SJSU | CENTER FOR INNOVATION IN APPLIED EDUCATION POLICY

# From AI Adoption to Pedagogical Practice: A Learning-Centered Framework for Evaluating AI Policy in K–12 Education

## Introducing the Five Pillars of Practice (5POP) and AI Review Cycles Evaluation Framework

Brent Duckor, Ph.D.
Executive Director
Center for Innovation in Applied Education Policy (IAEP)
San José State University

Carrie Holmberg, Ed.D.
Assistant Director
Center for Innovation in Applied Education Policy (IAEP)
San José State University

Jason Ravitz, Ph.D.
President and Founder
Evaluation by Design

Release date August 9, 2026

Revised August 10, 2026

Recommended citation: Duckor, B., Holmberg, C., & Ravitz, J. (2026). *From AI adoption to pedagogical practice: A learning-centered framework for evaluating AI policy in K–12 education* (IAEP White Paper No. 2026-01). Center for Innovation in Applied Education Policy (IAEP), San José State University.

# Acknowledgements

The development of this white paper was guided by a simple idea: the future of artificial intelligence in education should be shaped by the people closest to teaching, learning, public policy and educational improvement.

We are grateful to the educators, researchers, policymakers, technology leaders, and community partners who generously shared their expertise throughout the development of this work. Their questions, critiques, and insights strengthened the analysis presented here and helped clarify how responsible AI implementation can better connect policy, pedagogy, assessment, and continuous improvement.

We extend our sincere appreciation to the following colleagues for their thoughtful review and constructive feedback during the preparation of this white paper:

- **Annamarie Francois**, University of California, Los Angeles
- **Michael W. Kirst**, Professor Emeritus, Stanford University; former President, California State Board of Education
- **Wes Kreisel**, Orange County Department of Education
- **Matt Navo**, California Collaborative for Educational Excellence
- **Glen Price**, Glen Price Group; former California Department of Education
- **Mary Vixie Sandy**, California Commission on Teacher Credentialing

Their willingness to engage critically with these ideas significantly strengthened the final version of the white paper. Responsibility for the interpretations, recommendations, and any remaining errors rests solely with the authors.

We also extend our special thanks to **Dr. Lorri Capizzi**, Co-Director of the Center for Innovation in Applied Education Policy (IAEP), San José State University, for her ongoing collaboration, feedback, and encouragement throughout this project, and to **Jean Shiota**, Emerita, Center for Faculty Development, San José State University, for her invaluable assistance in preparing the final PDF version and supporting the production of the accessible publication materials.

# Preface

Artificial intelligence is transforming education at a pace rarely seen in public schooling. In just a few years, educators have moved from asking whether AI belongs in classrooms to determining how it can be used responsibly to support teaching, learning, assessment, accessibility, and student development. California has responded to this challenge by developing one of the nation's most comprehensive bodies of guidance for artificial intelligence in K–12 education. As implementation accelerates, however, a new opportunity emerges: to connect responsible AI governance more explicitly with what decades of educational research tell us about how people learn.

This white paper was developed to contribute to that next stage of the conversation. It builds upon California's existing leadership while proposing a learning-centered framework for evaluating AI policy and practice through the Five Pillars of Practice—Accuracy, Agency, Accessibility, Assessment, and Authenticity. Rather than replacing existing guidance, the paper seeks to extend it by connecting policy, pedagogy, assessment, professional learning, and continuous improvement through a common educational language grounded in the learning sciences.

The intended audience includes leaders across California's educational system, including the California Department of Education, county offices of education, school districts, institutions of higher education, state agencies, professional learning organizations, researchers, educational technology developers, and community partners. Although the paper draws extensively on California's policy context, the ideas presented are intended to contribute to broader national and international conversations about responsible AI-enabled student learning.

The recommendations offered here should be understood as an invitation rather than a prescription. Artificial intelligence will continue to evolve, and educational policy must evolve with it. The Five Pillars of Practice and related approaches described in this paper are intended to support ongoing educator inquiry, collaborative problem solving, and evidence-informed decision making rather than establish fixed rules for implementation. Ultimately, the future of AI in education will depend less on technological capability than on the collective capacity of educators, policymakers, researchers, and communities to design learning environments that preserve human judgment, deepen understanding, and expand opportunities for every learner to flourish.

# Executive Summary

Artificial intelligence has rapidly shifted from an emerging educational technology to a transformative force within K-12 teaching, learning, assessment, and school administration. Within only a few years, educators have moved from asking whether AI belongs in schools to determining how it should be implemented responsibly. State agencies, school districts, and educational organizations have responded by developing guidance that emphasizes data privacy, procurement, acceptable use, academic integrity, professional learning, and equitable access.

Convened in response to Senate Bill 1288 (Becker), the California Department of Education's (CDE) Artificial Intelligence in Education Working Group has helped position California as a national leader through the development of the 2023 *Artificial Intelligence Guidance*, the 2024 *Artificial Intelligence Resource Kit*, and the recommendations of the 2025 California Public Schools Artificial Intelligence Working Group.

These efforts at AI in K-12 education policy formation have expanded the state's attention from awareness and responsible use toward governance, professional learning, data privacy, procurement, accessibility, equity, and instructional implementation. They provide an important foundation for helping educators and educational leaders navigate an increasingly AI-mediated future.

These developments constitute meaningful progress. Yet they also reveal a fundamental limitation shared by much of the current AI policy landscape. Existing guidance is increasingly sophisticated about governing AI technologies but remains comparatively underdeveloped in explaining how AI should support learning itself. Discussions of privacy, procurement, bias, accessibility, and responsible use are essential, but they do not fully answer the educational question that ultimately determines the success or failure of AI implementation:

*What counts as learning in AI-mediated environments?*

As AI becomes more and more embedded within classroom practice, however, California faces a new policy challenge. The central questions are no longer primarily technological; they are educational.

- What counts as genuine learning when AI increasingly participates in thinking, writing, problem solving, assessment, and knowledge production in the lives of K-12 students?
- How should educators distinguish between AI that assists and extends learning and AI that substitutes for and offloads the cognitive work necessary for genuine student development?
- How should schools preserve trustworthy evidence of student growth and understanding of the curriculum while taking advantage of upskilling students who will eventually interact with increasingly capable intelligent AI systems?

These questions cannot be answered through governance frameworks, procurement policies, or acceptable use guidelines alone. They require an explicit educational framework capable of connecting responsible AI implementation to what decades of research tell us about ambitious teaching, deeper learning, classroom assessment, and human development.

California is uniquely positioned to address this challenge because it already possesses one of the nation's most comprehensive systems of continuous educational improvement. Rather than constructing an entirely new AI policy infrastructure, the state can build upon existing systems—including School Accountability Report Cards (SARCs), Local Control and Accountability Plans (LCAPs), the Local Control Funding Formula (LCFF), Learning Management Systems (LMSs), Professional Learning Communities (PLCs), and statewide networks for educator professional learning. The opportunity is not to create additional layers of compliance, but to embed a richer learning architecture within the systems California already uses to improve teaching and learning.

This white paper proposes such an architecture by connecting three complementary components. First, the learning sciences provide the theoretical foundation for understanding how learners develop knowledge, expertise, and agency in AI-mediated environments. Second, the Five Pillars of Practice—Accuracy, Agency, Accessibility, Assessment for Deeper Learning, and Authenticity—provide a learning-centered research-based framework for evaluating whether AI supports meaningful educational experiences (Duckor & Holmberg, 2026). Third, the paper illustrates how AI Review Cycles can provide a practical methodology for transforming educator inquiry into structured evidence that informs instructional improvement, procurement, professional learning, policy refinement, and continuous improvement. Together, these components connect educational theory, classroom practice, and public policy through a common language focused on learning rather than technology alone.

This white paper is intended to support California's next generation of educational leadership as the state's AI policy continues to evolve. While written in response to the rapid emergence of generative artificial intelligence, its recommendations are designed to inform long-term decision making across the California Department of Education, Geographic Leads (Geo Leads), county offices of education, school districts, institutions of higher education, California Subject Matter Projects, professional learning organizations, and other educational partners committed to strengthening teaching, learning, and continuous improvement in the classroom.

Rather than prescribing a single implementation model, the paper offers a learning-centered framework that educational leaders can adapt to their local contexts while contributing to a broader statewide conversation about responsible AI-mediated learning. It argues that the next generation of AI policy should move beyond responsible technology adoption toward responsible AI-enabled student learning.

Building on California's existing leadership, we recommend **six complementary priorities** for strengthening the state's AI policy infrastructure:

- **Integrate the learning sciences into future AI guidance** by explicitly incorporating concepts such as productive struggle, scaffolding, metacognition, developmental progression, transfer, self-regulated learning, and the Zone of Proximal Development as foundations for evaluating AI-supported learning.
- **Develop practical implementation resources** that operationalize the Five Pillars of Practice through district policy review protocols, classroom observation tools, instructional design resources, facilitation guides, student reflection protocols, and criterion-based evaluation instruments.
- **Embed AI within California's existing continuous improvement infrastructure** by integrating learning-centered AI evaluation into SARCs, LCAPs, LCFF reporting, LMS platforms, and Professional Learning Communities rather than treating AI as a separate educational initiative.
- **Continue broad stakeholder engagement** by convening educators, researchers, policymakers, students, families, and technology developers to refine a shared vision for AI-mediated education that remains responsive to technological innovation while grounded in enduring educational purposes.
- **Establish a statewide infrastructure for continuous AI learning** through implementation of the Five Pillars of Practice and AI Review Cycles Evaluation Framework, enabling educator inquiry to generate shared evidence, professional knowledge, and continuous improvement across California's educational system.
- **Invest in sustained research-practice partnerships** among teacher educators, preparation programs, school districts, county offices of education, higher education institutions, nonprofit organizations, industry partners, and state agencies to generate robust evidence regarding AI-supported teaching, learning, assessment, accessibility, and student development.

