways that respect individual rights, promote fairness, and prevent potential harm, an AI task force needs to intentionally center ethics. The Department suggests educational leaders provide ethics training to a task force, framing the ethical adoption, application, and use of AI as a cornerstone of any school's or district's AI strategy. This includes familiarizing task force members with broadly applicable ethics concepts and frameworks as well as those specific to AI and education. Then, task force members should make explicit the ethical considerations they want to hold at the center of their process. A shared ethical purpose is critical for the complex work of adopting AI in learning communities.

In February 2024, researchers at NIST [suggested](#) building on long-standing concepts in the [1979 Belmont Report](#), which provides and develops guidelines on the basic ethical principles that should underlie the conduct of human research. These long-standing concepts include beneficence, respect for persons, and justice, which can be used to organize an approach to ethics in the age of AI. Researchers and educators have specifically been working together to develop ethical guidelines, and the Department expects this work to continue.

Additionally, the Department's review of the variety of ethics frameworks for implementing AI solutions, including the [2024 NIST AI Risk Management Framework](#), found that general ethics and AI ethics concepts apply equally to education: equity, sustainability, transparency, justice and fairness, non-maleficence, responsibility, privacy, beneficence, freedom, and autonomy.

Value-Centered Design Ethics in Education

General Ethics Themes	Education Ethics Themes
<ul style="list-style-type: none">• Transparency• Justice and fairness• Non-maleficence• Responsibility• Privacy• Beneficence• Freedom and autonomy	<ul style="list-style-type: none">• Pedagogical appropriateness• Children's rights• AI literacy• Teacher well-being• Equitable instruction & access

Beyond adapting generally accepted ethics themes to more closely fit education settings, the Department's review of the 2024 NIST AI Risk Management Framework identified four additional principles specific to education: pedagogical appropriateness, children's rights, AI literacy, and educator well-being. Student well-being was not specifically listed in the review; nonetheless, the Department finds student well-being to be of paramount importance. These ethical principles are core components of designing AI systems that interact with children.

Creating a Charter

The Department found that freely available GenAI tools can create an initial outline for a task force charter. Below is a sample charter that was first drafted using Google's Gemini, and then substantially modified by the authoring team. The Department chose to include this example to illustrate the steps that educators may take when using AI to help with their work. The detailed content within the example is not directly from GenAI but reflects additional edited content and expertise from the authors.

Box: Example of an Artificial Intelligence (AI) in Education Task Force Charter

I. Introduction

The [School District Name] School Board recognizes the growing potential of Artificial Intelligence (AI) to create opportunities to improve education but also to introduce both anticipated and unanticipated risks. To ensure responsible and effective integration of AI tools, the School Board hereby establishes the AI in Education Task Force.



II. Mission

The mission of the AI in Education Task Force is to document the opportunities and risks of AI in our district. The Task Force will develop a plan for the ethical, equitable, and sustainable use of AI to enhance student learning (and may later address support for educators and contributions to district operations). The Task Force will respond to issues that arise as students and educators use AI for learning.



III. Goals

- Identify current and potential uses of AI for student learning in the district.
- Gather evidence on the uses both from published research and from experience in the district.
- Collect information on topics related to the risks of these uses of AI including data privacy, algorithmic bias, content moderation that impacts learning, and equity risks.
- Create an initial risk management plan based on the information collected in the previous bullet point.
- Recommend professional development opportunities for educators and staff regarding using AI to support student learning.
- Create a communication plan to keep our community informed about the Task Force's work.



IV. Membership

The Task Force will be composed of a diverse group from our community, including:

- District Administrators
- Curriculum Specialists
- Technology Experts
- Educators (representing various grade levels and subjects)
- Students
- Families, Caregivers, and Community Representatives

The Task Force may invite additional members with specific expertise as needed.



V. Meetings and Deliverables

- The Task Force will meet [frequency] for a period of [duration].
- Meetings will be open to the public (optional, depending on meeting format).
- Minutes will be recorded and posted publicly (with sensitive information redacted if necessary).
- The Task Force will produce these deliverables:
 - An inventory or classification of the uses of GenAI for student learning in our district
 - An analysis of what quality of evidence for these uses exists and what kinds of evidence is most strongly needed

- A discussion of risks and an initial risk management plan
- Recommendations for professional development and other next steps



VI. Support and Resources

The School Board will provide the Task Force with the necessary resources to fulfill its mission, including:

- Staff support
- Meeting space
- Compensation to members in the form of [...]
- Access to relevant data and research
- Budget for potential expert consultation



VII. Review and Revision

This charter may be reviewed and revised at any time by the Task Force and with the approval of the School Board.

VIII. Conclusion

By harnessing the power of AI responsibly, we aim to create a more engaging and effective learning environment for all students.

MAXIMIZING OPPORTUNITY: GUIDING THE EFFECTIVE USE AND EVALUATION OF AI



Maximizing Opportunity: Guiding the Effective Use and Evaluation of AI

This section addresses the implementation and use of AI in schools. Following the development of a strategy and plan for pacing and introducing AI to a school community, these three modules provide guidance on key actions that educational leaders can take when using, managing, evaluating, and scaling AI in their districts and schools.

- [Module 8](#) supports educational leaders to **build AI literacy for educators**.
- [Module 9](#) discusses top level considerations for educational leaders **updating their policies for acceptable use of technology**.
- [Module 10](#) suggests how educational leaders can create an **organization-wide action plan** for implementing AI in their educational communities.

Each module includes corresponding scenarios in [Appendix 3](#) for educational leaders to practice applying what they have learned in real-world contexts. An educational leader who is reading this section can ask: What key ideas should be included in AI literacy, acceptable use, and risk management policies?

Module 8: Building AI Literacy for Educators

The responsible use of AI requires building a strong foundation and knowledge of AI literacy. This includes the knowledge, skills, and attitudes needed to engage with AI safely and effectively.

AI literacy for educators was a top recommendation from educators in the Department's listening sessions and is also a priority for leaders across Federal, state, and local levels.

Developing AI literacy is essential both for students and for educators. This toolkit recommends prioritizing AI literacy and capacity first for educators, then empowering them to do the same

for students and others in their school communities. Developing educator AI literacy is important because:

- With greater knowledge about what AI is, how it works, and which uses of AI are recommended in classrooms, educators can make **informed decisions** about how to integrate AI in their classrooms.
- With greater awareness of how issues of ethics and risks arise when using AI, educators can **advocate for safety** among their students and broader communities and undertake their role as humans-in-the-loop as AI becomes more prevalent in their students' experiences.
- With a forum to develop understanding of how AI may impact the workforce, and the world students will enter, educators can better **guide students toward successful futures in a world where change is occurring at a fast pace.**

Defining AI Literacy

AI literacy includes the knowledge, skills, and attitudes needed to engage with AI safely. (Mills et al., 2024). This skillset is about understanding AI's strengths, limits, and impacts to make informed decisions about its integration and to prepare learners for the AI-driven future (Kulesa et al., 2024).

As an AI literacy initiative progresses, all educators should be able to **actively recognize** AI in multiple forms, make sense of their own **interactions with AI**, make informed decisions based on an **understanding of how AI is built** and its limitations, and view AI through a lens that includes **ethics and social impact**.

Key activities to address in AI literacy programs include **Understanding, Evaluating, and Integrating into the Classroom**. Each activity is discussed in a section below. Educational leaders are expected to make their own decisions, based on their specific context, about how to best cover the breadth of content over time. Activity 1: Understanding AI

