Getting Down to FACTS

Research Brief | May 2026

Artificial Intelligence, Learning, and the Future of California Schools

Chris Agnew, Cristina Barnard, Elizabeth Kozleski, Victor Lee, Susanna Loeb, Lauren Ziegler

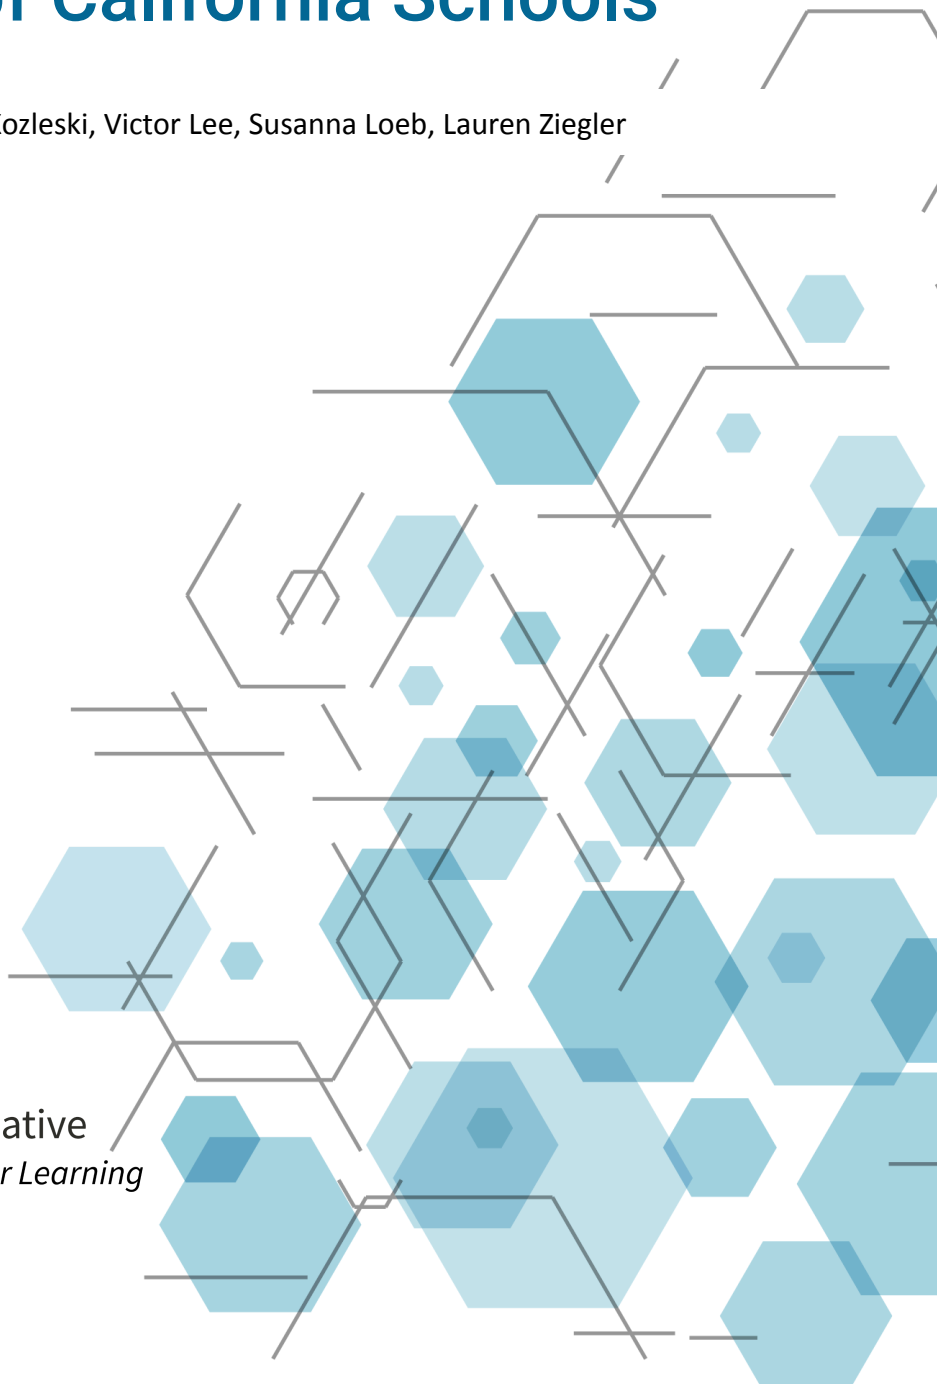

Stanford | SCALE Initiative
Accelerator for Learning

Introduction

AI is already part of daily life in California schools. Teachers are using it to draft lesson plans, adapt materials, and manage administrative tasks. Students are using it to brainstorm, summarize, solve problems, and draft writing. School and district leaders are exploring how these tools might support instruction, communication, data use, and operations. This adoption has moved quickly, often faster than the development of clear policies, shared expectations, and sustained professional learning.