Collectively, these recommendations position California to move beyond governing artificial intelligence toward cultivating a statewide learning system capable of continually improving how AI supports ambitious teaching, deeper learning, educator professional judgment, and student development. Rather than viewing AI policy as a static set of regulations, this paper envisions it as a dynamic process of organizational learning in which educator expertise, classroom evidence, and continuous inquiry become the foundation for ongoing policy refinement.

The next generation of AI policy will not be defined by the sophistication of its technology guidance alone. It will be defined by the extent to which it strengthens the intellectual, social, relational, and professional conditions that make meaningful learning and human development possible. California now has an opportunity to lead not only in governing artificial intelligence, but also in defining what counts as learning in an age of increasingly capable intelligent systems—and in building the knowledge infrastructure through which educators, policymakers, students, families, and communities can continually improve the ways AI serves ambitious teaching, deeper learning, and the flourishing of every learner.

# I. Introduction

## From Early Adoption and Use Cases of AI in Schools Toward Designing Educative Experiences for Systemic Approaches to Ambitious Teaching and Deeper Learning

Artificial intelligence has entered education with remarkable speed. In only a few years, generative AI systems have evolved from experimental technologies into everyday companions for teachers, students, administrators, and families. Writing assistants, intelligent tutoring systems, adaptive learning platforms, automated feedback systems, multimodal content generators, and increasingly sophisticated agentic AI applications now influence how lessons are designed, assignments are completed, assessments are administered, and student work is evaluated. The pace of technological innovation has challenged educational systems to respond with equal urgency, prompting a wave of guidance documents, procurement frameworks, policy statements, and professional learning initiatives intended to help schools navigate an increasingly complex instructional landscape.

The first generation of educational AI policy understandably concentrated on responsible adoption. Policymakers have been seeking to answer foundational questions about what artificial intelligence is, how it functions, where it appears within educational settings, and what risks accompany its use. Across states, guidance documents tend to emphasize student data privacy, cybersecurity, algorithmic bias, digital citizenship, academic integrity, equitable access, educator professional learning, and acceptable use (Dwyer, 2025). These efforts sought to reduce uncertainty by helping educators demystify AI while establishing safeguards that would support thoughtful experimentation rather than uncritical adoption.

California emerged as one of the nation's leading examples of this early policy work. The California Department of Education's 2023 *Artificial Intelligence Guidance* established an explicitly human-centered orientation toward AI in schools by emphasizing both "learning about AI" and "learning with AI." Rather than promoting technology as an educational end in itself, the guidance framed AI as a tool whose value depends upon human judgment, ethical use, equitable implementation, and strong relationships between educators and students. The companion *AI Resource Kit* released in 2024 reinforced these commitments through practical resources focused on foundational beliefs, safe use, AI literacy, digital citizenship, accessibility, educator professional learning, and community engagement. Throughout these materials, California consistently emphasized that technology cannot replace caring educators and that AI should strengthen, rather than diminish, human connection and instructional expertise.

## The Work of the California Department of Education

The work of the California Department of Education and its educational partners represents an important milestone in educational policy. Rather than approaching AI primarily as a technological innovation, California framed it as an educational challenge requiring thoughtful governance, collaborative leadership, public engagement, and continuous professional learning.

In 2025, the California Public Schools Artificial Intelligence Working Group released its guidance substantially expanded the state's recommendations regarding human-centered AI, AI literacy, equitable access, responsible use, procurement, academic integrity, assessment, and organizational governance. Collectively, these documents position California among the nation's most comprehensive efforts to support responsible AI implementation across transitional kindergarten through grade twelve.

This evolution reflects an important shift in educational policy. Early guidance focused primarily on whether schools should adopt AI and under what conditions. More recent guidance increasingly assumes that AI will become embedded within the instructional core of schooling and therefore asks how educational systems should govern its implementation responsibly.[i] This transition—from adoption to governance—represents a significant advance in educational policy thinking.

Yet it also reveals a new and increasingly important challenge.

As guidance becomes more sophisticated about governing AI technologies, it simultaneously exposes a foundational educational question that remains comparatively underdeveloped. We must now ask:

> *What counts as real learning and actual development for students in an AI-mediated K-12 classroom?*

Throughout this white paper, this question is examined not only by asking what counts as learning, but also for whom, under what instructional conditions, and based on what evidence. These questions recognize that educational equity is not a separate policy objective, but a cross-cutting consideration that informs every dimension of responsible AI practice. Across each of the Five Pillars, educators and policymakers are encouraged to more deeply examine:

*Who benefits from this use of AI in K–12 education? Who may be burdened, excluded, or underserved in ways that open up a new digital divide? Whose knowledge, experiences, and evidence of learning are recognized as authentic achievement? What do disaggregated patterns of evidence reveal about learning opportunities, participation, progress, and outcomes?*

These questions about educational excellence and equity differ fundamentally from earlier concerns regarding adoption, procurement, or legal compliance. It asks educators not simply whether AI tools are safe, equitable, or ethically deployed, but whether they strengthen the developmental processes through which students learn to think, reason, create, communicate, solve problems, and exercise sound judgment.

We have known for decades that artificial intelligence does not merely automate instructional tasks; it redistributes cognitive work (Risko & Gilbert, 2016) among students, teachers, and increasingly capable computational systems (Gerlich, 2025). Questions of authorship, assessment, accessibility, feedback, and academic integrity therefore become inseparable from broader questions concerning human learning itself (Hardwig, 1985).

Current policy guidance frequently references personalization, engagement, feedback, accessibility, and student support. These are unquestionably valuable educational goals. However, they are rarely situated within an explicit theory of learning or development. Concepts that have long guided ambitious teaching—including scaffolding, productive struggle, internalization, metacognition, developmental progression, artificial means and the Zone of Proximal Development—receive comparatively little attention within contemporary AI policy discourse (Duckor & Holmberg, 2026).

*Without these conceptual anchors, educators may find themselves evaluating increasingly sophisticated technologies without a shared language for determining whether those technologies actually contribute to deeper learning and deeper equity.*

The distinction is more than conceptual; it is profoundly consequential for educational policy. Schools have adopted transformative technologies before, including calculators, word processors, internet search engines, learning management systems, and adaptive instructional software. Each altered aspects of classroom practice while leaving the fundamental purposes of schooling largely unchanged (Cuban, 1986, 2001).[ii]

# Challenges and Opportunities in the Evolving Age of AI in Education

Artificial intelligence is different because it increasingly participates in activities once regarded as uniquely human. AI now generates explanations, composes arguments, provides formative feedback, synthesizes information, creates multimedia artifacts, analyzes data, and engages learners in extended dialogue.

But it does more. It offers continuous, "24/7 access" to forms of affective, psychological, and emotional interaction—providing encouragement, reassurance, personalized responses, and conversational engagement at moments when learners encounter uncertainty or struggle. This creates new opportunities for support but also raises deeper questions about how learners develop persistence, independence, judgment, and agency when intelligent systems can increasingly manage not only cognitive tasks but also the emotional experience of learning.

> *As these generative AI capabilities continue to expand and evolve, educational systems must distinguish between technologies that extend human learning and those that substitute for it.*

The tradeoff between interactive and interpassive learning is now a question for educational policymakers because AI challenges long-standing assumptions about where thinking occurs, how understanding develops, and what forms of cognitive engagement remain essential for learning. Systems that provide immediate answers, generate polished products, or remove moments of uncertainty may increase efficiency while simultaneously reducing opportunities for students to engage in the productive struggle through which reasoning, judgment, creativity, and independence develop.

We describe this distinction as the difference between *interactive* and *interpassive* learning (Duckor & Holmberg, 2026). Interactive uses of AI expand opportunities for inquiry, reflection, revision, and intellectual growth; interpassive uses shift essential dimensions of thinking, judgment, and meaning-making away from the learner and onto the technology itself.

More research is needed to determine if "24/7" generative AI assistance is inherently problematic for learning and development. Scaffolding has always been central to effective teaching and learning (Vygotsky, 1978). The critical question is whether AI functions as a scaffold that supports learners as they develop greater capability, independence, and agency, or whether it becomes a substitute that performs the cognitive and emotional work through which learners develop those capacities.

## A Critical Inflection Point: Moving Policy towards Pedagogy

This white paper argues that educational AI policy has therefore reached a critical inflection point. The next generation of policy must move beyond responsible AI adoption toward responsible AI-enabled student learning–beyond broad governance and procurement guidance. Educational policy dedicated to continuous improvement must also articulate how AI should support the development of expertise, preserve authentic intellectual engagement, strengthen ambitious teaching, protect the relational dimensions of teaching and learning that foster trust, belonging, and human development, and generate trustworthy evidence of student learning.