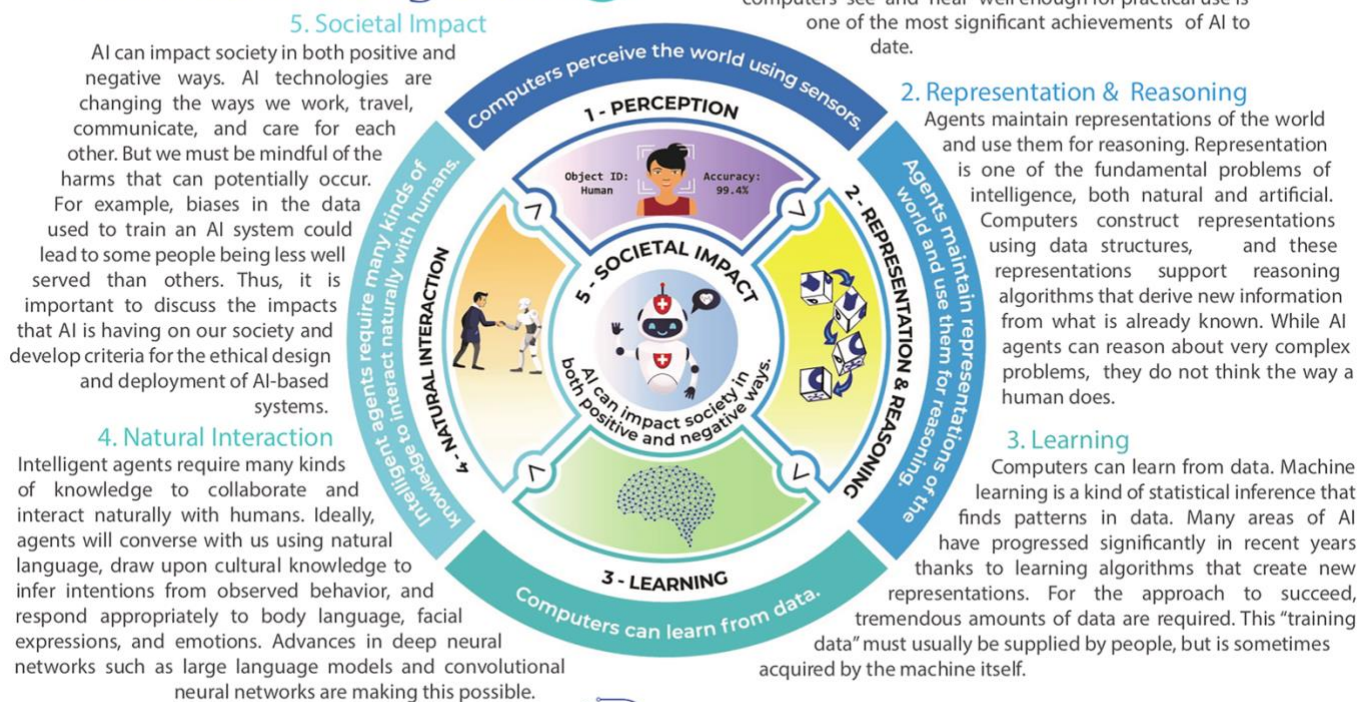
The above definition of AI literacy calls for comprehending a wide variety of topics. Based on our review of recommendations in the field (Druga et al., 2021; MIT STEP Lab, 2024; World Economic Forum, 2024) as well as the content of the [National AI Literacy Day](#), the Department recommends that AI literacy initiatives cover the following topics. These topics are best delivered over time to give educators time to develop understanding and may be delivered in any order.

- **Defining AI.** In its AI Report, the Department suggested conceptualizing modern AI as “automations based on associations” and critically discussed other definitions that rely on analogies to how people think—because machines generate outputs in different ways than people do. Many definitions of AI are now available, including the [Executive Order on AI](#), which defines AI as “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.” The Department recommends looking for AI literacy offerings

and facilitators who can work with educators to critically examine a variety of definitions, and who can work with educators to demystify the fundamentals of AI.

- **AI's History and Origins.** Educators should be aware that AI's history extends back to the earliest days of computers in the 1940s (e.g., the work of Alan Turing), and that humanistic and ethical issues have accompanied progress in the field throughout. AI arose from mathematics and computer science and has moved through a series of approaches that precede today's GenAI approaches. The educational use of AI dates back to 1970, and many prior forms of AI have resulted in useful educational applications (Wooldridge, 2021). The Department recommends looking for AI literacy offerings that provide a sense of history, both for AI and for AI's prior uses in education.
- **Data and Machine Learning.** Modern AI is engineered by identifying patterns in large volumes of data, by a set of techniques broadly called "machine learning." Although educators may not need to understand the differences among machine learning algorithms, they should understand the associative, statistical, and probabilistic nature of machine learning. Doing so will help educators understand why AI sometimes produces erroneous outputs. Without understanding the likelihood of generating erroneous outputs, users may utilize AI-generated information with no review, which can lead to incomplete understanding or false information being learned. Likewise, understanding the process of using data to build AI models can inform educators about the human role in decision making at each stage. Steps such as collecting, cleaning, curating, labeling, analyzing, and evaluating can be presented in ways that build on educators' existing knowledge. Such transparency can support educators to understand how and where biases and inaccuracy can arise.
- **Breadth of AI Capabilities.** Educators should be aware AI is not just chatbots, nor is it confined to textual input and output. Indeed, many experts see "AI" as an umbrella term that encompasses different subfields with distinctive approaches. One way to describe the different aspects of AI is the Five Big Ideas image.

Five Big Ideas in Artificial Intelligence v.2



The AI for K-12 Initiative is a joint project of the Association for the Advancement of Artificial Intelligence (AAAI) and the Computer Science Teachers Association (CSTA), funded by National Science Foundation award DRL-1846073.



This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/4.0/>.



These and other aspects of AI are rapidly being incorporated in a wide range of educational products and services. Educators are likely already familiar with a range of AI applications, such as mapping, robotics, voice assistants, self-driving cars, recommendation systems and more. The Department recommends increasing educator awareness of the range of AI applications and their current use in real life, so they are prepared for the future.

- Existing and Emerging Uses of AI in Education.** AI literacy should build educators’ knowledge of types of algorithms, automations and AI techniques that have been used in education or may be used in the future. One decades-old example is the use of intelligent tutoring systems to support students as they practice mathematics. Algorithms also have been used to identify students who are “at risk” and to score essays for decades. In adaptive learning systems, algorithms determine what content is assigned to students and what feedback they receive. Educators will benefit from coverage of examples with longer histories (like these), examples that are appearing now (such as uses of Chatbots), and future-oriented examples (e.g., where educators or students may use AI to create movies or interactive game-like experiences to pursue learning objectives). It would be beneficial to accompany examples like these with discussion of the types of machine learning or AI used in the application, with some attention both to the capabilities of the technology and its limits or risks.

- **Ethics and Social Impact.** Building on such examples, educators should consider how ethical concepts such as fairness, accountability, and transparency apply to the ways that educators and students select and use AI tools. It is important to acknowledge the potential risks of bias and discrimination that AI technologies may carry, particularly against historically marginalized groups. Understanding these risks and their implications for civil rights is critical for designing equitable learning environments. Students and communities often bring up concerns with the social impact of new technologies, including issues related to bias, fairness, and discrimination, and AI is expected to have broad impacts for the workforce, digital citizenship, and in everyday life. The Department recommends that AI literacy offerings provide ample time to discuss ethics and social impacts, both as they impact educators personally and as they shape the world their students currently live in and will enter in the future.

As an AI literacy initiative progresses, all educators should be able to actively recognize AI in multiple forms, make sense of their own interactions with AI, make informed decisions based on an understanding of how AI is built and its limitations, and view AI through a lens that includes ethics and social impact.

Activity 2: Evaluating AI

Educators will often need to evaluate current or proposed uses of AI, which includes both formal school adoption and informal individual choices by educators and students to use what is available to them for educational purposes. Building on the understanding developed in the prior section, AI literacy should prepare educators to evaluate AI uses in educational settings at multiple levels.

The uses of AI in a district setting are likely to be broad. For example, the Department's AI Report contains chapters on adaptive learning, supporting educators, and improving assessments. AI may also be used in logistics and operations such as lesson planning. An AI literacy curriculum should prepare educators for the potential breadth of use, and not be limited to narrowly evaluating chatbots or any other specific form of AI.

Earlier modules in this toolkit also support educators in evaluating AI. The Department recommends outlining how educators should evaluate AI by considering the following factors previously addressed:

- **Privacy and Data Security.** AI literacy can build on what educators already know about privacy and data security, but also should include attention to new or intensified challenges such as issues related to collecting and using biometric data.
- **Civil Rights and Digital Equity.** This topic encompasses protections from bias and algorithmic discrimination, as well as ways in which AI may better support access and learning for all students. Evaluating AI's capabilities to support diverse learners should also consider issues in the National Ed Tech Plan corresponding to the access, use and design divides.

- **Evidence.** AI literacy should reinforce the importance of looking for and considering evidence as technology is evaluated for educational use. Types of evidence relating to rationale, efficacy and safety are discussed in the earlier module.
- **Alignment to Opportunities and Mitigation of Risks.** AI literacy should emphasize choosing uses of AI that align to an overall educational vision and strategy and avoiding being distracted by misaligned alternatives. Ethics should also be considered during evaluation. [Module 4](#) lists additional risks.

The table below provides some guidance for evaluating edtech. As an exercise, you can also [read the scenarios](#) and consider how you would react in these situations.