AI creates real possibilities as well as real concerns. These tools may help schools provide more targeted instruction, more timely feedback, and better access to learning for students with diverse needs. They also raise concerns about misinformation, bias, privacy, and patterns of use that may weaken revision, persistence, and sustained thinking. The research base is still limited, and current evidence does not yet show the long-term effects of AI on student learning, motivation, critical thinking, or well-being across settings and student groups.

The technology studies in Getting Down to Facts begin with student development. Schools are responsible for helping students build knowledge, reasoning, curiosity, agency, persistence, and the ability to work with others. The studies summarized here consider AI in relation to those capacities and to the conditions that support their development. Starting with student capacities keeps the focus on the most important educational experiences for California schools.

This brief draws on the Getting Down to Facts III technical reports to describe AI use in California schools. The central policy question is how schools and systems will shape AI use. AI may help make learning environments more responsive, flexible, and centered on students' needs, but those possibilities depend on stronger guidance, professional learning, data protections, and implementation capacity than most schools currently have. California schools are already making decisions about instruction, privacy, academic integrity, communication, and procurement without strong shared infrastructure to guide them.

Key Findings

1 **New technologies, particularly AI, present both substantial opportunities and substantial risks for California's education system.**

Recent advances in AI, including generative tools, adaptive systems, and data-driven platforms, could change teaching, learning, and school operations in significant ways. These tools may help schools provide more personalized instruction, faster feedback, and more flexible supports for students. They also raise serious concerns about misinformation, bias, privacy, overreliance, and possible effects on students' motivation, critical thinking, and well-being.

2 **If developed and used effectively, AI tools can help California address persistent barriers that limit the spread of effective educational practices.**

Research has already identified the learning experiences that best support student development. In many schools, those practices remain hard to sustain because of organizational constraints, limited capacity, and uneven access to resources. AI could help reduce some of those barriers by supporting coordination, adaptation, and access to information that teachers and schools often lack the time or staff to provide consistently.

3 **California schools are already using powerful AI tools, often without sufficient guidance, policies, or system-level support.**

AI use is already widespread in California schools, especially for instructional planning and administrative work. Many schools and districts, however, still do not have formal policies for AI use. Guidance on safe, effective, and equitable implementation remains limited, leaving educators and local leaders to make consequential decisions with uneven support.

4 **Educators and system leaders are seeking significantly more guidance and professional learning, both on how to use AI effectively and how to mitigate its risks.**

Across the state, many educators report that they do not feel well prepared to use AI in teaching and learning. They are asking for clearer policies, stronger professional development, and more support on issues such as academic integrity, privacy, bias, and the role AI should play in student learning. Students also have limited access to AI literacy instruction, leaving many to use these tools without much formal preparation.

The Evidence Behind These Findings

New technologies, particularly AI, present both substantial opportunities and substantial risks for California's education system

These tools create meaningful opportunities and substantial risks. On the opportunity side, Barnard Gonzalez, Agnew, and Loeb identify five learning experiences that decades of research link to student success: targeted instruction responsive to individual trajectories, real-world learning connected to authentic problems, autonomy-supportive environments with meaningful student choice, rich discussions that build reasoning and perspective-taking, and sustained relationships with caring adults. AI has potential to make these experiences more consistently available through adaptive instruction, continuous feedback, and coordination tools that reduce the logistical burden on teachers.

Evidence from special education illustrates this potential. Kozleski documents how AI-enabled tools such as real-time progress monitoring, adaptive content delivery, speech-to-text supports, and multimodal materials can help schools serve students with disabilities more responsively, especially in settings where specialist shortages and labor-intensive documentation have long constrained what is possible.

The risks are just as documented. AI systems can generate confident misinformation. They can embed the deficit-based assumptions and historical biases present in their training data, a concern Kozleski flags specifically for tiered support systems that already show patterns of racial and linguistic disproportionality. Teachers in California surveys identify overreliance as their leading concern for student use. They worry that students will turn to AI in ways that short-circuit the practice, revision, and independent thinking that learning actually requires (Lee). And schools serving the students with the greatest needs are often the least positioned to implement AI-integrated models well, meaning careless adoption risks deepening the very inequities it might otherwise help address.

If developed and used effectively, AI tools can help California address persistent barriers that limit the spread of effective educational practices

California has many examples of good teaching. The problem it faces is how to create the organizational conditions of schooling to make those practices sustainable at scale.

