

Council of State Science Supervisors

Augmenting Safety Culture Through Better Decisions

How AI, Design, and Deliberate Choice Strengthen Safer STEM Practices: A nationally aligned framework supporting professional judgment, decision quality, and defensible safety leadership in K–12 science and STEM education

Authors

Brian P. Collins, Ph.D.
James Palcik, CHO
Edward Charles Cohen, Ed.D.

In Partnership With
Council of State Science Supervisors (CSSS)



CSSS AUGMENTING SAFETY CULTURE THROUGH BETTER DECISIONS

2026

Brian P. Collins, Ph.D., is the Chief Innovation Officer at Safer STEM, focusing on the design of safety decision support tools, instructional risk systems, and professional learning models for science and STEM educators.

James Palcik, CHO, is the Chief Executive Officer of Safer STEM and a nationally recognized school science safety authority with decades of experience advising K–12 and higher education institutions on laboratory safety, risk management, and professional practice.

Edward Charles Cohen, Ed.D., is the Associate Director of the Center for Mathematics, Science, and Computer Education at Rutgers University, NJ, and a leader in science program development and educational systems design.

These materials may include third-party content provided for informational purposes only. CSSS and Safer STEM LLC make no warranties, express or implied, regarding the accuracy or completeness of this content and disclaim all liability to the fullest extent permitted by law. Users are responsible for consulting applicable federal, state, and local laws, regulations, and current safety guidance to ensure compliance with required safety practices.

ACKNOWLEDGMENTS

The Council of State Science Supervisors (CSSS) extends appreciation to Brian P. Collins, Ph.D. (Safer STEM), James Palcik, CHO (Safer STEM), and Dr. Edward Charles Cohen (Rutgers University) for developing this leadership-focused guidance on strengthening safety culture through improved decision-making systems.

This publication was created to help state science supervisors, district administrators, safety leaders, and educators move beyond compliance-based models of safety and toward systems that improve judgment under real STEM instructional conditions.

The authors and the Council of State Science Supervisors gratefully acknowledge Ward's Science, a division of VWR, for their valued partnership in supporting the development of the CSSS Safety Reference Manuals. Their sustained commitment to advancing science education and strengthening laboratory safety continues to benefit STEM classrooms nationwide. We are especially appreciative of their collaboration in providing high-quality scientific illustrations, imagery, and technical insight, which enhance clarity, reinforce safer STEM practices, and support effective, engaging instruction across Science, STEM, and CTE programs.

About CSSS

The Council of State Science Supervisors is a national organization composed of leaders who oversee science education at the state level. CSSS members guide standards implementation, statewide policy alignment, professional learning systems, and program evaluation across K–12 education.

Through collaboration and research-informed guidance, CSSS supports science education systems that are rigorous, inclusive, and aligned with professional responsibility. This includes strengthening institutional safety culture through improved decision infrastructure, documentation systems, and leadership practice.

TABLE OF CONTENTS

Executive Summary	06
Section I. Safety Culture Is a Decision System	07
Section II. AI Decision Support in STEM Safety	09
Section III. Understanding AI’s Limits: Why Verification Is Essential	14
Section IV. Cognitive Flexibility and AI Design	16
Section V. VERA AI: Purpose-Built Decision Support for Safer STEM Instruction	19
Section VI. How Physical Design Improves Decisions	25
Section VII. Building a Safety Culture That Learns and Protects	27
Disclaimer	30
References	31

LIST OF FIGURES

Figure 1. Safety Culture as a Decision System	08
Figure 2. Framework for AI-Assisted Safer STEM Instructional Planning	11
Figure 3. Unrestrained Automation vs. Responsible Decision Support	13
Figure 4. Deterministic vs. Probabilistic Output in AI Systems: Confidence vs. Accuracy	15
Figure 5. Customized AI Tools Examine Activities Through Multiple Safety Lenses	18
Figure 6. Yesterday’s Plan vs. Today’s Conditions	24
Figure 7. Classroom-Adapted Hierarchy of Controls for Instructional Planning	25
Figure 8. Without Visual Cues vs. With Visual Cues	26
Figure 9. Proactive Safety Investment	29

EXECUTIVE SUMMARY

Safety in science and STEM education is often treated as a checklist. Wear goggles. Follow procedures. Complete the required documentation.

Checklists matter. But checklists alone do not prevent accidents.

Most safety failures occur when conditions change, and decisions do not. Class size increases. Equipment is unavailable. Ventilation systems fail. Students arrive distracted or unprepared. Under these conditions, educators must decide whether a planned activity still fits the moment. Safety culture, therefore, depends less on rules than on decision quality.

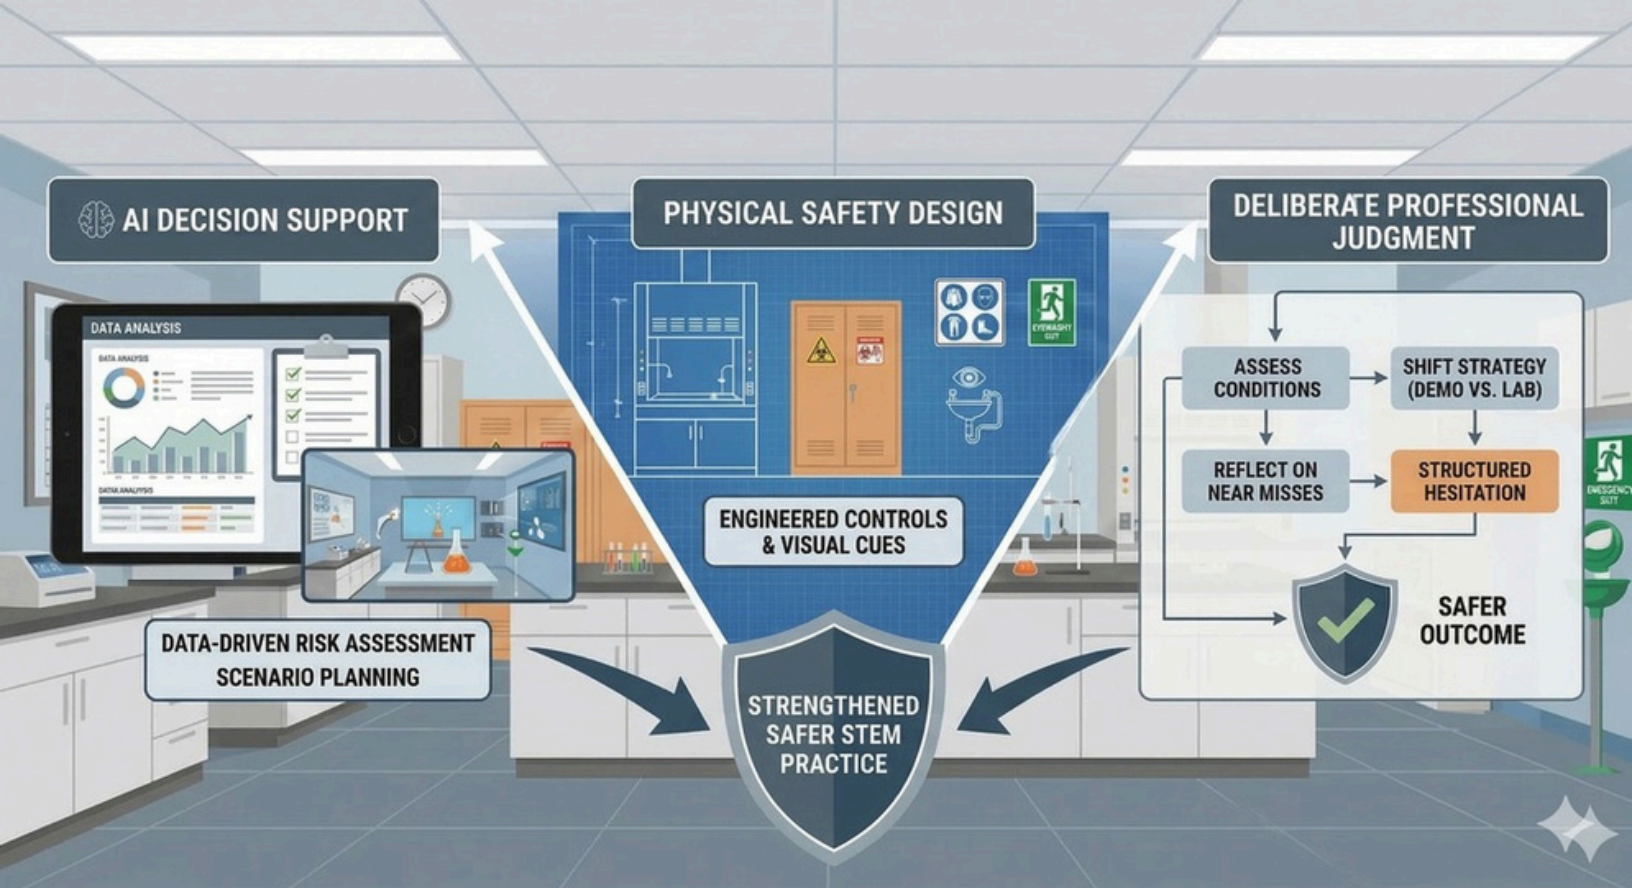
This guide examines three factors that improve safety decision-making in K-12 STEM programs:

- ▶ AI-supported decision tools that surface risks and assumptions
- ▶ Physical design features that reduce cognitive burden in real time
- ▶ Organizational systems that encourage learning from near misses

Together, these elements strengthen what high-reliability organizations understand well.

In practice, safety culture develops across three reinforcing layers: the cognitive tools that help educators make better decisions, the physical environments that make safer choices easier, and the organizational systems that turn experience into learning.

Safety is not only about compliance. It is about how people think, decide, and adapt when reality diverges from the plan.



I. SAFETY CULTURE IS A DECISION SYSTEM

Checklists remind us what to do. Judgment determines whether today is the day to do it differently.

In many schools, safety is framed primarily as ‘compliance’. Educators are expected to follow procedures and complete required documentation. While these practices are important, they capture only part of the picture. Safety failures rarely occur because rules are missing. They occur because decisions are made under uncertainty.

High reliability organizations address this challenge directly. Rather than relying only on procedures, they invest in decision quality. Professionals are trained to notice weak signals, question assumptions, and adjust when conditions shift (Weick & Sutcliffe, 2015).

One of the most valuable signals in these systems is the near miss. A near miss is an event that could have resulted in a serious injury, exposure, or damage but did not because of chance or timely intervention. Near misses provide early warnings about hidden potential hazards and resulting safety risks in our science / STEM instructional spaces.

In K–12 science and STEM programs, the same principle applies. Safety culture is not defined by what educators know. It is defined by the decisions they make when instructional conditions change. AI tools, planning protocols, and physical design features matter only insofar as they improve those decisions. Used well, they slow judgment down long enough to sharpen it.

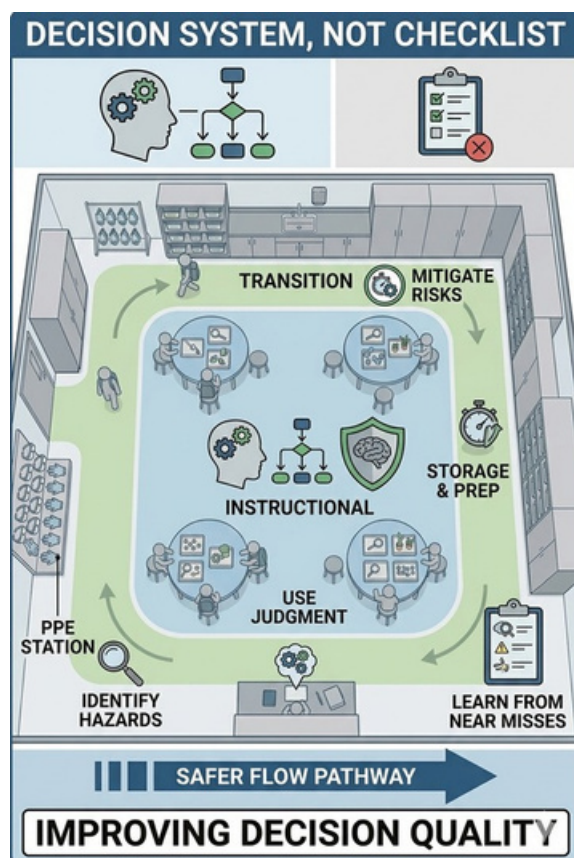
Over time, better decisions compound. Potential hazards are identified earlier, and resulting safety risks are mitigated sooner. Near misses become learning opportunities and prevention for future potential safety incidents.

Together, these elements form what can be described as a decision infrastructure, the tools, processes, professional norms, and physical conditions that shape how safety decisions are made in science / STEM instructional spaces.

That is how durable safety culture develops.

Figure 1

This figure emphasizes that safety in science / STEM instructional spaces is not a checklist, but a continuous cycle of decision-making that includes identifying potential hazards, mitigating risks, and learning from near misses. As a teacher, it reinforces the importance of actively adjusting instruction based on real-time conditions rather than relying solely on prewritten procedures.



© Safer STEM. All rights reserved.

II. AI DECISION SUPPORT IN STEM SAFETY

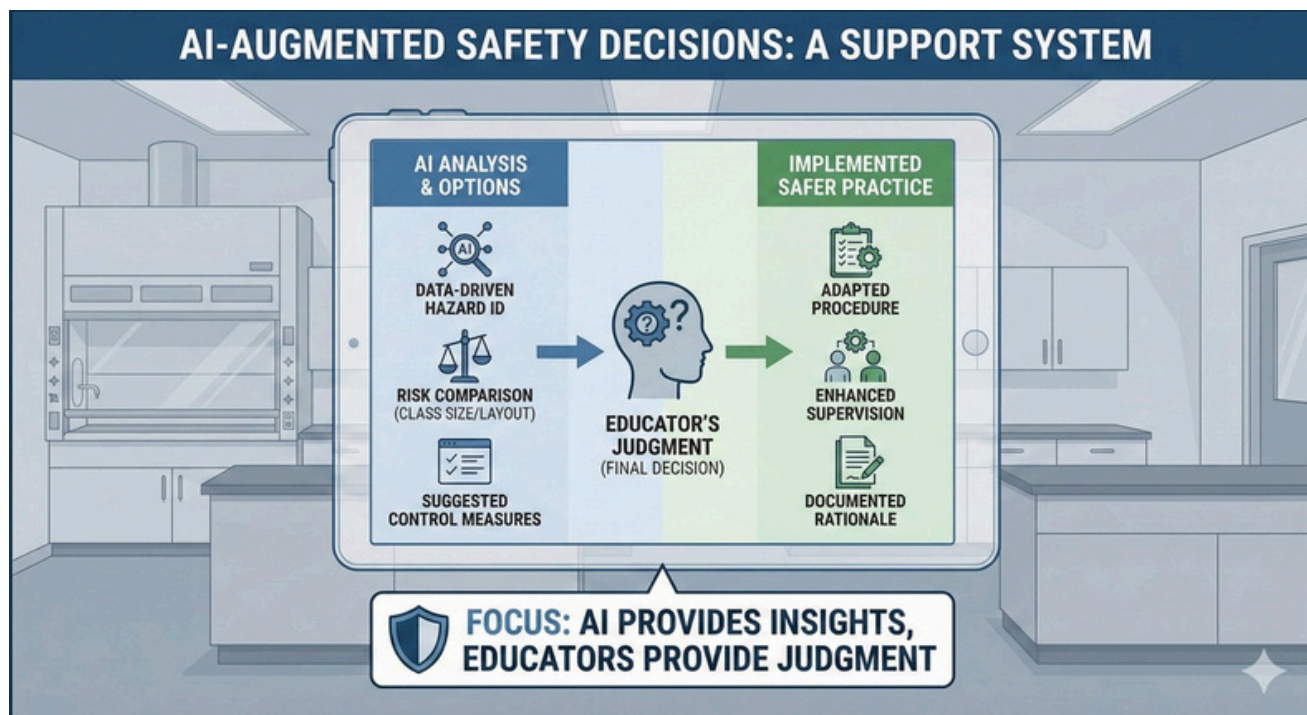
AI should not replace professional judgment. It should pause it long enough to help us notice what confidence or enthusiasm might conceal.