LEARNER-CENTERED EMERGING EDTECH EVALUATION RESOURCES	
Is there reliable evidence that the AI-based product or service works?	<ul style="list-style-type: none"> • What Works Clearinghouse • Kentucky Department of Education Evidence-based Practices • The EdTech Evidence Exchange • Ed Tech Adoption
How does the AI-based product promote equity and protect students' civil rights in its design and on-going support?	<ul style="list-style-type: none"> • OCR Reading Room • The CRDC's Data on Equal Access to Education • CASEL Equity Resources • Inclusive Intelligence: Impact of AI • Infrastructure in Inclusive Technology Systems
What safety measures are incorporated into the AI-based product and how do those align with the school's policies and requirements?	<ul style="list-style-type: none"> • U.S. Department of Education's Student Privacy Policy Office • NASSP's Student Data Privacy • The Student Data Privacy Consortium
Do reviews and oversight from trusted parties support the AI-based product's value and trustworthiness?	<ul style="list-style-type: none"> • Common Sense Education • ISTE/EdSurge Product Index • Gartner Education • G2 Education • Capterra Education

Activity 3: Integrating AI into Classrooms

Local educators and educational institutions make decisions about which applications of AI to integrate into classrooms and will design professional development experiences accordingly. This section focuses on five broader considerations in planning how an AI literacy initiative addresses classroom integration.

1. **Balancing immediate implementations vs. time for collaborative planning and support.** The Department's AI Report observed a tension between racing to implement AI in classrooms and taking the time to understand context and create the right conditions for implementation. With regard to getting the context right, there are two issues: (a) uncovering specific factors in the classroom that may contribute to success or failure and (b) creating the right surrounding conditions for success. Likewise, [Module 6](#) strongly encourages educational leaders to establish a sense of pace and priority, so that AI implementation is not rushed. AI literacy should give educators hands-on opportunities to explore AI applications, while also providing time and space for discussions with fellow educators, coaches, and support staff about how to best integrate AI into their specific settings. Educators themselves are in the best position to appreciate how to connect the specifics of their classroom context to new technological approaches but need time to work with each other without the expectation that they will immediately translate an offering into classroom use.
2. **Balancing productivity uses and transformative uses.** Supporting classroom teachers through effective edtech should strike a balance between “doing things better” (productivity) and “doing better things” (transformation). AI could increase productivity on routine tasks (such as writing a first draft of an email to a parent) so that educators can spend more time with students. AI also might more deeply transform educational practice, for example, by revitalizing the process of formative assessment or by generating lesson plans and unique content on demand that more closely reflect their students' backgrounds, yielding increased engagement. Using AI in the most beneficial ways may require pedagogical change, not just pedagogical efficiency.
3. **Establishing student and educator agency.** Students and educators need guidance and access to learning experiences that help them develop and practice AI literacy skills. This practice will help them improve their ability to understand when and how to implement AI tools for school-related activities, including tool selection and evaluation. Without a deep understanding of the three modes of engagement (understand, use, and evaluate), students may rely on AI tools and trust their outputs without critical examination, which could impact their access to information and learning experiences. Additionally, if educators use AI tools to monitor student use of AI tools, students may interpret this as mistrust, which may reduce their willingness to engage in classroom activities.
4. **Developing an understanding of human bias and the strong potential for AI bias.** Educational leaders should provide guidance to educators to help them recognize and acknowledge their existing cultural and societal biases, especially regarding how students from historically marginalized backgrounds are disproportionately disciplined around

biased views about academic integrity, and how using AI tools to detect disingenuous student use can perpetuate and worsen said disciplinary inequities. These supports can come in the form of guidance in responsible use policies, professional learning tools, and enhancing existing equity trainings.

5. **Establishing feedback loops that lead to improved use of AI over time.** In a simple feedback loop, AI literacy should help educators recognize any shortcomings of the initial output they receive (such as a lesson plan) and improve the output before using it. A more systematic feedback loop would be a “plan-do-reflect-improve” cycle (or a similar variant). Educators often meet with each other repeatedly to figure out what’s working and how to make improvements. Educators are also likely to be in feedback loops with educational leadership to discuss how to align opportunities and mitigate risks.

Educational leaders must address current educator needs and plan for successful long-term implementation of AI in classrooms. Thus, educational leaders may ask: **How can I help educators develop the necessary skills and practices to improve the integration of AI into classrooms in ways that meet the specific goals and needs of this learning community?**

Put it into practice! How will your district react in situations that require applying AI literacy skills?

See [Appendix 3](#) for scenarios to review and share with your team.

Additional Considerations for AI Literacy

The three activities discussed above can organize the content of an educational leader’s AI literacy initiative. In the course of reviewing input from the Department’s listening sessions and publicly available documents about AI literacy, the Department also recognizes cross-cutting considerations, which are shared below:




Knowledge Considerations

Demystify AI

Explain how AI works and why it sometimes produces biased, wrong, or risky outputs. Use minimal jargon and find relatable insights to share.

Make Connections	Make connections and build on educators' specific roles (e.g., what subject and grade level they teach) as well as what the audience already knows and does to grow understanding. Connect why, what, and how to support reasoning.
Be Specific About Harms	Build awareness of how AI can cause various harms through an educational system. Examples include: <ul style="list-style-type: none"> • Intensifying cyberbullying • Giving erroneous information • Making decisions that limit students' opportunity • Falsely or unfairly flagging detections of cheating or use of AI content
Human Agency and Risk Management	Highlight human agency and oversight in technical processes, help teachers understand their role in risk management, and when and how to engage individuals in other roles.



System Considerations

Fit AI into an Educational Vision	Prioritize AI applications that align with an educational vision and strategy for AI, as well as specifics in the classroom setting. Be wary of chasing trends and be strategic about implementation.
Build on Prior Initiatives	Build upon prior efforts to develop capacity in digital citizenship, media literacy, data science, and computer science.
Engage Relationships	It is too much to ask teachers to shoulder the burden of responsible AI alone. Design a system that engages the educational institutions' relationships to researchers, policy makers, and vendors, so that all support responsible AI together.



People Considerations

Meet Teachers Where They Are

Teachers may have varying attitudes toward and comfort with AI. A district's AI literacy plan for educators should take into consideration a range of participants, including differences in grade level, subject matter, and student population.

Overcome Unequal Resources

Teachers may have differential access to information resources, as well as time and support for developing AI literacy. Plan to overcome preexisting barriers.

Emphasize a Human Agency

Emphasize that humans are guiding, overseeing, and verifying the outputs generated by AI – that humans are in the loop – and are ultimately accountable for how AI they have chosen to use impacts students. However, keep in mind the limits of what teachers (or any person) can reasonably do, given their existing workload.

Module 9: Updating Policies and Advocating for Responsible Use

As AI becomes more prevalent, educational leaders should revisit their existing policies and any related guidance in place regarding the technological tools and platforms used by teachers, administrators and students and consider how to integrate AI and other emerging technologies. The Department urges educational leaders to prioritize developing and updating clear policies that define expectations about the topics covered in this toolkit and to be advocates for responsible AI use in interactions with the local community and with the developers who supply technology used in their schools.


What Is a Responsible Use of Technology Policy?


A **Responsible Use of Technology Policy (RUP)** provides guidance and supporting information to ensure responsible uses of technology systems and tools in schools. The following guidance was created with entire school ecosystems in mind, including staff, students, parents and caregivers, and learning community members. Most schools and districts already have policies governing use of technology. While often referred to as “acceptable use of technology,” districts are also using the term “responsible use of technology” to emphasize user agency and voice in using emerging technologies ethically (Ruiz et al., 2024). User agency and voice is particularly important with AI because many tools are entering educational settings from everyday use by individuals but can impact the whole community. A RUP gives users choice but holds them accountable to community standards.


Most schools and districts already have acceptable or responsible use of technology policies in place. To account for the rapid changes in AI and other emerging technologies, the Department recommends that district leaders update their policies.

How to Refine and Implement a RUP

Below, we outline key steps educational leaders can take to develop and maintain responsible use guidelines.