To address this challenge, we introduce the **Five Pillars of Practice (5POP) Framework** as a learning-centered architecture for evaluating AI policy and practice. Organized around the interconnected domains of Accuracy, Agency, Accessibility, Assessment, and Authenticity, the framework extends existing guidance by connecting ethical AI implementation to foundational principles drawn from the learning sciences, classroom assessment, ambitious teaching, and continuous educational improvement.

Although each pillar highlights a distinct dimension of responsible AI practice, they are intended to be interpreted as an integrated framework for evaluating both educational excellence and educational equity. Together, the Five Pillars encourage educators and policymakers to examine not only whether AI enhances learning, but also how opportunities to learn are distributed across California's diverse learners, classrooms, schools, and communities.

Rather than replacing existing policy, an **AI Pedagogical Practice Priorities** approach such as the Five Pillars of Practice framework provides educators, school leaders, districts, and policymakers with a common vocabulary for asking not only whether AI is being used responsibly, but whether its use advances the educational purposes of equity and excellence K-12 public schools exist to serve.

**AI Pedagogical Practice Priorities** are the learning-centered principles that guide how artificial intelligence should be integrated into teaching, learning, assessment, and professional practice. They provide a common language for aligning AI implementation with California's existing statewide priorities for educational equity, instructional improvement, assessment, and continuous improvement while ensuring that AI strengthens rather than competes with local educational goals.

The white paper further argues that this learning-centered framework can be operationalized through mechanisms similar to AI Review Cycles, a methodology proposed by Evaluation by Design. By using a rating interface, this methodology builds on a broader body of work that treats educator professional learning as central to evaluation, evidence generation, and dissemination through structured review of educational products and practices, with examples including online

resources in professional development settings (Ravitz & Hoadley, 2005), technology-supported assessment systems (Ravitz, 2004), and evaluation design improvement (Fetterman & Ravitz, 2018).

Systems such as AI Review Cycles position educators as active investigators of AI-mediated learning environments rather than passive recipients of technological innovation. Through repeated cycles of evidence collection, reflection, and collaborative review, schools and districts can generate cumulative knowledge regarding how AI influences teaching, learning, assessment, accessibility, and student development. Over time, these sorts of on-line, "bottom up" distributed reviews have the potential to create a statewide knowledge infrastructure with potential for informing procurement, professional learning communities, policy refinement, and other areas of continuous improvement in the State's systems of support.

The enduring question of this millennium is not whether educators, schools, and school systems should use artificial intelligence. Rather, it is whether K-12 educational systems can ensure that–in an age of increasingly "intelligent" machines–schools remain places where students develop the human capacities upon which democratic participation, professional expertise, and lifelong learning depend. As Dewey (1938) and others understood the purposes of public schooling, progress happens if and only if our next generation of learners learns how to learn with others with dignity, intellectual challenge, and respect.

The future of AI policy is therefore not principally about artificial intelligence itself. It is about the resources (curricular, pedagogical, experiential) necessary to cultivate and develop our socio-educational intelligence: the collective capacity of educators, students, leaders, policymakers, and communities to design learning environments in which emerging technologies deepen understanding rather than merely accelerate performance.

# II. The Next Policy Challenge for AI in K-12 Education: What Counts as Learning and What Merely Poses as It?

As artificial intelligence becomes an increasingly ordinary feature of teaching, learning, and assessment, educational policy must evolve beyond questions of adoption and governance toward questions of learning itself. Early AI guidance appropriately emphasized issues such as privacy, equity, acceptable use, educator readiness, and technical understanding. These policy problems remain necessary but are no longer sufficient for the foundations of responsible AI use in K-12 setting

The next challenge is educational rather than technological. Schools must now confront a more fundamental question:

> *What counts as learning when AI increasingly participates in thinking, writing, problem solving, feedback, and knowledge production for our K-12 learners?*

This distinction marks an important transition in educational policy. Technology policy traditionally asks whether tools are secure, accessible, transparent, and ethically deployed. Learning policy asks whether instructional experiences promote intellectual growth, disciplinary understanding, adaptive expertise, and long-term development.

Artificial intelligence now occupies both domains simultaneously. Because AI systems can generate explanations, compose essays, solve mathematical problems, summarize research, provide formative feedback, analyze complex data, and engage students in extended dialogue, they no longer function merely as instructional technologies. They increasingly participate in the very cognitive work that schools have historically expected students to perform independently.

This shift creates a fundamentally new policy problem. For generations, educators have relied on observable performances—essays, examinations, projects, presentations, laboratory investigations, and classroom discussions—as reasonable proxies for student learning. Although no assessment has ever measured learning directly, these performances generally reflected students' own intellectual efforts and therefore provided useful evidence from which teachers could make instructional judgments.

Artificial intelligence complicates this relationship. Students may now produce increasingly sophisticated performances with varying degrees of AI assistance, making it more difficult to determine whether the quality of a product reflects the student's development of understanding, the capabilities of the technology, or some combination of both.

## When Polish is not Progress: Dilemmas of the New Age of AI in Education

Consequently, educational policy, like teachers today, can no longer assume that observable performance necessarily reflects students' underlying learning. A polished essay may no longer provide direct evidence of writing proficiency. A sophisticated explanation may not indicate conceptual understanding. Immediate AI-generated feedback may improve final products without strengthening reasoning. A well-designed presentation may reveal technical fluency while obscuring the learner's actual grasp of disciplinary ideas. As AI systems become increasingly capable, educational leaders must distinguish between technologies that support learning and technologies that substitute for it, as well as between evidence of performance and evidence of development.

The distinction is especially important because schools exist to cultivate capabilities that extend beyond the production of correct answers. The purpose of education is not simply to help students complete tasks more efficiently but to develop the intellectual, social, ethical, and creative capacities that enable lifelong learning, informed citizenship, and professional judgment (Darling-Hammond, 1997, 2006; Mehta & Fine, 2019; Tyack & Cuban, 1996).

Artificial intelligence can undoubtedly contribute to these goals. However, it can also obscure them if educational systems evaluate AI primarily in terms of convenience, productivity, or automation rather than its contribution to human development.

## The Role of the Learning Sciences

The learning sciences (NRC, 1999, 2001, 2018) provide decades of scholarship that help illuminate these distinctions. Research in cognitive development, sociocultural theory, formative assessment, self-regulated learning, and ambitious teaching consistently demonstrates that expertise develops through sustained engagement with meaningful problems over time (Lampert, 2001; Sadler, 1989; Shepard, 2000, 2021; Vygotsky, 1978; Zimmerman, 2002). Students deepen their understanding through cycles of scaffolding, productive struggle, feedback, revision, reflection, collaboration, and gradual internalization (Bransford, Brown, & Cocking, 2000).

These developmental processes cannot be reduced to the efficient production of correct answers or polished products now offered "on demand" by AI agents. Rather, they depend upon opportunities for learners to wrestle with uncertainty, justify decisions, explain reasoning, test competing ideas, revise misconceptions, and construct increasingly sophisticated understandings through interaction with others in a thinking curriculum (Alexander, 2008; Resnick & Resnick, 1992).

Remarkably, critical concepts related to learning and development remain largely absent from contemporary AI policy discourse. Terms such as *development*, *knowledge types, metacognition*, and *transfer* seldom appear within state guidance documents despite their longstanding importance within educational research and the learning sciences. Similarly key concepts in learning theory and research such as *Zone of Proximal Development, productive struggle, cognitive offloading, and artificial means* are absent from state guidance documents.

Instead, policy discussions emphasize the benefits of personalization, efficiency, and automation while offering caution about ethics, plagiarism, privacy, and accessibility. These are valuable educational aspirations in the age of AI, but they remain conceptually disconnected from broader theories explaining how learning actually develops in the California learner. As a result, educational systems possess guidance for governing AI technologies *without* a correspondingly sophisticated framework for evaluating the actual *educative* AI-mediated educational experience for students.

## Mind the Gap: Implications for Classroom Practice and Systems Improvement

This conceptual gap has important implications for classroom practice and system improvement.[iii] Without a shared language describing how learning develops, educators may struggle to determine when AI appropriately extends student thinking and when it unintentionally replaces the cognitive work necessary for growth. Likewise, districts may successfully procure legally compliant and technically sophisticated AI systems while lacking common criteria for determining whether those tools strengthen ambitious teaching, improve assessment practices, preserve student agency, or simply automate existing instructional routines. Responsible procurement does not necessarily produce meaningful learning.

The implications extend beyond individual classrooms. If educational policy cannot distinguish between learning and increasingly convincing simulations of learning, then AI- driven accountability systems, assessment practices, curriculum design, and professional learning may miss the mark. A subtle but gradual shift toward AI "solutions" that evaluate student work products (assignments, exams, tasks, and now even exit slips as Anthropic had proposed) will pass for gauging actual learning and genuine development.

Machine work flows and just-in-time classroom and systems solutions will dazzle users rather than shed light on the developmental processes or cognitive shifts these school work products are intended to represent. Efficient scoring, marking and grading quietly prevails over State's interest in actual growth and what makes achievement possible in K-12 settings.

The next generation of AI policy must therefore move beyond responsible technology adoption toward responsible learning design. Meeting this challenge requires a conceptual framework capable of reconnecting AI implementation to the enduring purposes of schooling.

# III. California's Role in Evolving AI Knowledge Networks in K-12 Education

California offers a contemporary example of how educational AI policy is evolving alongside rapidly changing technologies and expanding networks of educational research, professional practice, and public governance.