The most responsible use of artificial intelligence in school science is not automation. It is informed decision support.

Recent national guidance reinforces this approach. Rather than treating AI as a fixed solution, experts increasingly emphasize situational use that adapts to context and task demands (Brookings Institution, 2026).

For science safety, this distinction matters.

When introduced without structure, AI can create a false sense of confidence or certainty. When designed thoughtfully, it can strengthen judgment by surfacing questions educators might not otherwise consider based on their own experiences and understanding.



The value of AI is not that it answers questions. It's that it helps us ask better ones.

Effective AI decision support tools help educators:

- ▶ Identify potential hazards and resulting safety risks under real instructional space conditions
- ▶ Compare risk across different supervision ratios and physical room layouts
- ▶ Determine when investigations should shift to demonstrations or simulations based on potential hazard and resulting risk assessment
- ▶ Reflect on near misses to improve future decisions in instructional programming

In practice, this looks less like answers and more like structured questions. Instead of asking whether an activity is allowed, educators are prompted to ask:

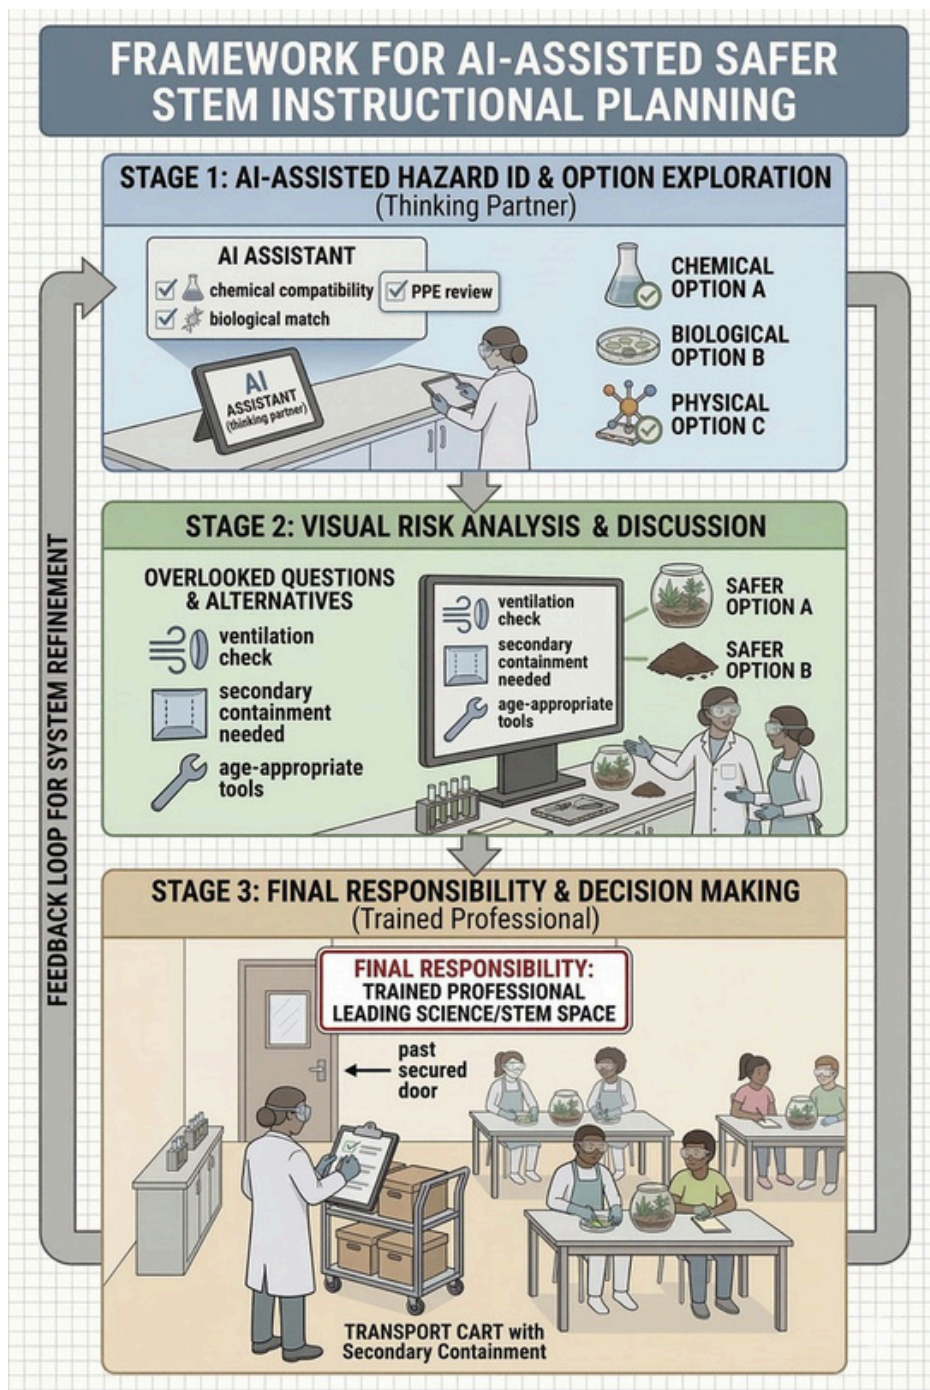
What potential hazards and resulting safety risks remain uncontrolled if class size increases?

Which engineering controls must be in place before relying solely on PPE?

Strong judgment rarely comes from certainty. It develops from asking better questions.

Figure 2

This figure outlines a three-stage process where AI supports hazard identification and option exploration, but final decisions remain with the educator. Teachers can use this model to structure planning conversations, ensuring that instructional choices are both pedagogically sound and aligned with safety expectations.



© Safer STEM. All rights reserved.

Professional Guardrail

AI systems must always operate within professional guardrails.