<div>Establish a baseline</div> <div></div>	<div><div>1. Review key components that constitute responsible use of AI:</div><div><div>Existing</div>components of the district's RUP.</div><div><div>Current</div>Federal, state, and local privacy and civil rights regulations impacting the development of responsible use plans.</div><div><div>New</div>opportunities and risks specific to AI (as discussed in other modules), including how peer districts have addressed responsible use.</div></div> <div><div>2. Answer questions from the learning community including:</div></div>
--	---

	<ul style="list-style-type: none"> ○ How is AI being used now? What new uses are likely, whether arriving by student choice, through integration into existing platforms, or via newly acquired tools? ○ What appropriate uses of edtech and AI should continue? (Consider perspectives of students, educators, families, administrators, and others.) ○ What are inappropriate uses of AI that should be prohibited? (Consider multiple perspectives.) ○ What are high-risk uses of AI where additional risk management practices should be required? ○ In what ways can students and educators have agency and choice, while avoiding downsides and risks to the broader community? <p>3. Collect a set of scenarios to test the proposed policy. A starter set is included in Appendix 3 regarding cyberbullying, accessibility, and surveillance.</p>
<p><i>Design or update the RUP</i></p> 	<p>4. Incorporate key elements for a successful responsible use plan:</p> <ul style="list-style-type: none"> ○ Context for understanding the policy, including the district's vision, principles, values, and understanding of rights and responsibilities. ○ Definitions and/or expectations around AI literacy, including understanding of how technology use can impact well-being. ○ Specific responsibilities and expectations regarding academic behavior (for both educators and students) regarding, for example, how to cite sources, how to represent one's own work, how to respect copyright and avoid plagiarism, and protecting privacy and security. ○ Allowable or prohibited uses of AI, such as allowing the use of AI to generate initial ideas and prohibiting cyberbullying. ○ High-risk uses of AI, such as the use of AI to make decisions which may impact the rights or safety of students and staff. ○ Consequences for misuse, which specifies how the educational institution will monitor use and respond to misuse. ○ Recognition that students and staff with disabilities may have unique needs and might require something in addition to what the RUP provides, and that their needs will be considered consistent with the IDEA, ADA, and Section 504. <p>5. Include additional guiding language on key topics of responsible AI such as:</p>

	<ul style="list-style-type: none"> ○ Safety, including data privacy and security ○ Accountability ○ Transparency ○ Fairness (and avoiding bias) ○ Efficacy <p>6. Involve multiple roles in defining and reviewing the policy, and clearly identify how students, educators, administrators, families, and broader community members participate in enacting and supporting the policy.</p>
<p><i>Implement the RUP and periodically improve it</i></p> 	<p>7. Disseminate the policy, train community members in its use, and actively follow it.</p> <p>8. Create feedback loops with educators and students (and the task force recommended in Module 7) to address new or challenging situations as they arise.</p> <p>9. Establish the cadence in which the learning community will review and update the policies based on use, technological updates, and feedback.</p> <p>10. Monitor implementation and follow a process for reporting, investigating, and resolving violations of the policy. Collect data on student use and any discipline that may be related to AI to ensure that the use of AI does not result in disproportionate discipline of certain student groups.</p>

Advocating for Responsible Use with Technology Partners

RUPs provide guidance for educators and students to select the appropriate tools for educational tasks, and to use them ethically. RUPs can also be the basis for important discussions with technology partners—the vendors or suppliers of the technologies in use in the district. Technology partners can make it easier for students, educators, and other community members to engage in acceptable use (for example, by making it easier for students to disclose how they used AI in completing an assignment) and prevent prohibited uses.

The Department has produced a guide called [Designing for Education with Artificial Intelligence: An Essential Guide for Developers](#) that complements the guidance here. Both documents have foundations in the topics covered in [Modules 1-4](#) of this toolkit, for example. Organizations such as the Software Information and Industry Association have released [Principles for the Future of AI in Education](#), and EdSAFE AI released its [SAFE benchmarks](#). These documents can help educational leaders advocate for responsible AI development and deployment as well as responsible use.

Indeed, trust across educational institutions and technology partners benefits all:

- Both educational institutions and technology partners want meaningful and lasting use of AI to improve education, and they can work together to prioritize a clear vision and avoid distraction by the latest (but unrelated) advances.
- Both educational institutions and technology partners have roles in establishing the safety of AI for educational use, and neither can do all the work on their own.
- Both educational institutions and technology partners have an incentive to address access, use, and design divides so that the full diversity of student populations can benefit.
- Both educational institutions and technology partners strengthen their work by clear communication and feedback loops that can lead to further mitigating risks and making product improvements.

This toolkit recommends that educational leaders use their RUP and other strategic AI documents to advocate with their technology partners for improvements that support the district's vision and strategy for using AI.

Put it into practice! How will your district apply what you have learned in situations involving accessibility, cyberbullying, and surveillance?

See [Appendix 3](#) for scenarios to review and share with your team.

Module 10: Ambitious Action Plans

In this concluding module, the Department calls upon educational leaders to build an organization-wide action plan for implementing AI in their educational community. Educational leaders will be continuously called upon to vet new ways of using AI and to manage anticipated and emergent risks. Doing so responsibly will require strengthening capacities across many roles and responsibilities.

In building a plan, educational leaders should anticipate three ways that AI will enter their community:

1. **Procurement:** Some uses of AI will be formally procured and implemented, as has been the case for edtech for decades.
2. **Integration:** AI will be integrated into existing educational products and services that are already in use.
3. **Diffusion:** AI will spread widely in society and enter schools in the hands of students and educators who bring everyday tools into learning settings.

An organizational plan should address a varied and changing AI landscape and include ongoing reviews to account for AI integration and diffusion. Because most educational institutions have established ways of working with vendors, AI may enter via formal procurement as a separate product less frequently. Based on experience, the Department anticipates that AI will enter schools more frequently through diffusion and integration. This means that teachers, students, and parents will likely use AI-enabled tools that circumvent the traditional approval processes. Increasing awareness about the opportunities and challenges of AI-enabled tools will require the evaluation of these tools by not just IT and emerging technology experts, but also by the broader school community.

Why Build an Organization-Wide Plan for AI?

As discussed throughout this toolkit, the growing presence of AI in education requires addressing longstanding concerns such as [data security and privacy](#) as well as newer technical issues such as [algorithmic discrimination](#). AI-specific challenges also include managing new problematic student behaviors such as cyberbullying and academic dishonesty. A look at the history leads to a clear lesson: Educational leaders will need to have a comprehensive plan for responsible AI integration, and the time to start building that plan is now.

What Should be Included in an Organization-Wide Plan?

In its [Executive Order on AI](#), the administration wrote that “Harnessing AI for good and realizing its myriad benefits requires mitigating its substantial risks. This endeavor demands a society-wide effort that includes government, the private sector, academia, and civil society.” In an educational organization such as a school district, a responsible AI plan will convey how the district intends to coordinate its people, resources, and activities to realize strategic educational benefits while mitigating the most consequential risks. While the scope of

responsible AI plans is still evolving, some educational organizations have begun to create planning tools. Some examples include:

- The GenAI Readiness [checklist](#), authored by the Council of Great City Schools, CoSN and Amazon Web Services, covers readiness in the following areas: Executive Leadership, Operational, Data, Technical, Security, and Legal/Risk Management.
- Michigan’s Virtual Learning and Research Institute produced an [AI Integration Framework](#). This includes many of the categories listed for the checklist (above) and adds areas such as Curriculum and Instruction, Assessment, and Community Outreach.
- TeachAI’s [Policy Ideas](#) cover five areas – Leadership, Governance, Capacity, AI Literacy, and Supporting Innovation.
- The EDSAFE AI [Framework](#) emphasizes working cohesively to advance Safety, Accountability, Fairness and Transparency, and Efficacy.

Individual educators, students, and families often look to school and district leaders to ascertain which AI tools and resources are safe and appropriate to use, and many districts already maintain a list, often on a website, of technologies they have reviewed for use in classroom instruction or other educational activities. Revising and updating these lists is an important starting point for educational leaders. As emphasized throughout this toolkit, educational leaders should evaluate existing resources and build from there.

Beyond initial evaluation, tools like the Michigan Virtual AI Integration Framework recommend a development progression: investigating, implementing, and innovating. This toolkit is aligned with an early developmental stage, in the form of questions for educational leaders who are guiding their communities through an exploratory phase.

Below, we suggest potential steps for an iterative developmental process. To build a comprehensive roadmap for their organization, educational leaders should review all the modules in this toolkit as well as additional resources from organizations they work with and trust.