Since releasing its initial *Artificial Intelligence Guidance* in 2023, the California Department of Education (CDE) has positioned itself as a potential national leader in developing a human-centered approach to artificial intelligence. Rather than framing AI primarily as a technological innovation, California

has situated the future of AI in K-12 education within broader commitments to educational equity, professional judgment, democratic accountability, student safety, and responsible governance.

At the same time, California's work to link responsible AI policy in K-12 settings illustrates a broader international phenomenon. AI policy in education is no longer being developed by isolated companies, research labs or government agencies working independently. Instead, states, provinces, ministries of education, universities, nonprofit organizations, professional associations, technology developers, and international research communities are playing a role. Ideas about AI guidance and policy, more and more draw from evolving knowledge networks that continuously exchange ideas, conceptual frameworks, research findings, implementation strategies, and policy guidance (OECD, 2025; UNESCO, 2024).[iv]

California's current guidance for responsible AI policy draws upon an emerging patchwork of scholarship in the learning sciences, computer and data science, educational technology, and digital citizenship while still searching to position the state to contribute to broader national and international networks of educational policy development.

Eventually, California's guidance may increasingly influence AI policy conversations in K–12 education well beyond its own borders: But only if it succeeds in connecting AI governance to what we know about ambitious teaching, classroom assessment, deeper learning, and human development.

California is well positioned to emerge as a policy innovator in the K–12 space, but that depends on its positionality within a growing global ecosystem of educational knowledge creation informed by advances in artificial intelligence across multiple domains (Kissinger, Schmidt, & Mundie, 2024).[v]

## Where We Currently Stand in the Policy Cycle

The next stage of policy development requires more than stronger governance. It requires a coherent educational architecture capable of connecting policy, pedagogy, assessment, and continuous improvement.

The California Department of Education (CDE 2023) guidance established several principles that continue to provide an important foundation for policy innovation. **First**, it distinguished between learning about AI and learning with AI, recognizing that students require both technical understanding and opportunities to use AI responsibly within authentic learning environments. **Second**, it emphasized that human relationships remain central to education and that technology should strengthen rather than replace educator expertise. **Third**, it introduced AI literacy as a statewide priority encompassing digital citizenship, ethical reasoning, bias awareness, media literacy, accessibility, and responsible participation in an increasingly AI-mediated society.

These principles have been reinforced through the 2024 *Artificial Intelligence Resource Kit*, which translated broad policy recommendations into practical resources for districts, educators, and

professional learning communities. In particular, the work of the California Public Schools Artificial Intelligence Working Group has further broadened the scope of state guidance to include detailed recommendations regarding AI literacy progressions, equitable access, academic integrity, responsible use, procurement, accessibility, data privacy, governance, educator professional learning, and implementation.

This progression represents genuine policy maturation. Compared with guidance emerging from many other states and educational jurisdictions, California's framework demonstrates sustained attention to implementation, educator capacity building, stakeholder engagement, procurement processes, and equity considerations.[vi] It recognizes that AI is not simply another instructional technology but a systemic educational issue requiring coordinated responses across curriculum, assessment, professional learning, leadership, governance, and continuous improvement.

California continues to move beyond asking whether schools should adopt AI and instead begin addressing how AI-mediated educational systems might govern increasingly sophisticated technologies in ways that remain consistent with public values, enduring pedagogical goals connected to deeper learning, and democratic educational purposes that emphasize academic excellence and equity.

## Where Policy Meets Pedagogy: Guiding Questions for California's AI K-12 Learning Ecosystem

California's next challenge is therefore not simply to produce stronger AI guidance, but to ensure that responsible AI implementation becomes integrated within the state's existing educational improvement infrastructure rather than emerging as a separate initiative. Achieving this goal requires coherent leadership across California's educational system.

State policymakers, educational agencies, county offices of education, district leaders, school administrators, teacher educators, and classroom practitioners each occupy distinct positions within the state's educational ecosystem, yet all share responsibility for advancing ambitious teaching, deeper learning, and continuous improvement in an era of artificial intelligence. Table 1 illustrates how system coherence depends upon shared but role-specific guiding questions across California's educational system.

Table 1

*Guiding Questions that Various Educational Partners in the State of California Should be Asking About AI in K-12 Education*

| Educational Partner | Guiding Leadership Question |
|---|---|
| **State Policymakers and Legislative Staff** | What statewide conditions will enable artificial intelligence to strengthen teaching, learning, assessment, and human development across California's diverse educational communities? |
| **California Department of Education and State Agencies** | How can statewide guidance strengthen coherent implementation while respecting California's tradition of local educational decision making? |
| **California's Statewide System of Support** (County Offices of Education, CCEE, Geographic Leads, Statewide Leads, California Subject Matter Projects, technical assistance providers, and related partners) | How can California's existing implementation infrastructure help educational systems integrate AI into ongoing improvement efforts rather than treating it as a separate initiative? |
| **District Leadership Teams** | How can AI strengthen our district's existing instructional vision rather than compete with it? |
| **School Leaders** | How can our school create the conditions under which AI supports ambitious teaching, deeper learning, and professional collaboration? |
| **Teacher Educators** | How should educator preparation programs cultivate the professional judgment needed for AI-mediated teaching and assessment? |
| **Classroom Educators** | How can AI extend students' opportunities to think, learn, create, and demonstrate understanding without replacing the cognitive work that promotes learning and development? |

Because California's educational system distributes leadership responsibilities across multiple organizations and levels of practice, each educational partner must grapple with a common challenge through questions appropriate to its role. Appendix A "Guiding Questions for California's K-12 Educational Systems" summarizes the guiding questions that each stakeholder should be asking now, as California prepares for the next generation of AI-enabled teaching and learning in K-12 education.

Yet California's current AI policy evolution is not occurring in a vacuum.[vii] Nor has research kept pace with corporate offerings. A startling gap is growing:

> *The preliminary architecture of AI guidance and governance has advanced more rapidly than the educational theories from the learning sciences needed to explain how AI contributes, if at all, to the genuine development of learners in K-12 classrooms.*

This imbalance between policy and pedagogy should not be interpreted as a weakness of California's work.[viii] Rather, it reflects the historical stage of AI policy development itself. The first generation of policy necessarily concentrated on responsible adoption, legal compliance, ethical safeguards, and technical implementation because these questions demanded immediate attention. As AI becomes increasingly embedded within the everyday work of teaching and learning, however, the policy conversation is shifting.

California therefore serves not only as a model of responsible AI governance but also as a potentially important case study for understanding the next stage of educational policy development that focuses on what we know about how and why people learn (Bruner, 1960, 1990; Dewey, 1933; Vygotsky, 1978). The state's experience suggests that governance frameworks, while necessary, are insufficient without complementary frameworks that explain how AI should support learning, teaching, and the development of expertise among K–12 learners.

The challenge, then, is not simply to strengthen AI governance but to provide educators and policymakers with a learning-centered framework capable of evaluating whether AI advances the kinds of ambitious teaching and deeper learning California seeks to promote. The next section introduces the Five Pillars of Practice as one such conceptual architecture for reconnecting AI policy with the enduring purposes of schooling.

# IV. From Governance and Guardrails Toward Learning and Development: The Five Pillars of Practice Framework for Ethical AI Use in K–12 Classrooms

If the first generation of AI policy focused appropriately on governance, compliance, and responsible adoption, the next generation must focus on learning and development.

The central challenge is no longer simply determining whether AI systems are secure, equitable, or legally compliant. Rather, educators and policymakers must determine whether AI strengthens the developmental conditions under which deeper learning occurs. Meeting this challenge requires a conceptual framework that reconnects AI implementation to the enduring purposes of schooling (Figure 1).