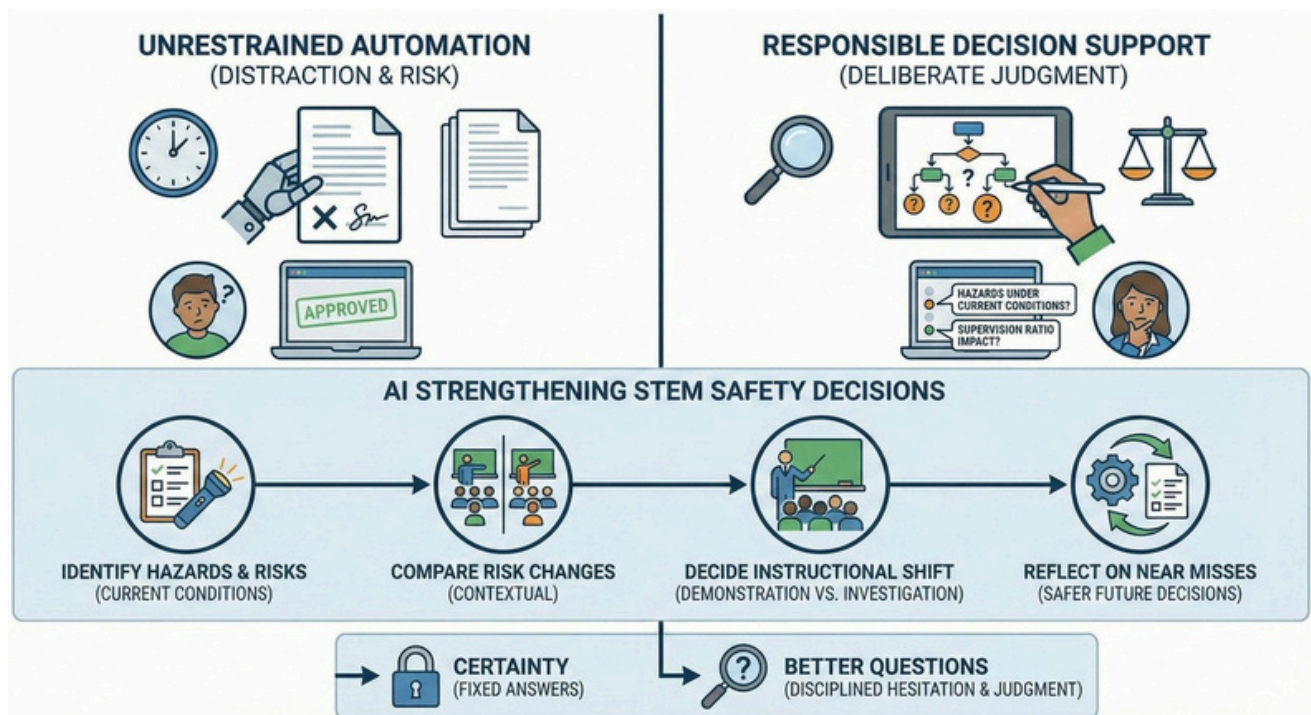
AI output should be reviewed by a trained STEM educator or Chemical Hygiene Officer before it is applied in instructional settings. AI does not assume duty and standard of care. Responsibility for safety decisions remains primarily with the educator in the room.

Overreliance on automated recommendations can increase potential hazard and corresponding risk. From a legal perspective, undocumented or unexamined AI outputs are not a defense. They represent a potential point of liability exposure for educators and administrators.

The most effective systems combine human expertise with AI-supported insight. Each strengthens the other. This should be a human-led, technology-enabled professional risk management system.

Figure 3

This figure contrasts passive reliance on AI-generated outputs with active, teacher-led decision-making supported by AI insights. It reminds educators that professional judgment, not automation, is the critical factor in managing potential hazards and resulting safety risks.



© Safer STEM. All rights reserved.

One practical discipline for leaders is learning to distinguish between what is required and what is recommended. Some safety expectations originate in enforceable regulations. Others reflect widely accepted professional benchmarks that represent responsible practice even when not written into law. Both matter. Regulations establish the minimum. Professional standards often define the level of care expected when decisions are reviewed after an incident.



III. UNDERSTANDING AI'S LIMITS: WHY VERIFICATION IS ESSENTIAL

AI systems often appear authoritative. Their responses are fluent and confident. That confidence can create the impression of accuracy.

But AI systems do not know the answer. They generate responses by predicting patterns in language based on training data.

This distinction matters.

Unlike calculators, which produce deterministic outputs, AI models estimate what an answer is likely to look like. As a result, responses may occasionally contain fabricated citations, incomplete procedures, or outdated guidance.

Training data introduces additional limitations. AI systems learn from vast collections of internet content that include both high quality information and unreliable material.

Bias can also appear when models replicate patterns present in their training data.

For educators and safety leaders, the principle is straightforward.

AI output should be treated as a starting point for professional judgment, not a final answer.

**Do not trust automatically. Verify.
Ask critical questions.**

Verification means checking AI-generated information against trusted professional sources, including regulatory standards (e.g., OSHA), adopted consensus codes (e.g., NFPA), and research-based guidance (e.g., NIOSH).

Not every authoritative source carries the same legal status. OSHA regulations may establish enforceable workplace requirements in certain jurisdictions. NFPA publishes consensus codes that become enforceable only when adopted by state or local authorities.

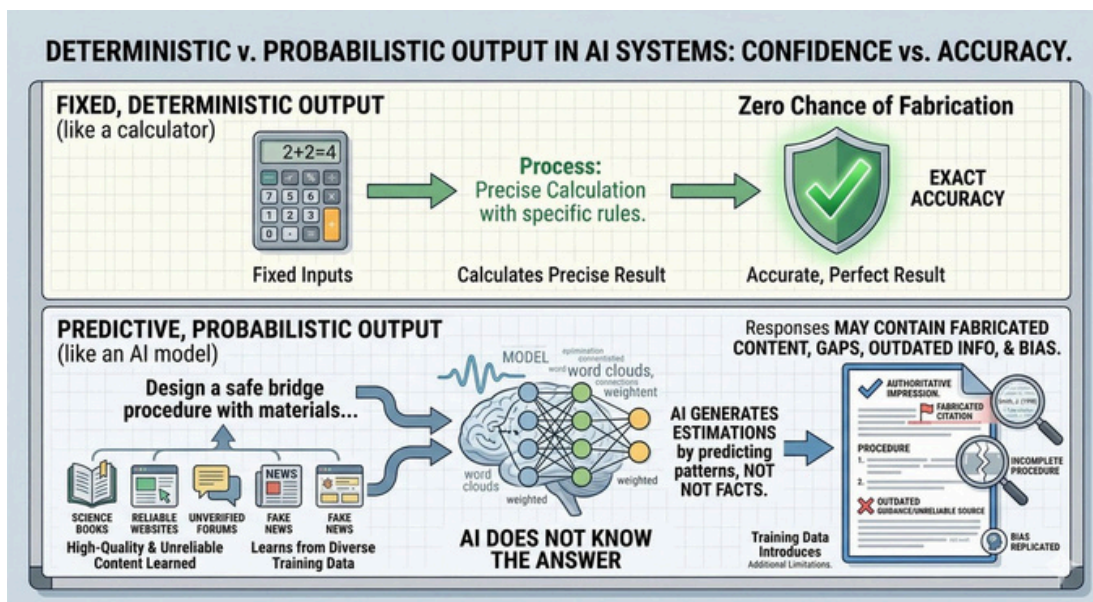
NIOSH provides research-based safety guidance that informs professional practice but is not itself regulatory. Effective safety verification requires recognizing the difference between legal requirements, adopted codes, and professional guidance.

And reliable associations such as the ACS, NSELA and NSTA accepted better professional safety practices and guidance.