1. **How does AI impact data privacy and security?** In an initial phase, educational leaders should review their existing data privacy and security policies, as well as guidance and sample policies from their states and from nonprofit organizations before developing new policies. We note three key directions for development:
 - a. Guiding educators, students, and community members about the privacy and data security risks that arise in technologies that enter the community, especially through diffusion.
 - b. Addressing the increasing breadth of private data that may be collected as edtech becomes more multimodal.
 - c. Clarifying the increasing role of biometrics in access and monitoring in educational settings.
2. **How could AI enhance or undermine digital equity and civil rights?** In an initial phase, educational leaders should map the access, use, and design issues (see discussion in

[Module 2](#)) already in play in their community. Important topics for development include:

- a. Investigating how algorithms in AI and emerging technologies can lead to discrimination as platforms increasingly use them to suggest resources for students and make recommendations about student placements and evaluations to educators, raising concerns about biased surveillance and inequitable outcomes.
 - b. Understanding that teachers are becoming heavily reliant on school-sanctioned AI content detection tools, and there is an increase in student discipline as a result. These AI detection tools have also been found to be ineffective.
 - c. Considering how the school will monitor AI usage for discrimination, and how it can be reported, investigated, and remedied when it happens.
 - d. Pursuing greater inclusion and use of AI for students with disabilities, multilingual learners, and for other student strengths and needs, especially as AI-enabled technologies could provide students with more ways to effectively engage in learning environments.
 - e. Addressing differential access to AI by students and their communities, especially in light of pricing models that may work for some families and not others.
3. **What evidence supports the effective use of AI in education?** In an initial phase, educational leaders should review the available evidence on AI to support their vision and approaches. Three key directions for growth include:
 - a. Developing their community's healthy skepticism of industry-based rationales for why and how to use AI in education, which may have emotional appeal but little basis in research.
 - b. Strengthening the degree to which evaluation of the rationale for a proposed AI use is based on research, or similarly, on trustworthy summaries of research such as in the What Works Clearinghouse [Practice Guides](#).
 - c. Continuing to require high-quality evidence from vendors for the efficacy of their products (U.S. Department of Education, 2024), such as independent research that would meet the Department's standards for Strong or Moderate evidence, if available, including whether any integrated AI features add value and have addressed risks (such as bias). Districts can require evidence from vendors that includes:
 - i. Third party certifications
 - ii. The creation of a data collection and analysis plan to check outcomes status before renewing a contract
 - iii. A logic model and annotated bibliography to understand the research basis
4. **Which opportunities and risks should an educational leadership team prioritize?** Educational leaders should focus their community on opportunities that matter most to their educational vision and approach, and avoid distraction by alluring but misaligned and untested possibilities. Risks are addressed in [Module 4](#) and below. Educational leaders should understand that risks may appear in technologies that enter through any

of the routes (procurement, integration, diffusion), and thus risk management should be comprehensive rather than just a component of formal adoption of technology.

How can educational leaders define an organization-wide plan?

Few educational institutions become “ready” for AI overnight, and most will strengthen their plans iteratively over years. Innovative educational leaders guide their communities through developmental growth in capabilities and structures over time. The Department suggests using the structure of the NIST AI Risk Management Framework (see [Module 4](#)) to guide this ongoing process, but with a modification to transform into both an opportunity vetting and a risk management framework. We suggest expanding the four elements of the NIST AI Risk Management framework as follows:

- **Govern:**
 - Develop a culture that assesses opportunities that align with the core of the district’s vision and strategy and avoids distraction.
 - Promote a culture where people in all roles in the district have training about risks and contribute to managing risks.
- **Map:**
 - Recognize how opportunities are specific to the varied contexts in the district, such as grade levels, subject areas, and other strengths and needs of groups of students and district community members.
 - Identify and document use, access, and design divides that have impacted equal opportunity to learn in the past, and how use of AI-enabled technologies might address these divides.
 - Anticipate both existing risks (e.g., that an ‘adaptive’ product necessarily collects information about each student) and a broader set of potential risks for each type or category of use of AI, including risks that go beyond privacy and data security (e.g., bias and algorithmic discrimination).
- **Measure:**
 - Establish relationships with developers of AI-enabled tools that specify and measure both the evidentiary strength of an opportunity and the quality of risk mitigation, starting by advocating for more transparency and accountability from vendors and suppliers.
 - Continue to strengthen leadership within the district with regard to gathering and using evidence (and research) to measure efficacy, as well as for whom and under what conditions an AI-enabled system mitigates risks, so that opportunities can be realized.
- **Manage:**
 - Strengthen capacity to manage implementation of new uses of technology so that benefits to students, educators, and the community are realized equitably.
 - Increase opportunities for professional learning for educators, including their capacity to manage both anticipated and emergent risks, across AI-enabled technologies that arrive in the hands of students and educators.

We give an example below of how an educational leader might use these four elements in a first round of iterative development, understanding that educational leaders will adapt the elements and goals to meet the needs of their districts.

In this scenario, an educational leader might decide to pursue an initial vision such as “every teacher will be better able to specify uses of AI that are supportive of their academic goals, recognize risks of using AI in their academic context, and know how to work with our task force to ensure student safety throughout the school year.” With this vision in mind, the educational leader might propose key actions in their first-year plan, such as:

- **Govern:** Our leadership team will establish an AI task force on the academic use of AI (see [Module 7](#)) that will be charged with supporting classroom educators in making decisions about how to protect student safety as they explore responsible use of AI, and with reporting to leadership on both the opportunities and challenges that emerge.
- **Map:** Our IT leadership will identify the 12 most common AI tools that educators and students are using for academic content and how uses vary in our district (for example, in different grade levels, in different schools, or for different content). Given these uses and our context, they will also start to identify important risks and communicate about the risks to district leaders, building leaders, and educators.
- **Measure:** Our climate team will upgrade their survey to measure how GenAI uses may be positively or adversely impacting school climate.
- **Measure:** Our IT team will engage vendors of our most used platforms or systems to understand whether they are incorporating AI, in what ways, and what steps they are taking to ensure transparency and accountability.
- **Manage:** We will deliver AI literacy training to all educators in our district, and the training will prepare them to have discussions with students about AI use on academic tasks. The training will support educators to achieve our initial vision and increase their awareness of risks.
- **Manage:** Our task force will respond to risks that arise as educators pursue academic uses of AI, and will draw district leadership’s attention to important, unresolved risks.

Put it into practice! How will your district apply what you have learned about informed adoption?

See [Appendix 3](#) for scenarios to review and share with your team.

Appendix 1: Key Terms and Acronyms

The more granular our knowledge base is on any topic, the richer our collaboration and the tighter our bonds are through shared experiences. Mountaineering and climbing have hundreds of terms for gear, technique, and conditions as does the specific computer science subfield of artificial intelligence. Communication is more effective when all parties have a baseline of such terms, and the following list is intended to help build AI literacy among your teammates.

- **Artificial Intelligence (AI):** The term “artificial intelligence” or “AI” has the meaning set forth in 15 U.S.C. 9401(3): a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments, abstract such perceptions into models through analysis in an automated manner, and use model inference to formulate options for information or action.
- **Generative Artificial Intelligence (GenAI):** The term “GenAI” means the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content.
- **Educational Technology (EdTech):** Any technologies designed for and employed by educators that can be used to transform school operations, but more so, teaching and learning for the sake of engagement, post-graduation readiness and success, and to make everywhere, all-the-time learning possible for early learners through PK-12, higher education, and adult education.
- **Large Language Models (LLM):** Large language models form the foundation for GenAI systems and are artificial neural networks. At a very basic level, the LLM detected statistical relationships between how likely a word is to appear following the previous word in their training. As they answer questions or write text, LLMs use the model of the likelihood of a word occurring to predict the next word to generate. LLMs are a type of foundation model, which are pre-trained with deep learning techniques on massive data sets of text documents. Sometimes, companies include data sets of text without the creator’s consent.
- **User Interface/User Experience (UI/UX) Design:** The human-centered design processes for creating interfaces for both explicitly digital and physical, analog, or other non-digital applications to yield solutions that are both effective and aesthetically pleasing to the end user.
- **Guardrails:** Provide specific boundaries for the use of AI systems and tools. These include specific examples and use cases and are different from policies which provide a framework for the use of AI.

- **Personally identifiable information (PII):** Personally identifiable information (PII) refers to direct identifiers, such as a student’s name or address; indirect identifiers, such as a student’s date of birth; and any information that, alone or in combination with other information, is linked or linkable to a specific student that would allow a reasonable person in the school community, who does not have personal knowledge of the relevant circumstances, to identify the student with reasonable certainty.
- **Algorithmic Discrimination:** “Algorithmic discrimination” occurs when automated systems contribute to unjustified different treatment or impacts disfavoring people based on their race, color, ethnicity, sex (including pregnancy, childbirth, and related medical conditions, gender identity, intersex status, and sexual orientation), religion, age, national origin, disability, veteran status, genetic information, or any other classification protected by law. Depending on the specific circumstances, such algorithmic discrimination may violate legal protections. Throughout this framework the term “algorithmic discrimination” takes this meaning (and not a technical understanding of discrimination as distinguishing between items).
- **Sensitive Data:** Data and metadata are sensitive if they pertain to an individual in a sensitive domain (defined below); are generated by technologies used in a sensitive domain; can be used to infer data from a sensitive domain or sensitive data about an individual (such as disability-related data, genomic data, biometric data, behavioral data, geolocation data, data related to interaction with the criminal justice system, relationship history and legal status such as custody and divorce information, and home, work, or school environmental data); or have the reasonable potential to be used in ways that are likely to expose individuals to meaningful harm, such as a loss of privacy or financial harm due to identity theft. Data and metadata generated by or about those who are not yet legal adults is also sensitive, even if not related to a sensitive domain. Such data includes, but is not limited to, numerical, text, image, audio, or video data.