Barnard Gonzalez, Agnew, and Loeb identify five institutional barriers that consistently get in the way. Fixed schedules and age-graded classrooms limit responsive instruction. Standardized accountability systems reward what is easily measured, pushing teachers toward uniform content even when students need differentiated pacing and open-ended inquiry. Teachers are asked to differentiate, facilitate,

mentor, and coordinate, often with too little planning time or sustained coaching. Most instructional materials are built for standard grade-level entry points rather than for the range of readiness levels present in actual classrooms. Uneven access to technology, translation supports, and community networks further concentrates enriched learning in the schools already best positioned to provide it.

These barriers compound. A teacher working within a rigid schedule, narrow accountability pressures, limited materials, and thin support is operating under very different conditions from one with time, flexibility, and resources. That helps explain why promising practices remain localized and fragile.

AI has potential to serve as infrastructure that could help reduce some of these barriers. It could connect learning data with scheduling and staffing to support more flexible grouping. It could embed assessment into instruction so that teachers have a clearer view of what students understand while learning is still underway. It could support professional learning through simulation and feedback tools that are available more regularly than occasional observations. It could also help align content to authentic challenges while adjusting for varied readiness levels, without requiring teachers to create differentiated materials entirely on their own. Used in these ways, AI supports the work good teachers already do and may make those practices easier to sustain across more settings (Barnard Gonzalez et al.).

California schools are already using powerful AI tools, often without sufficient guidance, policies, or system-level support

AI use in California schools is already widespread. In the California Department of Education survey, roughly 84 percent of respondents reported AI use in their schools or districts, with instructional planning and administrative tasks as the most common areas of use (Lee). Teachers are using AI to prepare lessons, adapt materials, and manage paperwork. School and district leaders are also looking to these tools for efficiency and data analysis.

Policy and governance have not developed at the same pace. About three-quarters of survey respondents reported that their schools or districts did not yet have a formal AI policy, a pattern that is consistent with principal interview data collected for GDTF III (Lee). In many places, teachers have turned to personal accounts and unapproved tools because district guidance and approved options have not kept up with day-to-day demands. Lee describes this as "shadow AI use." The pattern raises concerns about privacy, procurement, and responsibility for implementation.

This gap between use and governance leaves schools making consequential decisions without much shared structure to guide them. Those decisions include how AI is used in instruction, what standards apply to academic integrity, how student data are handled, and which tools schools are willing to

adopt. The issue is not isolated misuse by individual educators. It is a system in which adoption is fragmented, variable, and has moved ahead of policy, guidance, and oversight.

Educators and system leaders are seeking significantly more guidance and professional learning, both on how to use AI effectively and how to mitigate its risks

Educators are asking for more training than schools are currently providing. Eighty-four percent of principals reported that their schools provided between zero and five hours of AI-related professional development in the past school year. National data show that 57 percent of teachers have received no AI training and would like it. California survey data also indicate low levels of staff preparedness, with a median readiness rating of 2 on a five-point scale, even among educators engaged enough to respond to a statewide survey (Lee).

The training educators want is fairly specific. In a California district survey of more than 1,400 teachers, respondents asked for support with instructional integration, evaluation of AI outputs, privacy, bias, and academic integrity (Lee). Elementary teachers emphasized foundational understanding of how AI works. Secondary teachers were more likely to emphasize helping students think critically about it. Across grade levels, educators were looking for guidance they could use in their classrooms immediately.

Students are also encountering AI both in and outside of school with limited formal preparation. Few California schools are teaching AI literacy systematically, and where it does appear, it is often brief and disconnected from the rest of the curriculum. One of the clearest patterns in the data is that schools with a formal AI policy are 7.5 times more likely to provide AI literacy instruction than schools without one (Lee). That pattern suggests that student preparation depends in part on whether schools have developed a broader structure for thinking about AI use. Schools that have written policies are also more likely to be attending to what students need in order to use these tools well.

Implications for California

The research points to five areas where the evidence has direct bearing on decisions California is now facing.

A clearer educational standard for AI use

AI use needs a clearer educational standard. Research on student development provides a stronger basis for evaluating AI than adoption rates or efficiency gains alone. Barnard Gonzalez, Agnew, and Loeb identify the capacities that matter for long-term student success: academic knowledge, higher-order thinking, autonomy, motivation, belonging, and self-regulation, along with the learning

experiences that help develop them. With the California State Board of Education currently working to finalize a statewide Portrait of a Graduate, these student capacity goals already sit at the center of California policy. A useful standard for evaluating AI tools is whether they expand student access to targeted instruction, authentic learning, rich discussion, and sustained adult relationships, or whether they displace those experiences.