Used responsibly, AI becomes a thinking partner rather than a final decision maker.

Figure 4

This figure illustrates the difference between tools that produce exact answers (like calculators) and AI systems that generate probable responses that may include errors. For teachers, it reinforces the need to verify AI-generated safety guidance against trusted standards before applying it in instruction.



© Safer STEM. All rights reserved.

IV. COGNITIVE FLEXIBILITY AND AI DESIGN

Strong decisions rarely come from a single view of a situation. They emerge when educators know how to look again through multiple safety lenses.

This habit is known as cognitive flexibility: the ability to revisit knowledge from multiple perspectives and apply it across changing contexts (Spiro et al., 1988). In science and STEM safety, cognitive flexibility means recognizing that a planned activity may appear appropriate through one lens, such as the learning objective, but look different when examined through others.



A teacher might first view an investigation through the lens of content: Does this activity help students understand the concept? That same activity must also be viewed through additional lenses, including supervision, classroom layout, engineering controls, student readiness, materials handling, and emergency preparedness. Each lens can reveal a different pattern of potential hazards and resulting risks.

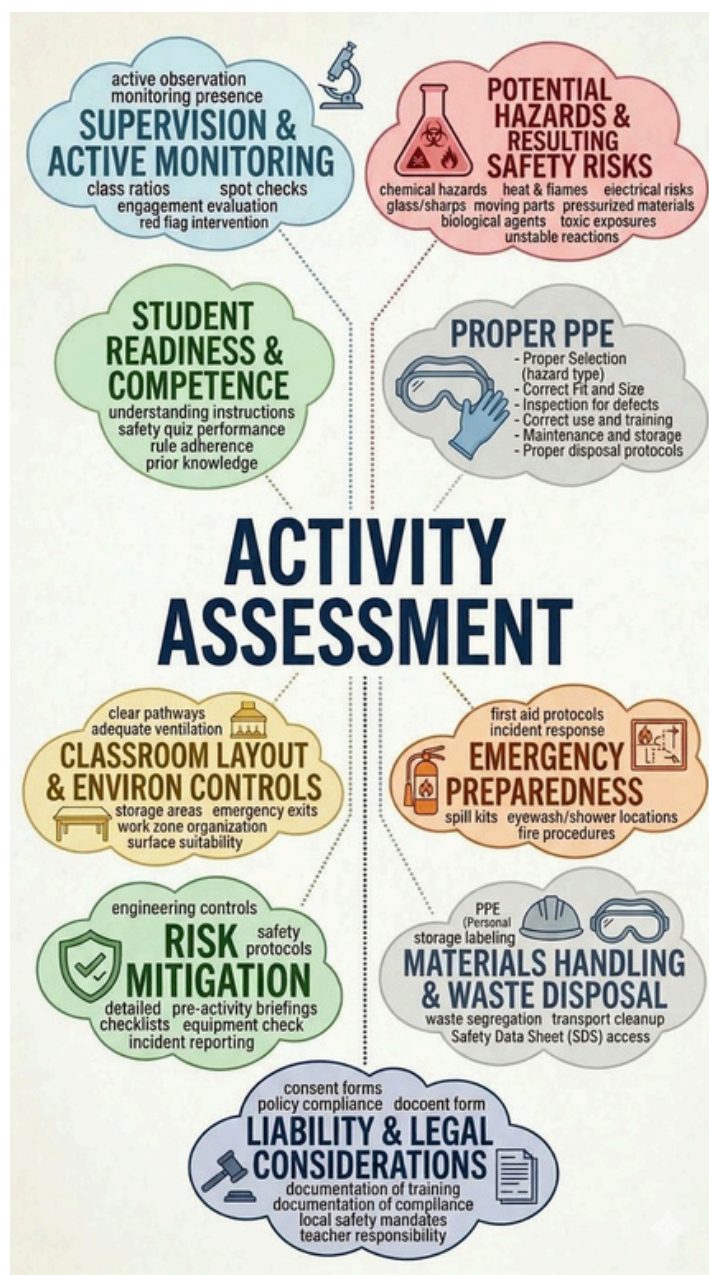
Science safety decisions work this way. The core elements of instruction may remain constant, but the safety profile changes as classroom conditions change. A procedure that is manageable with a small, prepared class in a properly equipped room may require adjustment when class size increases, equipment is unavailable, ventilation is reduced, or students need additional support.

Customized AI tools can strengthen this process by helping educators apply multiple safety lenses more consistently. The goal is not to make AI smarter than seasoned educators. The goal is to make educators' reasoning stronger by prompting them to notice conditions, questions, and trade-offs they might otherwise miss.

In this sense, safety reasoning becomes part of scientific thinking itself. Students and teachers alike learn that responsible scientific practice involves examining a problem through more than one lens before acting. Strong judgment comes from seeing the activity clearly, not from a single perspective, but through the full set of conditions that shape safer instruction.

Figure 5

This figure presents a comprehensive “activity assessment” model that includes supervision, student readiness, PPE, environment, and emergency preparedness. It helps teachers systematically evaluate an activity from multiple perspectives to identify potential hazards and reduce risk before instruction begins.



© Safer STEM. All rights reserved.



V. VERA AI: PURPOSE-BUILT DECISION SUPPORT FOR SAFER STEM INSTRUCTION

Off-the-shelf AI systems can generate useful drafts for unit planning, lesson generation, and even activity ideas. But science and STEM safety decisions require something more than fluent text. They require alignment with legal safety standards and accepted professional safety practices, awareness of instructional context, and professional judgment about potential hazards and resulting safety risks associated with instruction.

This is where purpose-built decision-support systems can become valuable. When designed around the realities of classroom instruction, these tools help educators examine potential hazards, evaluate alternatives, and document the reasoning behind instructional choices.

Purpose-built, education-aligned AI decision-support systems can provide structured guidance by integrating safety standards, instructional context, and professional practice expectations. These systems are most effective when they support hazard identification, evaluation of alternatives, and documentation of decision-making while keeping educators in the role of final decision maker.

VERA AI is designed as a decision-support environment for science educators, not simply a text generator. The system integrates general AI capabilities with curated safety resources, professional standards, and instructional frameworks so that the output educators receive is safer, more compliant, and more useful in real classrooms.

These added layers act as structured guidance. They do not change the underlying AI model.

Instead, they shape how the model retrieves information, evaluates risk, and explains recommendations.

Just as important, they keep the educator in the decision loop. VERA AI produces drafts and analyses, but the teacher verifies, refines, and adapts the output for the specific instructional context in their STEM program for their students and instructional space features.

In this way, AI becomes a scaffold for professional reasoning rather than a substitute for it.

Example 1: Activity Risk Assessment and Safer Substitutions

One of the most valuable applications of VERA AI is activity risk assessment.

A general AI system may generate a laboratory procedure that technically demonstrates a scientific concept but fails to consider safer alternatives. The system may simply reproduce a traditional experiment found somewhere in its training data or source library.

Purpose-built decision-support systems approach this challenge differently from general-purpose AI tools. Instead of producing generic responses, they are designed to evaluate instructional scenarios through multiple safety lenses relevant to classroom practice.

The platform evaluates the activity through multiple perspectives. These include safety standards, greener chemistry practices, instructional goals, and student readiness. The result is a formal hazard analysis and activity risk assessment that does more than identify hazards. It explores whether the activity itself can be made potentially safer.

Consider a common oxidation-reduction investigation.