Appendix 2: Building on Foundational Initiatives

Proficient hikers who learn to climb open up exponentially more opportunities for exploration around the world. The same is true for the cross-cutting themes in education we are bringing with us into this new era of AI-enabled edtech.

- **Build on foundational skills from existing initiatives.** Districts and schools likely have systems in place to learn capacity building on a variety of literacies:
 - Web and Media literacy, Mozilla Foundation (<https://foundation.mozilla.org/en/initiatives/web-literacy/>)
 - Digital Citizenship, Common Sense Media (<https://www.commonsense.org/> and <https://iste.org/digital-citizenship>)
 - Computational thinking (<https://iste.org/computational-thinking> and <https://digitalpromise.org/initiative/computational-thinking/>)
 - STEM and STEAM initiatives (<https://iste.org/steam-education>)
 - 21st Century Skills (<https://curriculumredesign.org/frameworks/framework-1-0/>)
 - Habits of Mind (<https://www.habitsofmindinstitute.org/what-are-habits-of-mind/>)
 - Portrait of a Graduate (<https://www.battelleforkids.org/> and <https://digitalpromise.org/2023/08/21/six-attributes-for-portrait-of-a-powerful-graduate/>)

Appendix 3: Reacting and Shaping Use Scenarios

The following real-life examples include possible topics that may arise as you consider how to use AI in your school or district.

Each scenario is accompanied by an example of what an educational leader may do in response to the situation. Educational leaders should consider the guidelines from this toolkit in the context of their specific community and audiences when determining how to address a scenario.



Accessibility

Because digitally enhanced learning design has prioritized accessibility for decades, navigating technology for individuals with disabilities is more manageable than ever before. But GenAI's influence upon digital accessibility is multivariate. Emerging technology may provide focused assistance to assist students with disabilities, such as by employing natural language processing and computer vision. But in other instances, technology may not be designed to be accessible to students who learn differently or who have an IEP or 504 plan. Though this work is still emerging and does come with inherent risks, administrators can expect more progress and innovation related to digital accessibility in the future.

Digital Tools Empowering Real-World Accessibility

Your director of special education forwards you an email from an edtech solution provider that claims to support students with vision impairments in a new way. The solution employs computer vision via the user's smartphone camera to help with navigating physical environments as well as interpreting written signage and other materials, giving students more freedom and agency during the school day. As an administrator, you recognize the potential accessibility value but also the human resource value: Such a tool could ostensibly free up staff to assist students in other ways.

What are your next steps?

An Educational Leader might: Ask the provider about: (1) what they will use the recorded data for ([Module 10](#)) and (2) how will they ensure and protect student (and family) privacy ([Module 2](#)). District leaders should also engage the community, particularly paraprofessionals and caregivers of students who participate in special education programs, to understand their concerns, recommendations, and needs. This can be done through a listening session ([Module 5](#)) or by developing a task force ([Module 7](#)).

Assessment Accommodations via Emerging Technologies

State testing is approaching, and a popular vendor already trusted with prior solutions has a new suite of GenAI-enabled tools that the vendor claims are aligned with the ESEA/ESSA provisions and provides assessment accommodations tailored to individual students' IEP requirements. With a high percentage of students who have IEPs, this seems like a valid opportunity to ensure students with disabilities can access the assessment.

What are your next steps?

An Educational Leader might: Ask the vendor to provide specific ESEA-aligned documentation of the tools' basis in research and evidence ([Module 4](#)) and to provide more details about what they plan to use the recorded data for ([Module 10](#)). The leader could also collaborate with special education educators to understand how the tools could be helpful to students with IEPs, while considering what risks should be mitigated for the populations they serve ([Module 3](#)).



Cyberbullying

Cyberbullying is an ongoing, challenging issue which has inflicted extreme harm upon students, even resulting in death. As emerging technologies with autonomous agents surface in the GenAI era, cyberbullying may become more volatile. It is the responsibility of educational leaders to maintain awareness and reach out to their community, including task force members, to understand how to address these issues as they come up.

Deep Fake Images from GenAI

A teacher brings to your attention that a student has been targeted by a cyberbullying effort utilizing deep fake images from a GenAI tool and that her parents are on their way to pick her up. You do not know where the images came from, who made them, or how they were shared. Before the first student has made it down to the office, another student is targeted in the same way. And all that is known is that the images are being circulated by other students on their personal devices.

What are your next steps?

An Educational Leader might: Prioritize addressing harms to the students to assure their safety and well-being. Then, proactive communication with students, families, and the greater district community will be important, for example, holding listening sessions with impacted students and families ([Module 5](#)). They should work with their IT team to investigate the source of the attack and consider enhancements to the school's data privacy and security policies ([Module 2](#)). They can also train and support the school community about responsible digital behavior, including steps to prevent, report, and address cyberbullying and related attacks ([Module 3](#) and [Module 8](#)).

Threatening Messages from GenAI

Your high school's administrative assistant comes to school in the morning and, while reviewing the general inbox for parent emails and regular business, finds a threatening email with explicit details about a threat to cause harm to students at the school. It isn't obvious who wrote and sent the threat because they apparently used a chatbot to write the language and multiple measures to cover their IP address and physical location. However, the administrative team is unsure about the extent of safety concerns.

What are your next steps?

An Educational Leader might: Prioritize student safety, privacy, and civil rights ([Module 2](#)) by taking immediate action to address risks while ensuring transparent communication with the community ([Module 10](#)). Leaders should collaborate with authorities, according to their emergency response plan and in accordance with applicable law, to fully evaluate the threat and keep families informed. They should explore if the email is covered by FERPA, and, if so, if PII from the email will be disclosed, with consideration to FERPA's general consent requirement and exceptions. If safety concerns persist and cannot be fully mitigated, erring on the side of caution may be necessary while respecting students' right to a safe learning environment. Additional information on responding to these types of incidents can be found in the recent [CDT report](#).

Impersonation Campaigns

A group of students receive convincing, but inaccurate, progress reports with personalized commentary allegedly from their school counselor saying they are at risk for failing in certain classes, causing a flurry of visits to the counseling center and inbound calls from parents who were also copied on the emails. In another attack, several students are impersonated by another student using chatbots and other means to create convincing social media posts directed at fellow classmates to obscure the source, convincing the victims that these posts came from real classmates.

What are your next steps?

An Educational Leader might: Assess the risks of these attacks to student well-being and privacy ([Module 3](#)) while leveraging evidence-based cybersecurity measures. A leader may also view this as an opportunity to enhance AI literacy programming ([Module 8](#)), educating the school community on recognizing AI-driven misinformation and fostering a culture of responsible tech use. Listening sessions ([Module 5](#)) could be another way to empower community members to understand and work to address these issues.



Informed Adoption

Some educators have access to the latest technology, resources, and support for sophisticated digital learning experiences. Others get by with the most affordable and readily available edtech solutions. Some even cobble together their own solutions for data tracking via online spreadsheets and forms, which may not be reliable or secure.

The GenAI paradigm shift may cause further separation between such educational leaders' approaches, but it may help others catch up by supplanting knowledge gaps with increasingly capable intelligent assistants. It is incumbent upon administrators to improve their adoption approaches as solution providers update their product roadmaps and marketing campaigns. As with innovations in adventure equipment, astute educational leaders will be able to discern which claims will help them accomplish their vision and mission and which are just smoke and mirrors.

The “Wow” Factor

Using edtech adoption resources available from trusted organizations, your edtech evaluation team has sifted through some of the latest solutions that fit your school’s needs and decided on a clear frontrunner. This tool has GenAI components incorporated into its core functionality. The vendor’s demonstrations are truly impressive in efficiency and reliability. However, the vendor cannot answer some of your questions about transparency for intellectual property or “trade secret” reasoning.

What are your next steps?