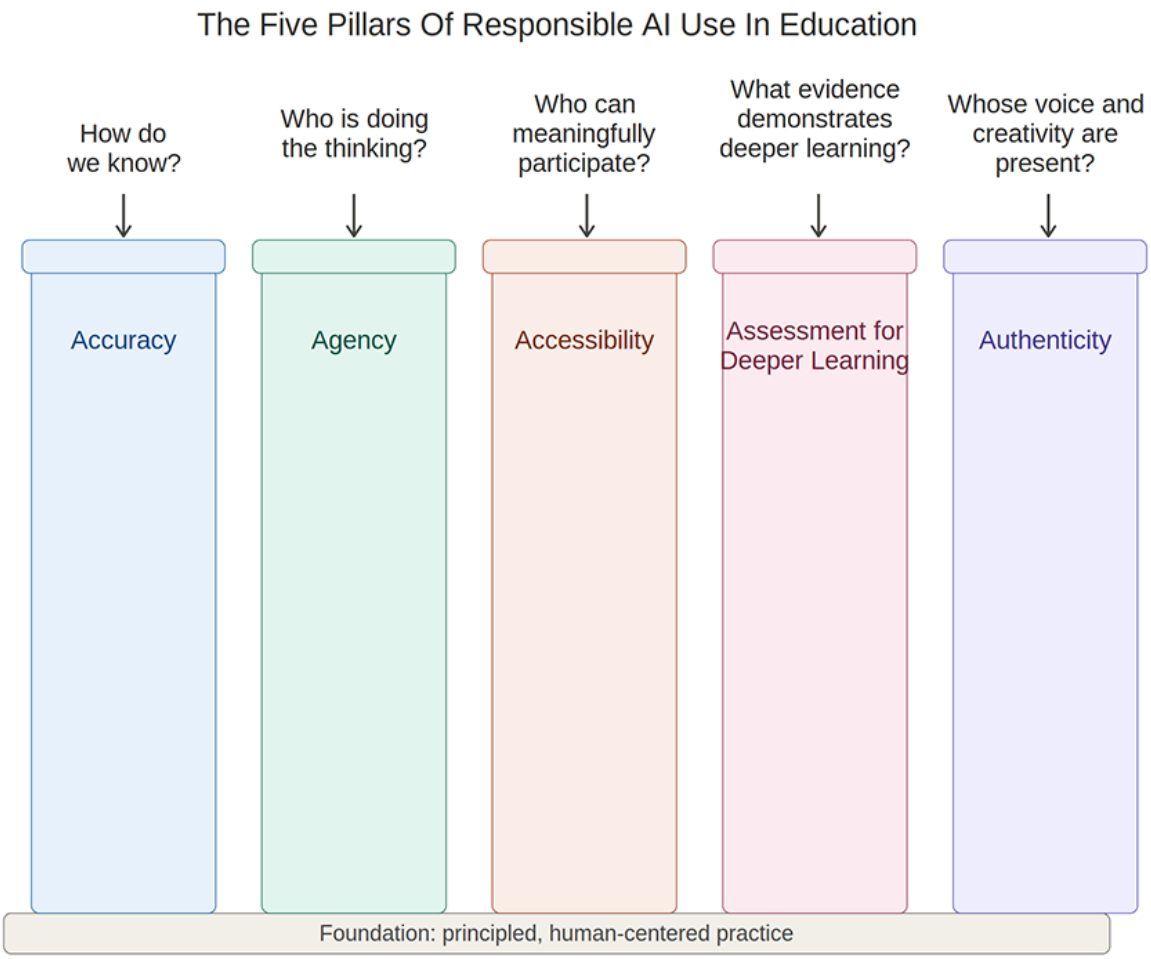

*Figure 1.* The Five Pillars of Practice for Responsible Human-Centered AI Use in K–12 Education. Adapted from Duckor, B., & Holmberg, C. (2026). *AI and Deeper Learning: Promises, Paradoxes, and Evolving Practices*. Harvard Education Press.

The Five Pillars of Practice—Accuracy, Agency, Accessibility, Assessment, and Authenticity—were developed to address precisely this need (Duckor & Holmberg, 2026). Rather than replacing existing guidance, the framework complements California's evolving AI policies by providing a learning-centered architecture through which educators, school leaders, districts, and policymakers can evaluate AI-mediated teaching and learning.

It is important to note that the Five Pillars of Practice framework is not intended to function as a separate educational initiative or as an additional layer of educational reform. Rather, they provide a common language through which California's existing educational infrastructure can examine how artificial intelligence is influencing teaching, learning, assessment, and student development. Because AI will continue to evolve more rapidly than traditional policy cycles, the challenge is not simply adopting new guidance but strengthening the capacity of educational systems to learn, adapt, and improve while remaining coherent around enduring commitments to ambitious teaching, deeper learning, educational equity, and human development.

The Five Pillars of Practice framework shifts attention from what AI systems are capable of doing to what learners are actually learning and developing as they use these tools and tech systems. Each pillar provides a common vocabulary for examining whether AI supports ambitious teaching, deeper learning, and continuous educational improvement. Moreover, the Five Pillars invite educators and policymakers to evaluate AI-mediated learning through a common set of questions about educational quality, equity, and evidence of continuous improvement.

## The Pillar of Accuracy: How Do We Know?

Accuracy focuses on the quality, credibility, and trustworthiness of information. AI systems increasingly generate text, images, recommendations, and feedback that appear authoritative but may contain inaccuracies, hallucinations, incomplete reasoning, or embedded bias.

Educational policy should therefore promote transparency, verification, disciplinary rigor, and critical evaluation rather than unquestioning acceptance of AI outputs. Triangulating multiple sources of evidence remains essential for developing *epistemic responsibility* among both educators and students.

This first pillar of responsible practice, **Accuracy**, extends well beyond factual correctness to encompass epistemic responsibility within and across subject matter disciplines. In AI-mediated K-12 environments, students must learn to verify claims, triangulate evidence, recognize uncertainty, identify hallucinations and bias, evaluate sources, and justify conclusions within disciplinary contexts.

Accuracy therefore concerns the development of informed judgment rather than passive acceptance of computational outputs. AI should strengthen students' capacity to ask, "How do we know?" before accepting what appears plausible or authoritative.

**Core questions** that support the alignment of policy to pedagogy with the **Accuracy Pillar of Practice** include:

- How is information verified?
- What evidence supports particular claims?
- What limitations, uncertainty, or bias should be acknowledged?
- How are learners developing habits of critical evaluation and epistemic responsibility?
- Are AI contributions represented honestly and transparently?

## The Pillar of Agency: Who Is Doing the Thinking?

Agency emphasizes human judgment, intellectual labor, and meaningful decision-making. AI should support rather than replace *productive struggle*, reflection, creativity, and professional expertise. Evaluating agency requires examining where cognitive work resides within AI-mediated environments and whether instructional designs preserve opportunities for students and educators to exercise judgment.

Educational policy should therefore promote instructional designs that preserve learner autonomy, encourage active engagement, and ensure that cognitive responsibility remains centered on students and educators rather than being delegated to AI systems. Creating opportunities for learners to exercise judgment, make decisions, and reflect on their thinking remains essential for developing agency in AI-mediated learning environments.

This second pillar of responsible practice, **Agency**, focuses attention on intellectual responsibility and cognitive ownership. Its central question is deceptively simple: Who is doing the thinking? To answer this question, the distinction between *interactive learning* and *interpassive learning* becomes especially salient. When AI assumes increasing portions of the intellectual work, educators must continually ask whether learners remain active participants in constructing understanding or are gradually delegating essential cognitive and affective processes to increasingly capable machines.

**Core questions** that support the alignment of policy to pedagogy with the **Agency Pillar of Practice**

- Who is directing the work and making important decisions?
- What cognitive work remains essential for learners?
- How does AI support rather than replace learning?
- How are students exercising metacognitive awareness and self-regulation?
- How is teacher professional judgment strengthened rather than diminished?

## The Pillar of Accessibility: Who Can Meaningfully Participate?

Accessibility extends well beyond access to technology. It encompasses equitable participation, appropriate instructional supports, human scaffolding, and opportunities for every learner to engage meaningfully in rigorous academic work. The objective is authentic inclusion rather than simply providing technological accommodations and *simulations of access*. AI should broaden participation while preserving high expectations and meaningful opportunities to learn.

Educational policy should therefore promote instructional environments that reduce barriers to participation while maintaining ambitious learning goals for all students. AI-supported accommodations should expand opportunities to engage in challenging academic work without diminishing intellectual expectations or replacing the essential work of learning. Human judgment remains indispensable for ensuring that accessibility serves development rather than dependency.

This third pillar of responsible practice, **Accessibility**, reframes equity beyond simple access to digital technologies. Meaningful accessibility in K–12 learning environments involves expanding opportunities for every learner to participate in intellectually demanding work while preserving appropriate expectations for growth and achievement.

**Core questions** that support the alignment of policy to pedagogy with the **Accessibility Pillar of Practice** include:

- Who benefits from AI-supported learning?
- What barriers to participation remain?
- Are supports developmentally appropriate and instructionally meaningful?
- Does AI expand participation without reducing intellectual challenge?
- Does increased access result in deeper understanding rather than superficial completion?

AI can reduce barriers through translation, adaptive supports, multimodal representation, and assistive technologies, but accessibility should never become the simulation of participation without corresponding development. Authentic accessibility strengthens learners' capacity to engage more fully in ambitious learning rather than merely making tasks easier to complete.

## The Pillar of Assessment for Deeper Learning: What Evidence Invites and Demonstrates Learning?

Assessment focuses on making student thinking visible and generating meaningful evidence of learning. AI should strengthen formative assessment practices by supporting feedback, revision, reflection, and instructional decision-making rather than merely automating grading or producing polished performance. Evaluation therefore centers on whether AI enhances educators' ability to understand student thinking, adapt instruction, and support continued growth over time.

Educational policy should therefore promote assessment systems that prioritize evidence of student reasoning, conceptual understanding, and developmental growth over efficiency and polished production. AI should strengthen human-centered formative assessment by enriching both fast and slow feedback processes while supporting educator professional judgment and meaningful dialogue about learning (Duckor & Holmberg, 2023).

Used uncritically, AI-assisted assessment and grading tools may shortcut productive struggle and obscure valid evidence of student learning. Rather than replacing careful interpretation of students' developing knowledge and skills, AI should enhance educators' capacity to make learning visible over time. Meaningful assessment remains grounded in the ongoing interactions among students, teachers, instructional tasks, and evidence collected throughout the learning process.

This fourth pillar of responsible practice, **Assessment for Deeper Learning**, reconnects AI policy to formative assessment and continuous instructional improvement in K-12 classrooms. AI-generated feedback possesses educational value only when it supports reflection, revision, dialogue, and increasingly sophisticated understanding. Assessment therefore focuses on making student thinking visible rather than simply evaluating final products.