A traditional version of the experiment might suggest the use of zinc powder, which presents inhalation and handling hazards in a classroom setting. VERA AI flags this concern and proposes safer substitutions such as mossy zinc, or in some cases alternative demonstrations that illustrate the same redox principles with significantly reduced chemical hazard and resulting risk.

This reflects a core principle of the hierarchy of controls: substitution reduces risk at the source.

Compliance frameworks built into the system ensure that these substitutions align with recognized guidance from organizations such as the American Chemical Society (ACS) and the National Science Teaching Association (NSTA).

Pedagogically, the substitution becomes a teaching opportunity.

Teachers can prompt students to ask:

**Why was this material changed?
What hazards were reduced?
What trade-offs were considered?**

Students learn chemistry content and safety reasoning at the same time.

Example 2: Department Safety Leadership and Professional Communication

VERA AI also supports Science and STEM educators in leadership roles.

Imagine a teacher preparing to present at a monthly STEM department meeting on laboratory safety practices.

A general AI system might produce a visually polished presentation. However, it may omit important compliance considerations or promote practices that are inconsistent with recognized legal safety standards and accepted professional guidance.

VERA AI approaches the task with additional structure.

Compliance frameworks built into the system ensure that recommendations align with OSHA or PEOSH expectations, ACS safety guidance, and NSTA safety recommendations.

Transparency tools provide explanatory notes showing why specific safety practices are highlighted. These notes allow the teacher to understand the reasoning behind the recommendations rather than simply presenting them.

The teacher remains the decision maker. VERA AI highlights options and explains the reasoning behind them, but the educator evaluates whether the recommendation fits the class, the facility, and the learning goals.

In this way, the tool supports both safer practice and deeper professional judgment.

Pedagogical guidance also shapes the presentation.

Rather than focusing only on science safety rules, the slides model how safety can be taught as part of scientific thinking. Teachers are encouraged to discuss hazard identification, decision making, and risk mitigation as part of scientific practice.

The result is a presentation that does more than share information. It strengthens and elevates the overall safety culture across the department and makes it professional, personal and profound.

For early-career educators, structured decision support can function much like a mentor, helping them recognize risks, ask better questions, and develop stronger professional judgment over time. For experienced educators, it becomes a way to reinforce accepted better professional practices without rebuilding materials from scratch each time.

Again, the teacher remains central to the equation. The educator reviews the draft, adjusts examples, and tailors the message to the department's specific context and topical needs.

Example 3: Scenario-Based Training for Educators

VERA AI can also function as a professional learning environment for teachers.

Many activity and safety decisions occur under conditions that differ significantly from the original lesson plan. Equipment fails. Class sizes change. Student readiness varies.

Scenario-based training helps educators practice navigating these situations before they occur in real science / STEM instructional spaces, from a planning and emergency response, as well as hazard and risk recognition and mitigation proactively.

VERA AI can generate structured scenarios that allow educators to explore decision making in a low-risk environment, based on unlimited variables that could be included found in typical and custom science / STEM instructional spaces.

For example, a teacher might be presented with the following situation:

A planned investigation involves using a known volatile solvent. The class size has increased by five more students. The room ventilation system is functioning but operating at reduced capacity. Less than ideal instructional settings.

The system asks the educator to evaluate several potential options:

Should the activity proceed as planned? Should the experiment shift to a teacher demonstration? Would a simulation better support the learning objective under current conditions?

VERA AI then explains the safety implications of each choice and references relevant safety guidance. This allows the educator to make an informed determination and choose the option that meets curricular objectives as well as safer conduct in their science / STEM instructional spaces.

The goal is not to provide a single correct answer. The goal is to help educators practice the reasoning process that strong safety decisions require.

Through repeated exposure to these scenarios, teachers strengthen the habit of evaluating instructional conditions before proceeding with an activity.

In this way, VERA AI becomes a scaffold for professional learning. It supports educators as they develop the judgment that defines strong and defensible safety leadership.

Why Teacher Judgment Remains Central

Across all of these examples, one principle remains constant. AI can surface insights, but it cannot assume responsibility for safety decisions. Teachers verify the output, refine the recommendations, and determine whether the guidance fits the classroom reality in front of them. This partnership between human expertise and structured decision support is where AI becomes most valuable. When used thoughtfully, systems like VERA AI do not replace professional judgment. They strengthen it.

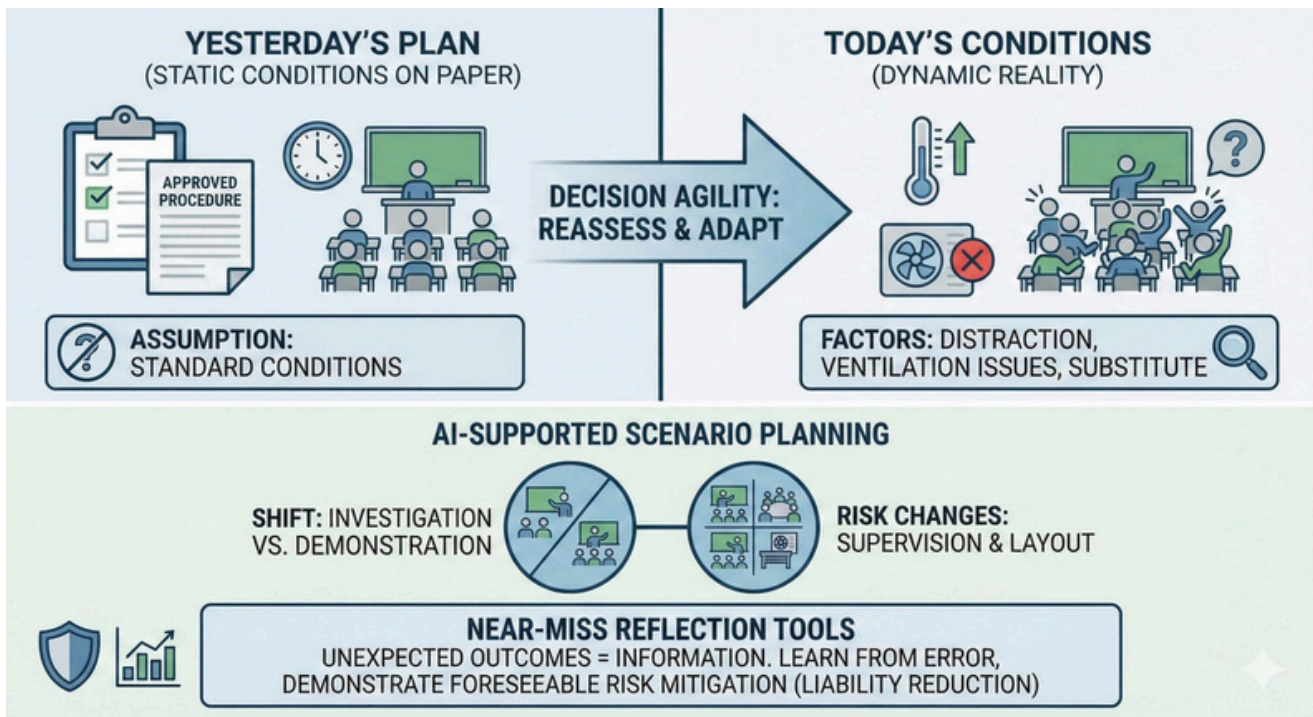
Tools like VERA AI can also contribute to professional learning. By repeatedly surfacing tradeoffs, supporting case-based learning, highlighting patterns of hazard and risk recognition and mitigation, and prompting personal reflection on near misses, the decision-support systems help educators develop the habits of mind associated with strong safety leadership.