An Educational Leader might: Continue to advocate for clarity from the vendor ([Module 10](#)), which is critical for making an informed purchasing decision. If the vendor cannot provide satisfactory answers regarding intellectual property, that may indicate potential risks that outweigh the tool’s benefits. With many edtech solutions available, leaders can feel empowered to explore alternative options that align with both educational goals and transparency standards ([Module 10](#)).

The New Kid on the Block

You’re walking the exhibit hall at a national edtech conference and are drawn into a booth for a new product that clearly uses AI based upon their signage. A rep shows you a few features and mentions machine learning and natural language processing which draws you further into the demonstration because your team has been complaining about not wanting to lag behind in innovation. However, the team on the floor cannot answer simple questions about the underlying AI model or the training data that were used, and thus cannot refute any potential ethical biases that the system could introduce into your school or district.

What are your next steps?

An Educational Leader might: Remain cautious when edtech vendors cannot answer fundamental questions about their AI models or the training data used, especially given the risk of introducing ethical biases ([Module 3](#)). Leaders should prioritize products that align with their school’s ethical standards ([Module 3](#)) and community priorities ([Module 5](#)) and feel confident exploring other options if a vendor cannot provide requested information ([Module 10](#)).

“...Now with GenAI!”

Your colleagues are excited about advertised GenAI updates coming to the student information system you’ve been using for some years, because it will allegedly make certain tasks, such as entering grades and preparing end-of-term reports, more efficient and accurate. However, the vendor is offering this updated package at a premium price, increasing the per-student rate by several dollars.

What are your next steps?

An Educational Leader might: Ask for a demonstration and work with their district/school community to carefully evaluate whether the efficiency gains of the tool justify the increased cost. The team should also ask the vendor for clear evidence of benefits and what safeguards are in place to prevent bias ([Module 1](#), [Module 3](#), and [Module 10](#)).

Cost-Effectiveness

You are a seasoned administrator, leading a city-based LEA that is relatively underfunded compared to neighboring suburban districts. Staff turnover is above the state average and your team is considering a new counseling solution, which employs a popular LLM with a custom chatbot interface, that is reportedly attuned to the needs of students needing school counseling across a variety of domains, both academic and social-emotional. Early reports and endorsements praise the developer for its innovation. The tool in question is considerably cheaper than hiring even one more counselor.

What are your next steps?

An Educational Leader might: Evaluate the LLM-based chatbot solution by assessing its civil rights implications, ensuring equitable access and privacy protections, while managing risks such as bias and data security ([Module 3](#), [Module 10](#)). Gather feedback from the school community, for example through listening sessions ([Module 5](#)), and prioritize building AI literacy among staff to ensure the tool's responsible and effective use ([Module 8](#)). Above all, the leader and team would keep in mind that while such tools may be able to enhance or supplement the work of human counselors, risks and biases need to be addressed and AI tools should never replace human team members.

It's About Assessment

Educators are concerned about students using GenAI to cheat on formative assessment formats, including short answer questions and essays. Despite efforts to prevent student cheating, there is an ongoing debate about how to fully verify that a student's work is truly their own. On the other hand, educators well versed in modern instructional methods such as project-based and real-world learning are seeing the emergence of GenAI in edtech as an opportunity to redesign common assessment strategies across content areas and grade levels. Perhaps their standards-aligned approaches, with rubrics for measuring student performance and growth, could prove exemplary for updating assessment practices for teachers who have been employing more traditional approaches.

What are your next steps?

An Educational Leader might: Ensure that assessments using GenAI include human safeguards to prevent biased evaluations, particularly for students from historically marginalized groups ([Module 3](#)). Encouraging more authentic, qualitative assessments can promote equitable and meaningful measures of student growth. They can provide optional AI literacy workshops for educators and families to support fair and informed assessment practices ([Module 8](#)).

Hello, Neighbor!

You hear about a neighboring district that has taken an aggressive approach to exploring and implementing GenAI in its learning experience design. So have your colleagues, as they regularly talk about the innovative opportunities on the other side of town. Yet the culture of your learning community has historically been more conservative and risk averse.

What are your next steps?

An Educational Leader might: Be transparent about their slow and purposeful approach to integrating GenAI, prioritizing solid research and evidence ([Module 4](#)) and aligning with the unique needs and values of their community. By engaging the students, families, and educators, they can listen to community concerns and explore the potential of GenAI while fostering trust alongside innovation ([Modules 5-7](#)).



Monitoring and Surveillance

Sophisticated solutions have existed for monitoring and surveillance in our schools for years, and AI is already capable of making inferences based on audio and visual data to highlight and describe events involving people, accelerating privacy risks. While educational leaders hope to use such early warning systems to support educators and students, this segment of edtech has potential for bias and misuse.

Students' civil rights, privacy, and well-being are potentially vulnerable, so students need to be able to trust the policies their educators and administrators implement and enforce surrounding online behaviors at school. An educational leader engaging students based on the automated analysis of digital artifacts needs to be weighed against documented evidence of bias in AI training data sets that disproportionately single out students from historically excluded cohorts based on race, ethnicity or nationality, or socioeconomic status.

Educators must always consider the margins for error with AI in monitoring and surveillance scenarios.

Keeping an Eye on Them

One of your assistant principals alerts you that your content moderation solution has reported some suspicious online behavior from “a usual suspect” and that she will be keeping an eye on this particular student. Though well intended, you have also heard about this issue at an education conference and are having second thoughts about that software’s value based on ethics alone.

What are your next steps?

An Educational Leader might: Work with a team including the student’s teachers to better understand the context of the accusation ([Module 7](#)). Additionally, the educational leader can reach out to the content moderation solution provider to better understand their system, including how it flags and what it counts as suspicious activity, as well as how they recommend educational leaders interpret the reports ([Module 2](#) and [Module 3](#)). The educational leader can then speak to the student and their family to hear their perspective and use this information to develop guidance and guidelines for how the content moderation solution can be leveraged in the broader student support and well-being ecosystem ([Module 9](#)).

Digital Profiling

Your Chief Technology Officer has been in conversation with an edtech sales representative with an innovative solution that uses AI to interpret the physical presence data from multiple sources including security cameras, Wi-Fi access points, and student ID badge terminals. The solution can accurately track a student’s movements and recognize patterns aligned with staff documentation and indicate which students are most likely to misbehave.

What are your next steps?

An Educational Leader might: First, educational leaders must check the laws in their state and city to better understand the legality of this type of surveillance within a school setting ([Module 2](#)). They must also consider student privacy and data rights, and whether or not they have permission to use these systems for these purposes ([Module 2](#)). Importantly, educational leaders should work to build a shared understanding of how this use of emerging technology works within their context, bringing in community members including students ([Module 8](#)) and the technology task force ([Module 7](#)).

De-anonymizing Anonymous Posts

A frustrated high school student anonymously posts online a verbose criticism of school policy or even your leadership for being unfair in some manner. While pointed, the manifesto isn’t particularly dangerous or threatening. The teenager is using a social media account to share their heated perspective that may have resulted from stress, an isolated incident, or another trigger. However, a recent monitoring tool that has been updated to employ an LLM “de-anonymizes” the student’s writing based on other posted language and shares that student’s name with you and their advisory teacher.

What are your next steps?

An Educational Leader might: An important first step is to help the district community learn about and better understand the increasing availability of tools that allow for triangulation of data and what that means for them in various school and personal settings ([Module 8](#)). Many students and community members do not realize that the large amount of data they make available to companies can be connected easily to identify them even when they are under the impression that they are acting anonymously. This situation also requires the educational leader to work with the appropriate school team to address this situation with the student and their family directly ([Module 7](#)).