The framework distinguishes assessment for deeper learning from what we have termed *faux-mative assessment*—feedback that appears formative because it is immediate or personalized but ultimately fails to promote meaningful conceptual growth or instructional decision-making. The educational value of AI lies not in accelerating grading, generating more dashboards, or speeding up scoring of student work products, but in strengthening the quality of evidence available to students and teachers as they learn to learn and make learning processes and skills visible during instruction.

**Core questions** that support the alignment of policy to pedagogy with the **Assessment for Deeper Learning Pillar of Practice** include:

- How is student thinking made visible?
- What evidence informs instructional decisions?
- How does feedback promote revision and continued learning?
- How are students engaged in reflection and self-assessment?
- Does AI strengthen assessment for deeper learning rather than simply increasing efficiency?

## The Pillar of Authenticity: Whose Voice and Creativity Are Present?

Authenticity focuses on student voice, intellectual ownership, provenance, and integrity. AI should amplify rather than obscure human creativity while maintaining transparency about how work is produced. Evaluating authenticity requires examining learning processes as well as final products, ensuring that evidence of reasoning, originality, and intellectual contribution remains visible throughout AI-mediated learning.

Educational policy should therefore promote instructional practices that preserve authentic evidence of student thinking, creativity, and authorship while encouraging transparent and responsible use of AI. Learners should be supported in documenting how ideas develop, acknowledging appropriate AI assistance, and demonstrating intellectual ownership throughout the learning process. Authenticity depends not on eliminating AI from learning but on ensuring that human judgment, creativity, and voice remain visible within AI-mediated educational environments.

This fifth pillar of responsible practice, **Authenticity**, addresses perhaps the defining challenge of visible and invisible AI-mediated education. In K–12 settings, the goal is to cultivate visible evidence of human intellectual work among novices who are learning to master new concepts, skills, and dispositions. Rather than relying exclusively on plagiarism detection technologies or punitive academic integrity policies, Authenticity emphasizes *provenance*, transparency, authorship, intellectual ownership, *creative labor*, and documentation of learning processes.

The central concern shifts from identifying misconduct to preserving meaningful evidence of human reasoning, creativity, revision, and voice throughout AI-supported learning. Authenticity thus becomes a developmental concept concerned with cultivating integrity and intellectual contribution rather than simply enforcing compliance.

**Core questions** that support the alignment of policy to pedagogy with the **Authenticity Pillar of Practice** include:

- Whose ideas and perspectives are represented?
- Who contributed to the work and in what ways?
- How are sources, influences, and AI assistance acknowledged?
- How is provenance documented throughout the learning process?
- Where is the boundary between creative remixing, collaboration, and unattributed borrowing?

Viewed together, the Five Pillars of Practice reposition AI policy within a broader theory of ambitious teaching, 21st century skills, and deeper learning. These pillars of AI practice in K-12 schooling recognize that responsible AI use cannot be separated from responsible instructional design and that technological innovation should always remain subordinate to educational purpose.

The Five Pillars of Practice framework demonstrates that responsible AI implementation cannot be separated from thoughtful instructional design and educative experiences in the curriculum. Policies establish expectations, but classroom practice determines whether those expectations ultimately improve learning. By embedding the Five Pillars of Practice within California's existing systems of accountability, professional learning, procurement, and continuous improvement, educational leaders can move beyond responsible AI adoption toward responsible AI-enabled learning.

The Five Pillars of Practice framework establishes a conceptual bridge between policy and pedagogy. Each pillar provides educators and policymakers with a common language for examining how decisions about governance, procurement, professional learning, assessment, and classroom practice collectively shape students' opportunities to think deeply, exercise judgment, construct knowledge, and develop expertise over time. In doing so, the framework extends California's existing guidance from responsible AI adoption toward responsible AI-enabled learning.

*The question now is: how can such a learning-centered framework for AI in education be embedded within California's existing systems of accountability, continuous improvement, professional learning, and existing educational governance.*

The following section explores how the Five Pillars of Practice can be operationalized through the state's existing educational improvement infrastructure–driven by end-users who encounter these new generative tools and solutions–rather than requiring the creation of entirely new policy systems.

# V. Embedding the Five Pillars of Practice into California's System of Educational Improvement

California possesses an unusual advantage in advancing the next generation of educational AI policy. Unlike many states that would need to construct entirely new accountability structures, California already operates one of the nation's most comprehensive systems of continuous educational improvement. School Accountability Report Cards (SARCs), Local Control and Accountability Plans (LCAPs), the Local Control Funding Formula (LCFF), statewide curriculum frameworks, professional learning systems, and local governance structures collectively provide an existing infrastructure capable of supporting learning-centered AI implementation.

> *The AI in K-12 education policy challenge, therefore, is not to create new systems of accountability but to embed richer pedagogical architectures within those already in place.*

## Centering the Role of System Coherence

The Five Pillars of Practice are not intended to function as a separate educational initiative. Rather, they are designed to provide a common language through which California's existing educational infrastructure can engage in coherent conversations about the educational implications of artificial intelligence. Because AI will continue to evolve more rapidly than formal policy cycles, the challenge is not simply adopting new guidance but strengthening the capacity of educational systems to learn, adapt, and improve while remaining anchored to enduring commitments to ambitious teaching, deeper learning, educational equity, and human development.

California's existing Statewide System of Support—including the California Department of Education, the California Collaborative for Educational Excellence, county offices of education, Geographic Leads, Statewide Leads, California Subject Matter Projects, institutions of higher education, professional learning organizations, and local educational agencies—already provides a robust infrastructure for continuous improvement. The opportunity is not to create parallel AI systems, but to integrate AI-related inquiry into these existing structures so that implementation remains coherent with long-standing educational priorities.

Because each educational partner occupies a distinct role within this system, the questions that guide responsible AI implementation will necessarily differ across audiences. Legislators, state agencies, county offices, district leaders, school administrators, teacher educators, and classroom practitioners all contribute to the same statewide learning ecosystem, but each makes decisions at different levels of the educational system. Appendix A illustrates how the Five Pillars may be translated into guiding questions tailored to these different roles while maintaining a shared commitment to continuous improvement.

The Five Pillars of Practice provide a framework for accomplishing this integration. Rather than functioning as an additional layer of compliance, they offer a common language through which educators, school leaders, districts, county offices of education, and state agencies can evaluate whether AI is advancing ambitious teaching and deeper learning. By aligning existing improvement processes with common questions about Accuracy, Agency, Accessibility, Assessment, and Authenticity, California can move beyond regulating AI technologies toward strengthening the educational purposes those technologies are intended to serve.

## Focusing on Existing Policy Levers

One immediate opportunity exists through **School Accountability Report Cards (SARCs)**. Traditionally, SARCs communicate information about school performance, staffing, instructional programs, and student outcomes to parents and community members. As AI becomes increasingly embedded within teaching and learning, SARCs could also become an important mechanism for public transparency regarding AI implementation. Rather than simply reporting technology availability, districts could describe how AI supports curriculum, assessment, educator professional

learning, and student learning while documenting safeguards related to authorship, data privacy, accuracy, and academic integrity. Reporting could also identify the major AI platforms and vendors used across schools, helping communities better understand the technologies shaping students' educational experiences. In this way, the SARC becomes not merely a compliance document but a public accountability mechanism for responsible AI implementation.

**Local Control and Accountability Plans (LCAPs)** provide a second opportunity for integration. California's continuous improvement model already requires districts to establish measurable goals, allocate resources, engage educational partners, and evaluate progress across multiple state priorities. Rather than treating AI as an independent technology initiative, districts could explicitly articulate how AI supports goals related to student achievement, equity, multilingual education, engagement, college and career readiness, and educator professional learning. Embedding AI within existing LCAP planning would reinforce that artificial intelligence is not a separate reform effort but a cross-cutting condition influencing curriculum, instruction, assessment, and organizational learning. Such integration would also encourage districts to move beyond one-time technology adoption toward sustained investments in instructional improvement and educator capacity.

The **Local Control Funding Formula (LCFF)** offers an equally important policy lever. AI implementation often focuses public attention on purchasing software licenses, expanding digital infrastructure, or acquiring new technological capabilities. Yet California's experience consistently demonstrates that meaningful educational improvement depends far more upon educator expertise than technological sophistication. Greater transparency regarding expenditures on AI infrastructure, educator professional learning, instructional coaching, curriculum development, and implementation support would reinforce the principle that sustainable AI integration requires sustained investment in human capacity. At the same time, LCFF reporting can help address emerging forms of digital inequality by documenting disparities in access to high-quality AI tools, reliable broadband, local infrastructure, and family resources. AI equity is increasingly shaped not only by school funding but also by geography, household resources, and opportunities to learn with emerging technologies both inside and outside of school.

Perhaps the most underutilized policy lever lies within **Learning Management Systems (LMSs)**. Increasingly, LMS platforms document drafts, revisions, collaborative interactions, formative feedback, AI-supported activities, and student reflections throughout the learning process. Unlike traditional assessments that often capture only final products, LMS environments make it possible to examine learning as it unfolds over time. Properly designed, these systems can generate evidence regarding revision cycles, metacognitive reflection, patterns of AI use, collaborative learning, and the development of disciplinary understanding. Rather than evaluating only completed assignments, educators can begin examining the processes through which understanding develops. AI-supported learning analytics and process-oriented indicators may eventually provide richer evidence of deeper learning and twenty-first century competencies than traditional product-based assessments alone.