For district leaders, the key question is not whether AI will influence safety decisions; it already does. The more important question is whether those tools are intentionally selected, evaluated, and governed. District leaders should distinguish between general-purpose AI tools and district-approved, vetted systems. Clear guidance ensures that AI strengthens decision quality without introducing unmanaged risk.

AI tools used in STEM programs must align with district data privacy policies, student information protections, and approved technology use guidelines. Leaders should establish review and approval processes for AI systems used in instructional planning and safety decision support.

Figure 6

This figure highlights how static lesson plans must be adapted to dynamic classroom realities such as changes in class size, equipment, or environmental conditions. It encourages teachers to reassess and modify activities in real time to maintain safer instructional practices.



© Safer STEM. All rights reserved.

VI. HOW PHYSICAL DESIGNS IMPROVE DECISIONS

Under pressure, people do not rise to the level of policy.

They fall to the level of design. The physical environment of a STEM instructional space strongly influences safety decisions. Clear aisles, visible eyewash stations, labeled master shutoffs, and accessible PPE reduce cognitive load during instruction.

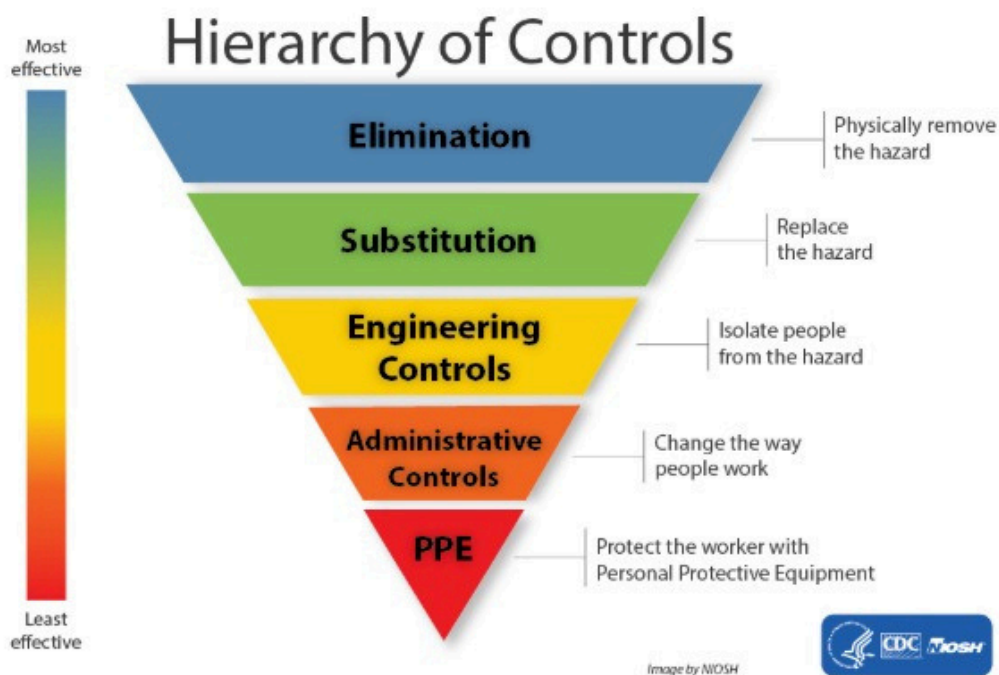
When engineering controls are visible and intuitive, educators intervene earlier. Students recognize expectations more quickly. Near misses are detected before they escalate.

This principle reflects the hierarchy of controls. Engineering controls rank above PPE because they reduce reliance on individual behavior.

Safer and intelligent design can absorb part of the decision burden.

Figure 7

This figure applies the hierarchy of controls to classroom settings, prioritizing elimination and substitution over reliance on PPE. Teachers can use this framework to design lessons that reduce risk at the source rather than depending solely on student behavior.



Visual Cues as Decision Anchors

Well-designed visual cues function as external judgment aids.

Signage, icons, and color-coded systems help standardize decisions across experience levels. They also reduce variability when substitutes or novice educators are present.

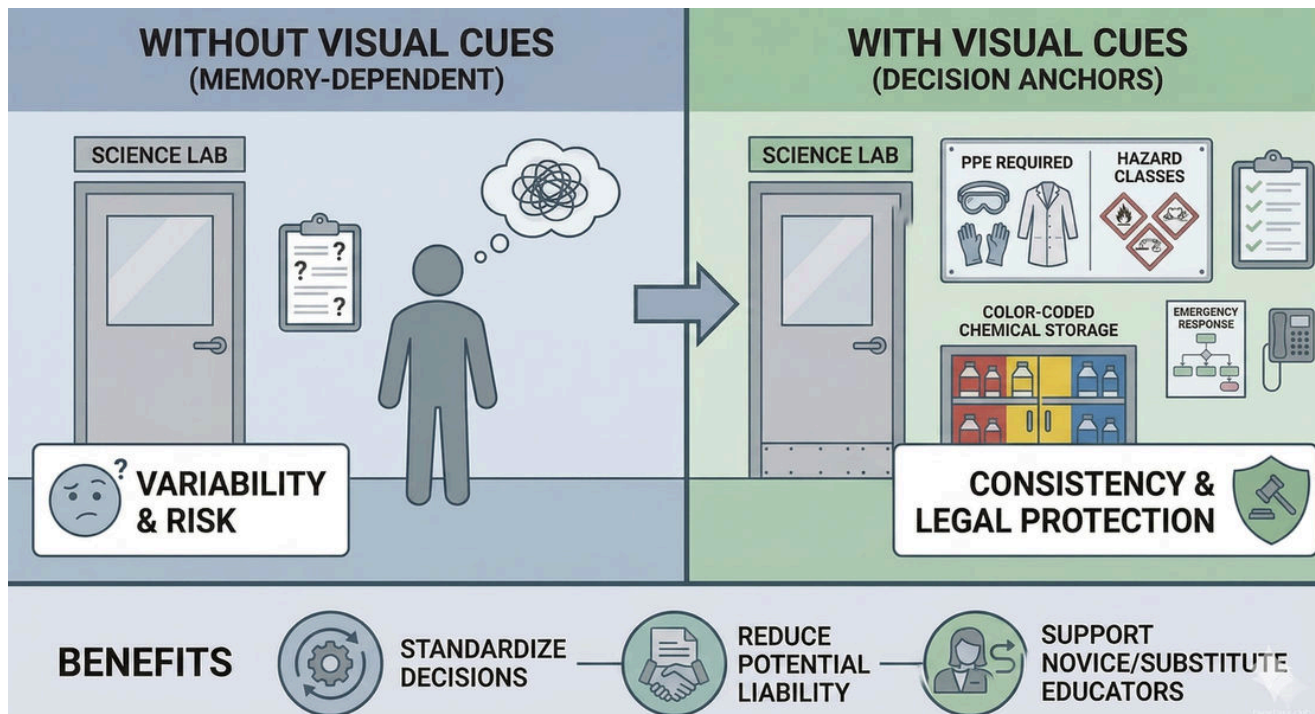
Examples include:

- ▶ PPE icons posted at laboratory entrances
- ▶ Color-coded chemical storage by hazard class
- ▶ Emergency response flowcharts near phones and exits


Consistency improves safety and provides legal protection by demonstrating that risk mitigation does not rely solely on memory.

Figure 8

This figure demonstrates how visual supports like signage, color-coding, and posted expectations improve consistency and reduce errors in science / STEM instructional spaces. It underscores for teachers that well-designed environments support safer decisions, especially for novice or substitute educators.