References

- The Association for Supervision and Curriculum Development, The International Society for Technology in Education, The National Association of Secondary School Principals, The National Association of Elementary School Principals; The American Association of School Administrators. (2023, July). Bringing AI to school: Tips for school leaders. https://cms-live-media.iste.org/Bringing_AI_to_School-2023_07.pdf
- California Department of Education. (2024, June 12). Learning with AI, Learning about AI. <https://www.cde.ca.gov/pd/ca/cs/aiincalifornia.asp>
- CAST (2024). Universal Design for Learning Guidelines version 3.0. Retrieved from <https://udlguidelines.cast.org>
- Center for Democracy and Technology. (2023, December 18). *The shortcomings of generative AI detection: How schools should approach declining teacher trust in students*. Center for Democracy and Technology. <https://cdt.org/insights/the-shortcomings-of-generative-ai-detection-how-schools-should-approach-declining-teacher-trust-in-students/>
- Chine, D.R. et al. (2022). Educational Equity Through Combined Human-AI Personalization: A Propensity Matching Evaluation. *Rodrigo, M.M., Matsuda, N., Cristea, A.I., Dimitrova, V. (eds) Artificial Intelligence in Education*. AIED 2022. Lecture Notes in Computer Science, vol 13355. Springer, Cham. https://doi.org/10.1007/978-3-031-11644-5_30
- Coburn, C. E., & Penuel, W. R. (2016). Research–practice partnerships in education: Outcomes, dynamics, and open questions. *Educational researcher*, 45(1), 48-54.
- Cohen, D. K., & Ball, D. L. (1999). *Instruction, capacity, and improvement*. CPRE Research Reports. Retrieved from <https://repository.upenn.edu/handle/20.500.14332/8434>
- Consortium for School Networking, Council of the Great City Schools. (2024b, April). CoSN/CGCS K-12 Gen AI Maturity Tool. https://www.cosn.org/wp-content/uploads/2024/04/Cosn_CGCS-K-12-Gen-AI-Maturity-Tool-Conference.pdf
- Consortium for School Networking. (2023, October). Readiness Checklist. <https://www.cgcs.org/genaichcklist>
- Digital Promise. (2024, June 19). District Resources. *Product Certifications*. <https://productcertifications.digitalpromise.org/resources/district-resources/>
- Digital Promise. (2024, March 22). A process to help run successful Ed-Tech pilots. *Edtech Pilot Framework*. <https://edtech.digitalpromise.org/>
- Druga, S., Yip, J., Preston, M., & Dillon, D. (2021). The 4As: Ask, adapt, author, analyze - AI literacy framework for families. In *Algorithmic Rights and Protections for Children*. <https://doi.org/10.1162/ba67f642.646d0673>

- The Edtech Equity Project. (2020, June). *School Procurement Guide: Buying Edtech Products with Racial Equity in Mind*. The Edtech Equity Project.
<https://www.edtechequity.org/work/school-procurement-guide>
- TNTP. (2018). Choosing the Opportunity Gap. *The Opportunity Myth*.
<https://opportunitymyth.tntp.org/choosing-the-opportunity-gap>
- Elmore, R. (2008). *Improving the Instructional Core*. Harvard University, School of Education.
Retrieved from
https://www.pps.net/site/handlers/filedownload.ashx?moduleinstanceid=19185&dataid=123381&FileName=Improving_the_Instructional_Core.pdf
- EngageAI Bill of Rights Task Force. (2024, May). An AI Bill of Rights for Educators. EngageAI Institute.
- Farrell, C. C., Penuel, W. R., Allen, A., Anderson, E. R., Bohannon, A. X., Coburn, C. E., & Brown, S. L. (2022). Learning at the boundaries of research and practice: A framework for understanding research–practice partnerships. *Educational Researcher*, 51(3), 197-208.
- Gao, G., Yang, X., & Chi, M. (2024). Get a head start: On-demand pedagogical policy selection in intelligent tutoring. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(11), 12136-12144. <https://doi.org/10.1609/aaai.v38i11.29102>
- Kentucky Department of Education. (2010). *Highly Effective Teaching and Learning in Kentucky Overview Guide*.
https://www.education.ky.gov/curriculum/standards/teachtools/Documents/Highly_Effective_Teaching_and_Learning_Overview_Guide.pdf
- Kentucky Department of Education. (2024). AI Guidance Brief.
<https://www.education.ky.gov/districts/tech/Documents/AI%20Guidance%20Brief.pdf>
- Kulesa, A. C., Croft, M., Brian, R. and Wells, M. K. (2024). Learning Systems: Opportunities and Challenges of Artificial Intelligence-Enhanced Education. Bellwether. Retrieved October 14, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_2_Bellwether_September2024.pdf
- Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), 42-78.
<https://doi.org/10.3102/0034654315581420>
- Laird, E., Dwyer, M., & Woelfel, K. (2024, September 26). In Deep Trouble: Surfacing Tech-Powered Sexual Harassment in K-12 Schools. Center for Democracy & Technology.
<https://cdt.org/wp-content/uploads/2024/09/2024-09-26-final-Civic-Tech-Fall-Polling-research-1.pdf>

- Mills, K., Ruiz, P., Lee, K., Coenraad, M., Fusco, J., Roschelle, J. & Weisgrau, J. (2024, May). AI literacy: A framework to understand, evaluate, and use emerging technology. <https://doi.org/10.51388/20.500.12265/218>
- MIT Media Lab Personal Robots Group and the MIT STEP Lab (2024). The Daily Curriculum for middle school students. <https://raise.mit.edu/daily/>
- National Center on Improving Literacy (2022). The Science of Reading: The Basics. Washington, DC: U.S. Department of Education, Office of Elementary and Secondary Education, Office of Special Education Programs, National Center on Improving Literacy. <https://improvingliteracy.org/brief/science-reading-basics/index.html>
- National Educational Technology Plan. Office of Educational Technology. (2024, January 24). <https://tech.ed.gov/netp/>
- National Institute of Standards and Technology. (2023). AI Risk Management Framework. *Artificial Intelligence Risk Management Framework* (AI RMF 1.0). <https://doi.org/10.6028/nist.ai.100-1>
- NGSS Lead States. (2013). Next Generation Science Standards: For States, By States. Washington, DC: The National Academies Press.
- Office for Human Research Protections (OHRP). (2024, July 10). *The Belmont Report*. Department of Health & Human Services. <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/index.html>
- Office of Management and Budget. (2024, March 28). Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (Memorandum M-24-10). The White House. <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>
- Roschelle, J., Fusco, J., & Ruiz, P. (2024). Review of guidance from seven states on AI in education. Digital Promise. doi.org/10.51388/20.500.12265/204
- Ruiz, P., Armstrong, A., Karim, S., Shell, A., Singmaster, H., & Giang, M. (2024) Emerging technology acceptable use policy [Google slides]. Digital Promise. <https://bit.ly/DPAUPLanguage>
- Ruiz, P., & Fusco, J. (2024, March 31). Glossary of Artificial Intelligence Terms for Educators. CIRCLS. <https://circls.org/educatorcircls/ai-glossary>
- Southern Education Foundation. (2024, January). Outcomes Based Contracting. <https://obc.southerneducation.org/>

- Steenbergen-Hu, S., & Cooper, H. (2013). A meta-analysis of the effectiveness of intelligent tutoring systems on K–12 students’ mathematical learning. *Journal of Educational Psychology*, 105(4), 970
- U.S. Commerce Secretary Gina Raimondo Announces Expansion of U.S. Ai Safety Institute Leadership Team. National Institute of Standards and Technology. (2024, April 23). <https://www.nist.gov/news-events/news/2024/04/us-commerce-secretary-gina-raimondo-announces-expansion-us-ai-safety>
- U.S. Department of Education, Office of Educational Technology, Designing for Education with Artificial Intelligence: An Essential Guide for Developers, Washington, D.C., 2024.
- United States Department of Education. (2023, May 24). *Artificial Intelligence and the future of teaching and learning*. Office of Educational Technology. <https://tech.ed.gov/ai-future-of-teaching-and-learning/>
- United States Government Accountability Office. (2022, October). *Critical Infrastructure Protection: Additional Federal Coordination Is Needed to Enhance K-12 Cybersecurity*. <https://www.gao.gov/assets/gao-23-105480.pdf>
- The United States Government. (n.d.). Blueprint for an AI Bill of Rights Definitions. The White House. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/definitions/>
- The United States Government. (2023, October 30). Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. The White House. <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>
- Weiss, A. (2024, July 31). From Vision to Practice: Part 1. Grounding AI for Schools in Learning Science with Dr. Kristen DiCerbo. Khan Academy. <https://blog.khanacademy.org/from-vision-to-practice-part-1-grounding-ai-for-schools-in-learning-science-with-dr-kristen-dicerbo/>
- What Works Clearinghouse. (n.d.). *Institute of Education Sciences*. U.S. Department of Education. Retrieved from <https://ies.ed.gov/ncee/wwc/>
- World Economic Forum. (2024, April). Shaping the Future of Learning: The Role of AI in Education 4.0. <https://www.weforum.org/publications/shaping-the-future-of-learning-the-role-of-ai-in-education-4-0/>
- Xu, Z., Wijekumar, K., Ramirez, G., Hu, X. and Ireys, R. (2019). The effectiveness of intelligent tutoring systems on K-12 students' reading comprehension: A meta-analysis. *Br J Educ Technol*, 50: 3119-3137. <https://doi.org/10.1111/bjet.12758>