California can further strengthen this infrastructure by expanding the conceptual vocabulary embedded within future guidance documents. Existing policy appropriately emphasizes ethics, governance, responsible use, procurement, accessibility, privacy, and professional learning. The next generation of guidance should explicitly incorporate concepts from the learning sciences—including productive struggle, scaffolding, developmental progression, metacognition, transfer, internalization, and ambitious teaching. These concepts would provide educators with a richer language for making principled instructional decisions about AI while clarifying the educational purposes that governance structures are ultimately designed to protect.

## Working Within the System

These emerging recommendations illustrate that California does not require an entirely new AI policy infrastructure. Rather, it requires a more explicit learning architecture capable of connecting governance to pedagogy, pedagogy to assessment, and assessment to continuous improvement. The challenge, therefore, is not simply writing better AI policies, but equipping educators with practical frameworks that translate broad policy principles into everyday instructional decision making and focused approaches to deeper learning.

Traditional AI governance establishes the conditions for responsible adoption, but governance alone cannot determine educational value. The Five Pillars of Practice (5POP) framework provides a learning-centered lens for evaluating AI-mediated environments, while AI Review Cycles generate practice-based evidence through systematic educator inquiry. Together, these complementary frameworks help educational leaders move beyond compliance toward evidence-informed implementation. Over time, the knowledge generated through repeated review cycles can inform professional learning, procurement decisions, policy refinement, and continuous improvement.

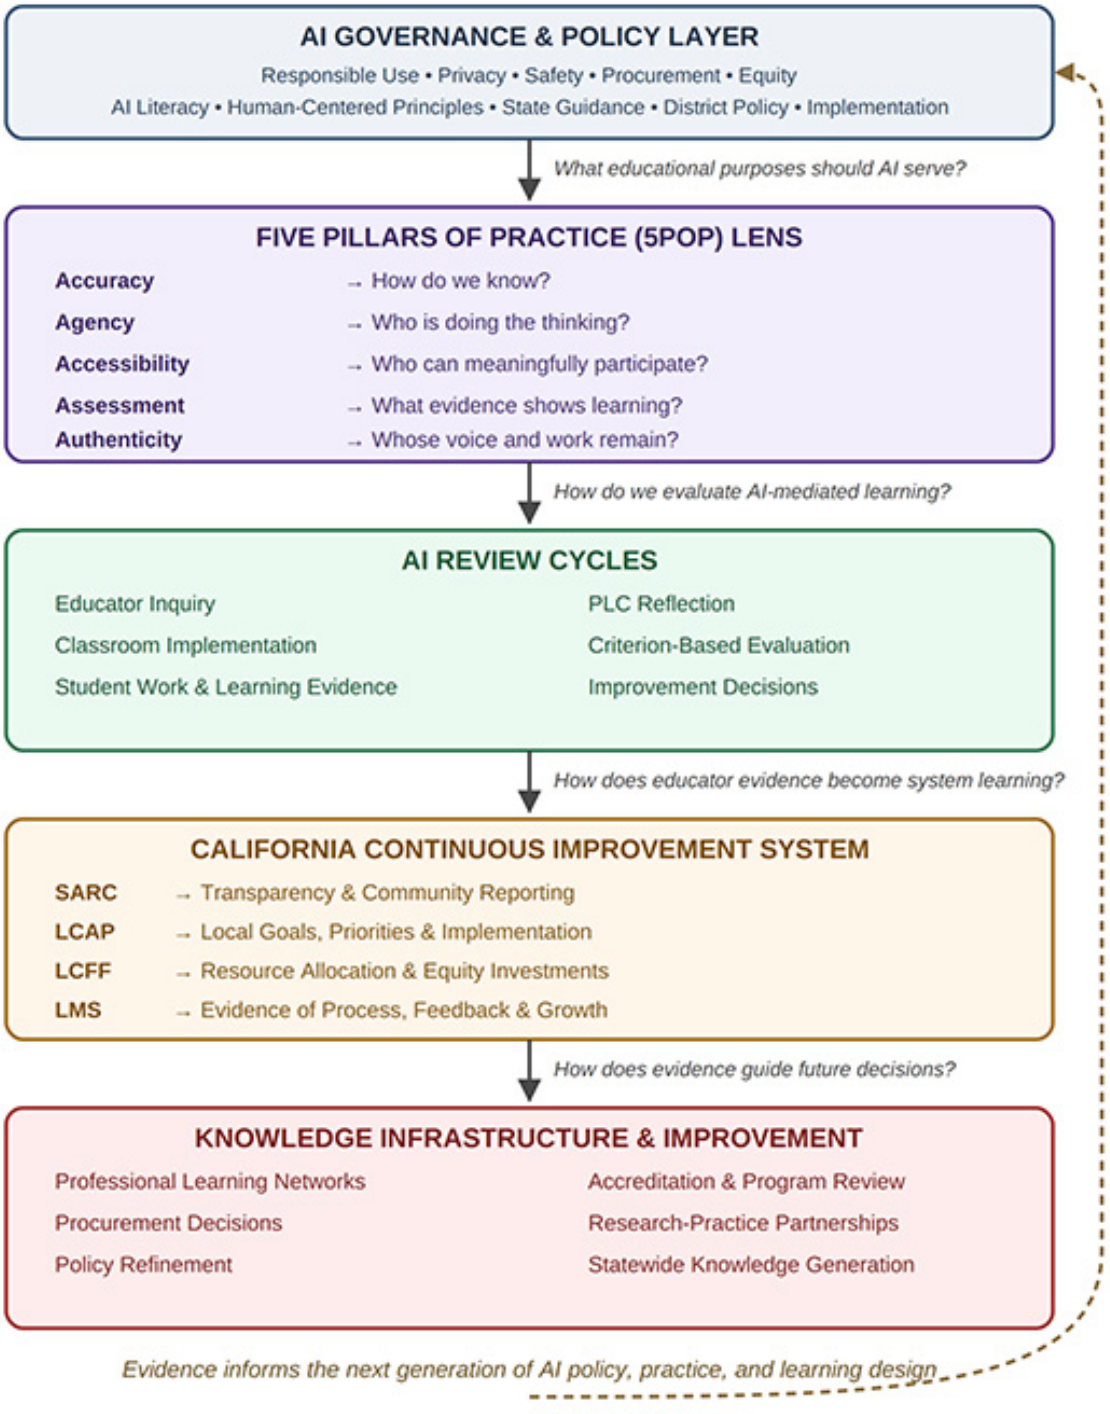


*Figure 2.* Integrating AI Governance, the 5POP/AI Review Cycles Evaluation Framework, and California's Continuous Improvement System (Source: IAEP Center, 2026).

Responsible AI governance establishes the conditions for safe and ethical implementation, but educational improvement requires evidence about learning. The Five Pillars of Practice (5POP) framework provides criteria for evaluating AI-mediated environments, while AI Review Cycles transform educator judgment into evidence as we discuss in section VI.

By connecting this evidence to existing California improvement structures—including SARCs, LCAP, LCFF, and LMS platforms—the system creates a statewide knowledge infrastructure capable of informing professional learning, resource allocation, policy refinement, and continuous improvement (Figure 2).

The Five Pillars of Practice framework provides that research-based, connective framework by translating broad commitments to responsible AI into shared questions about the quality of teaching, learning, evidence, and human development.

When paired with structured AI Review Cycles methodology, the Five Pillars of Practice framework creates a mechanism for transforming classroom experience into organizational knowledge. Emerging knowledge networks centered on pedagogical principles of deeper learning allow educators to make professional judgments, to share student learning evidence, and to garner implementation insights to inform decisions across classrooms, schools, districts, and the broader educational system.

At the systems level, this approach illustrates how AI policy can evolve from a collection of safeguards into a statewide system of educational learning. As districts, schools, and educators continually evaluate AI through a common learning-centered framework, their collective judgments become a source of organizational knowledge that informs procurement, professional learning, accreditation, policy refinement, and future instructional design. California's existing continuous improvement infrastructure is therefore capable not only of governing AI but also of learning from its implementation.

When embedded within California's existing improvement infrastructure, the Five Pillars of Practice framework transforms AI policy from a static set of recommendations into a dynamic process of evidence-informed educational improvement.

# VI. Operationalizing the Five Pillars of Practice Through AI Review Cycles

As California moves from AI guidance toward widespread implementation, educators and educational leaders face an increasingly practical question: How should schools evaluate the growing number of AI tools, systems, and machine learning applications entering classrooms? While current procurement guidance appropriately emphasizes privacy, security, legal compliance, accessibility, and

technical functionality, these considerations alone cannot determine whether a particular AI application contributes to meaningful learning.

Traditional AI governance establishes the conditions for responsible adoption, but governance alone cannot determine educational value. The Five Pillars of Practice (5POP) framework provides a learning-centered lens for evaluating AI-mediated environments, while AI Review Cycles generate practice-based evidence through educator inquiry. Over time, accumulated evidence informs professional learning, procurement, policy refinement, and continuous improvement.

The challenge is no longer simply selecting compliant technologies. Rather, it is determining whether AI tools improve teaching, support student development, and strengthen professional judgment. Educational systems require evaluation approaches that extend beyond vendor claims, product demonstrations, or isolated pilot experiences. They need systematic methods for examining instructional usefulness, identifying unintended consequences, and generating evidence that can inform continuous improvement (See Figure 3).