Organizational uses of AI

Some of the most consequential uses of AI may be organizational rather than instructional. The evidence suggests that AI's potential extends beyond individual classroom tools. Barnard Gonzalez, Agnew, and Loeb identify five institutional barriers that have historically limited the spread of effective practice: rigid structures, narrow accountability, limited professional capacity, inflexible materials, and unequal access. AI systems designed to address those barriers through coordination, adaptive assessment, professional learning support, and communication tools may have broader effects than any single tutoring or content-generation product. Kozleski makes a similar point in special education, where the opportunity lies in making support systems more flexible and responsive.

Evaluation and learning from implementation

The papers make clear that the evidence base on AI in California schools is still thin. Lee notes that data on AI use in schools are currently "scattered and scarce," and that rigorous studies on impact and efficacy for California schools are largely absent. Barnard Gonzalez, Agnew, and Loeb similarly observe that while near-term tools can be deployed within existing infrastructure, the more consequential question is whether they actually expand access to the learning experiences that matter. Answering that question requires better systems for gathering information about AI use, expenditures, and outcomes across schools and student populations. Structured opportunities for early implementation, through pilot sites, district partnerships, or research collaborations, could help California build the practical knowledge that policy, guidance, and professional learning will eventually need to draw on.

Governance, guidance, and oversight

Without clear state policy, the education technology market is shaping education instead of education shaping the technology market. Adoption has moved well ahead of formal policy, and that gap is consequential. The result is what Lee describes as "shadow AI use," with teachers relying on personal accounts and unapproved tools because official guidance has not caught up with day-to-day demand. That pattern raises concerns about student data privacy, procurement accountability, and consistency in how academic integrity is handled across schools and districts. The current policy gap leaves districts making decisions about privacy, procurement, academic integrity, and acceptable instructional use without much shared structure to guide them.

Clear state rules for procurement, interoperability, data privacy, and alignment with California's learning goals for AI could shape the education technology market. Those expectations could encourage developers to design tools that reflect the priorities and needs of California students. As AI becomes a larger part of young people's lives both inside and outside of school, K-12 policy alone cannot fully govern students' technology use. Agencies responsible for health and well-being, including California Health and Human Services and the California Department of Public Health, also shape the conditions under which students use technology. Coordination between the Department of Education and these agencies could help the state address concerns about data use, student safety, and mental health more fully.

Professional learning and student AI literacy

Professional learning and student AI literacy are both underdeveloped. Eighty-four percent of California principals reported providing zero to five hours of AI-related professional development in the past school year, while national data show that 57 percent of teachers have received no training and would like it (Lee). Staff readiness ratings from the CDE survey had a median of 2 on a five-point scale. Few schools are teaching AI literacy to students in any systematic way, though schools with a formal AI policy are 7.5 times more likely to provide AI literacy instruction than schools without one (Lee). Kozleski's analysis of special education reinforces the same concern. Without deliberate attention to preservice and in-service educator preparation and equitable access to tools and support, technology adoption may reproduce existing disparities rather than reduce them.

Conclusion

California is still early in this transition. AI use is already widespread in schools, but the surrounding systems are much less developed. Policies remain uneven. Professional learning remains thin. AI literacy for students remains limited. School and district staff are making decisions now about instruction, privacy, academic integrity, communication, and procurement under conditions that are still unsettled.

California already has a strong body of evidence on the capacities students need and the learning experiences that support them. The studies summarized in this brief assess AI against that evidence. The central question they raise is whether AI use will expand access to those learning experiences, particularly for students who have had the least access to them.

The California evidence also suggests that AI is unlikely to settle into a single role. In some places it will function mainly as a productivity tool. In others it may support more ambitious forms of personalization, accessibility, and coordination. In others it may deepen existing inequalities, because the conditions that support stronger implementation are already concentrated in better-resourced

settings. Which of these patterns becomes more common will depend in part on how state and local systems respond as current uses solidify into routine practice.

The evidence points to a state with substantial interest, widespread experimentation, and uneven guidance. How AI shapes California schools will depend on whether the infrastructure for policy, professional learning, and student preparation develops alongside adoption, or continues to trail it.

GDTFIII Technical Reports Referenced

Barnard Gonzalez, C., Agnew, C., & Loeb, S. (2026). *The learning experiences that matter and AI's role*. *Getting Down to Facts III*, SCALE Initiative, Stanford University.

Kozleski, E. (2026). *Leveraging technology for flexible, equitable special education*. *Getting Down to Facts III*, SCALE Initiative, Stanford University.

Lee, V. (2026). *AI and education across California schools: Calls for more professional development and AI literacy*. *Getting Down to Facts III*, SCALE Initiative, Stanford University.