© Safer STEM. All rights reserved.



VII. BUILDING A SAFETY CULTURE THAT LEARNS AND PROTECTS

A culture that punishes mistakes hides risk. A culture that studies mistakes reduces it.

Safety culture functions as a system: decisions, tools, environments, and institutional norms interact to shape outcomes.

Strong safety cultures treat incidents and near misses as valuable information.

Near misses, events where potential hazards and resulting safety risks were present but injury did not occur, serve as early warning indicators.

In practice, a near miss is an event that could have resulted in injury, exposure, or equipment damage but did not, often because someone noticed the risk in time or circumstances intervened. These moments are invaluable signals. When organizations treat them as learning opportunities rather than sources of blame, they become one of the most powerful tools for preventing future incidents.

Effective systems encourage reflection after investigations, provide a shared language for discussing risk, and support continuous improvement grounded in evidence.

From a legal perspective, learning organizations document adaptation. They can demonstrate how decisions evolved and how foreseeable risks were addressed.

What This Means for Leaders

For state science supervisors, district administrators, and safety and program leaders, the implication is clear.

Invest in decision infrastructure, not just compliance infrastructure.

Facilities and equipment matter. But the greatest reduction in hazards and corresponding risks comes from improving how decisions are made before something goes wrong.

This includes:

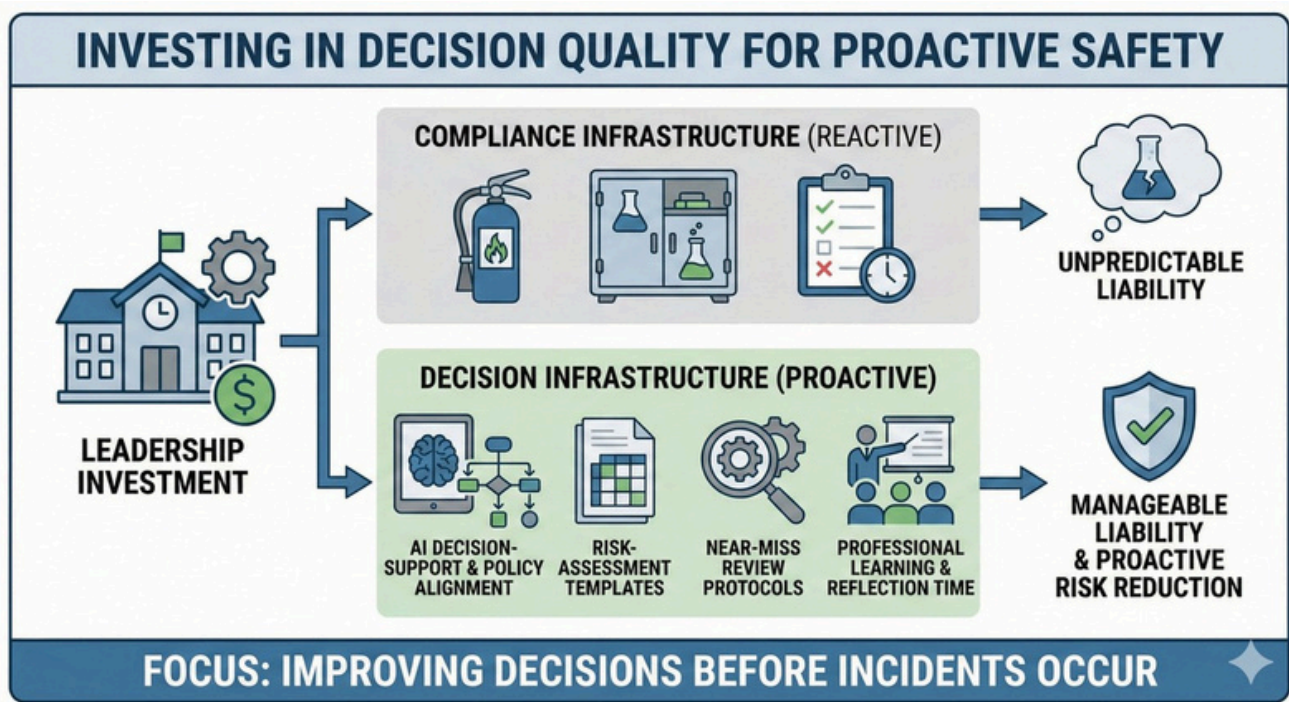
- ▶ AI decision support tools aligned with policy and safety standards
- ▶ Structured hazard analysis and resulting risk assessment templates
- ▶ Near-miss review protocols
- ▶ Professional learning on bias and judgment
- ▶ Time for educators to reflect and revise their professional practice

When leaders invest in decision quality, safety becomes proactive rather than reactive.

Safety culture is built one decision at a time.

Figure 9

This figure contrasts reactive compliance-based approaches with proactive investment in decision-making systems such as risk assessment tools and professional learning. It highlights for educators and leaders that improving decision quality before incidents occur is the most effective way to reduce risk and strengthen safety culture.



© Safer STEM. All rights reserved.

Disclaimer

These materials may include content provided by third parties and are made available solely for informational and convenience purposes. CSSS and Safer STEM LLC make no representations or warranties, express or implied, regarding the accuracy, completeness, or reliability of such content, and disclaim all warranties to the fullest extent permitted by law, including but not limited to implied warranties of merchantability, non-infringement, and fitness for a particular purpose. In no event shall CSSS and/or Safer STEM LLC or its officers, employees, or affiliates be liable for any loss, injury, or damages of any kind, including incidental or consequential damages, arising from the use of these materials. Users are advised to consult applicable occupational safety and health, environmental, and laws at the federal, state, and local regulations, as well as public health and safety guidelines to ensure compliance with current laboratory and STEM instructional space safety requirements.

The contents of this safety manual are proprietary to the authors at Safer STEM LLC, and are considered intellectual property (IP). If referencing this safety manual, please include the name and authors. All content contained in this document has been attributed to the original publisher in the Reference section and every attempt has been made to provide a comprehensive, current, and compliant resource to be used by the science & STEM department educators, School Building Administrators, Chemical Hygiene Officers and/or the designate in order to elevate the level of safety awareness across the school and school district. Almost every school in the United States must follow the OSHA 29CFR 1910.1450 Laboratory Standard (or state equivalent safety program) for their science department safety requirements. Always follow your approved Chemical Hygiene Plan (CHP) or equivalent safety plan in your local school district jurisdiction and consult with your Chemical Hygiene Officer for more details.

REFERENCES

Brookings Institution. (2026). A new direction for students in an AI world. Retrieved January 2026, from <https://www.brookings.edu/wp-content/uploads/2026/01/A-New-Direction-for-Students-in-an-AI-World-FULL-REPORT.pdf>

Edmondson, A. (2018). *The fearless organization: Creating psychological safety in the workplace for learning, innovation, and growth*. Wiley.

Grant, A. (2021). *Think again: The power of knowing what you don't know*. Viking.

National Science Teaching Association. (n.d.). Safety and school science instruction. Retrieved January 2026, from <https://www.nsta.org/nstas-official-positions/safety-and-school-science-instruction>

Spiro, R. J., Coulson, R. L., Feltovich, P. J., & Anderson, D. K. (1988). Cognitive flexibility theory: Advanced knowledge acquisition in ill-structured domains. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp. 375–383). Lawrence Erlbaum Associates.

Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the unexpected: Sustained performance in a complex world* (3rd ed.). Wiley.



©Safer STEM, 2026. All Rights Reserved