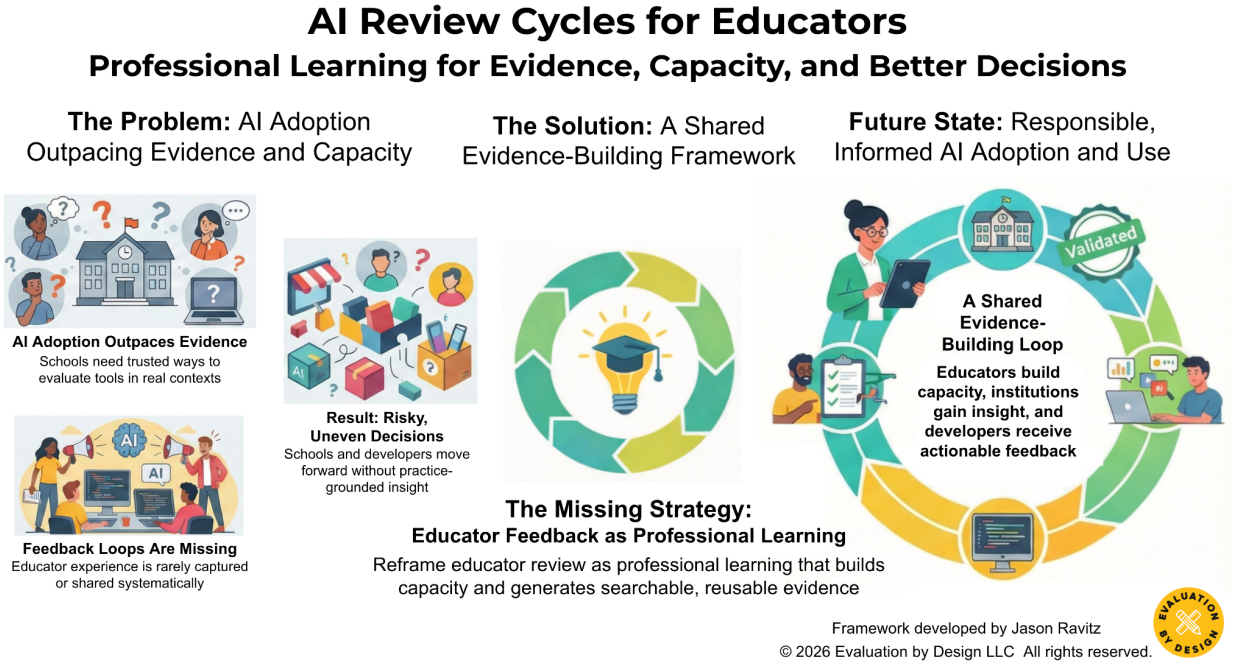


*Figure 3.* AI Review Cycles for Educators: A Framework from Evaluation by Design **(**Source: Evaluation by Design, 2026).

Emerging work on **AI Review Cycles for Educators** (Ravitz, 2026) offers a promising model for addressing this challenge, building on a research-based methodology developed by Ravitz and now being advanced through Evaluation by Design as a methodology that treats educator professional learning as central to evaluation, evidence generation, and dissemination. This methodology has been applied across several educational technology initiatives, including earlier work on technology-

supported assessment systems, standards-based project-based learning units, and structured review of online resources in professional development settings (Ravitz, 2004; Ravitz et al., 2004; Ravitz & Hoadley, 2005).

Rather than treating AI evaluation as a one-time procurement activity, AI Review Cycles position educator review as an ongoing process of professional learning and evidence generation. Educators can engage in structured cycles of reviewing AI tools within authentic instructional contexts, documenting both quantitative ratings and qualitative observations regarding usability, instructional value, equity, accessibility, and classroom implementation. Review findings are then synthesized into discussion protocols, dashboards, and decision-support materials that inform schools, districts, developers, and policymakers. This approach supports current technology infrastructure, including structured review forms, dashboards, discussion protocols, and searchable evidence systems.

The research-based Five Pillars of Practice approach to responsible AI use in K-12 education, coupled with AI Review Cycles, aims to support pedagogically-focused knowledge networks and represents an important conceptual shift. Instead of asking whether an AI product is simply "approved" or "effective," educators become active investigators of how AI functions under varying instructional conditions, for different learners, and across diverse educational settings.

The resulting evidence becomes cumulative and searchable, allowing patterns from classroom implementation to be identified, compared, and refined over time. Rather than leaving valuable educator experiences disconnected across individual classrooms or districts, AI Review Cycles create a shared, searchable and exchangeable knowledge base that supports continuous organizational learning and evidence-informed improvement.

Together, these approaches establish a professional learning infrastructure allowing educators to actively examine, vet among peers, and share with one another how AI tools are moving deeper learning forward. Combined, these two approaches move the evaluation of "what works" beyond technical performance toward the continuous assessment of instructional quality, evidence of learning, and the rich discussions of the developmental conditions necessary for ambitious teaching and deeper learning led by, for, and with California teachers.

## Professional Learning Communities in the Age of AI

In the age of AI, the notion of professional learning communities is also evolving. Linking AI policy implementation to *pedagogical capital* and capacity building requires educators to strengthen their own AI literacy through structured inquiry into classroom practice with other communities of adult learners. Reviewing AI becomes an opportunity to deepen understanding of pedagogy, assessment, accessibility, and instructional design while simultaneously generating evidence that benefits classrooms, schools, districts, and the broader educational community.

This occurs because meaningful evaluation requires educators to move beyond determining whether an AI tool works and toward examining how, when, why, and for whom it supports learning. As educators apply the Five Pillars of Practice, they engage in the same forms of professional reasoning central to effective teaching (Grossman et al., 2009; Lave & Wenger, 1991; Schön, 1983; Shulman, 1987): analyzing evidence of student thinking, evaluating the quality of feedback, identifying barriers to participation, examining where intellectual work resides, and determining whether learning experiences preserve student voice and agency. The review process therefore becomes a form of distributed, collaborative inquiry in which educators refine their understanding of both AI and the conditions that support deeper learning.

In many respects, approaches like those of AI Review Cycles can be understood as the next evolution of Professional Learning Communities (PLCs) for the age of artificial intelligence. For more than two decades, PLCs have encouraged educators to work collaboratively by examining student work, analyzing instructional practices, identifying challenges, and collectively refining teaching. AI-mediated learning environments introduce a new object of inquiry. Educators must now examine not only what students produce, but also how AI participates in learning, how cognitive work is distributed between learners and technologies, and whether instructional designs continue to promote ambitious teaching and deeper learning.

**AI Review Cycles** extend the collaborative spirit of PLCs by providing structured opportunities for educators to investigate these questions together using shared evidence, common evaluative criteria, and a common language grounded in the Five Pillars of Practice framework.

This orientation closely aligns with California's long-standing commitment to continuous improvement (Fullan & Quinn, 2015, 2018; Fullan & Rincón-Gallardo, 2017; Furger, Hernández, & Darling-Hammond, 2019). Just as formative assessment supports ongoing instructional refinement rather than one-time evaluation, AI Review Cycles establish iterative feedback loops in which educator experience informs future policy, procurement decisions, professional learning, instructional design, and continuous improvement.

Districts gain actionable evidence regarding implementation, educators contribute to a growing knowledge base about AI-supported teaching and learning as they learn, developers receive classroom-based feedback for improving tools, and policymakers gain insight into how guidance functions under authentic educational conditions.

Importantly, the knowledge generated through these review cycles is not confined to individual classrooms or isolated schools. The value of building pedagogically informed knowledge networks like the AI Review Cycles approach is the potential to augment and increase the power of educator judgment which can accumulate across teams, schools, districts, and eventually statewide networks.

Individual reviews of AI tools, practices and shared artifacts become collective professional (non-corporate) knowledge through a shared evaluation architecture: educators apply common criteria, document implementation evidence, analyze student learning artifacts, and contribute insights to a searchable and exchangeable knowledge base. Because reviews are organized around common constructs and evidence categories (the Five Pillars), findings can be examined across classrooms, grade levels, disciplines, and educational contexts.

Over time, these distributed reviews have the potential to create a statewide knowledge infrastructure capable of informing procurement, professional learning, policy refinement, and continuous improvement with other systems such as LCAP/LCFF (See Recommendations in Section VIII).

Through this dynamic distributed process of AI Review Cycles when coupled with the Five Pillars of Practice allow policy to focus transforming educator expertise into a living knowledge infrastructure by converting professional judgments and classroom experiences into structured, searchable, and exchangeable evidence across California's educational system.

## From AI Tool Adoption to Evaluating Learning and Development

*California policymakers must now ask: Does this AI technology improve learning and guide K-12 learner development? If so, how would we know?*

Learning-centered evaluation begins where technology evaluation ends. Rather than asking only whether AI works, it asks whether AI strengthens student thinking, supports ambitious teaching, enhances equitable participation, improves classroom assessment, and contributes to deeper learning over time. The former evaluates products. The latter evaluates educational value. One determines whether a technology performs successfully; the other determines whether that performance contributes to meaningful educational outcomes.

As Figure 4 illustrates, AI Review Cycles and the Five Pillars of Practice framework together represent a broader evolution in educational AI policy, one that generates its own evidence base.

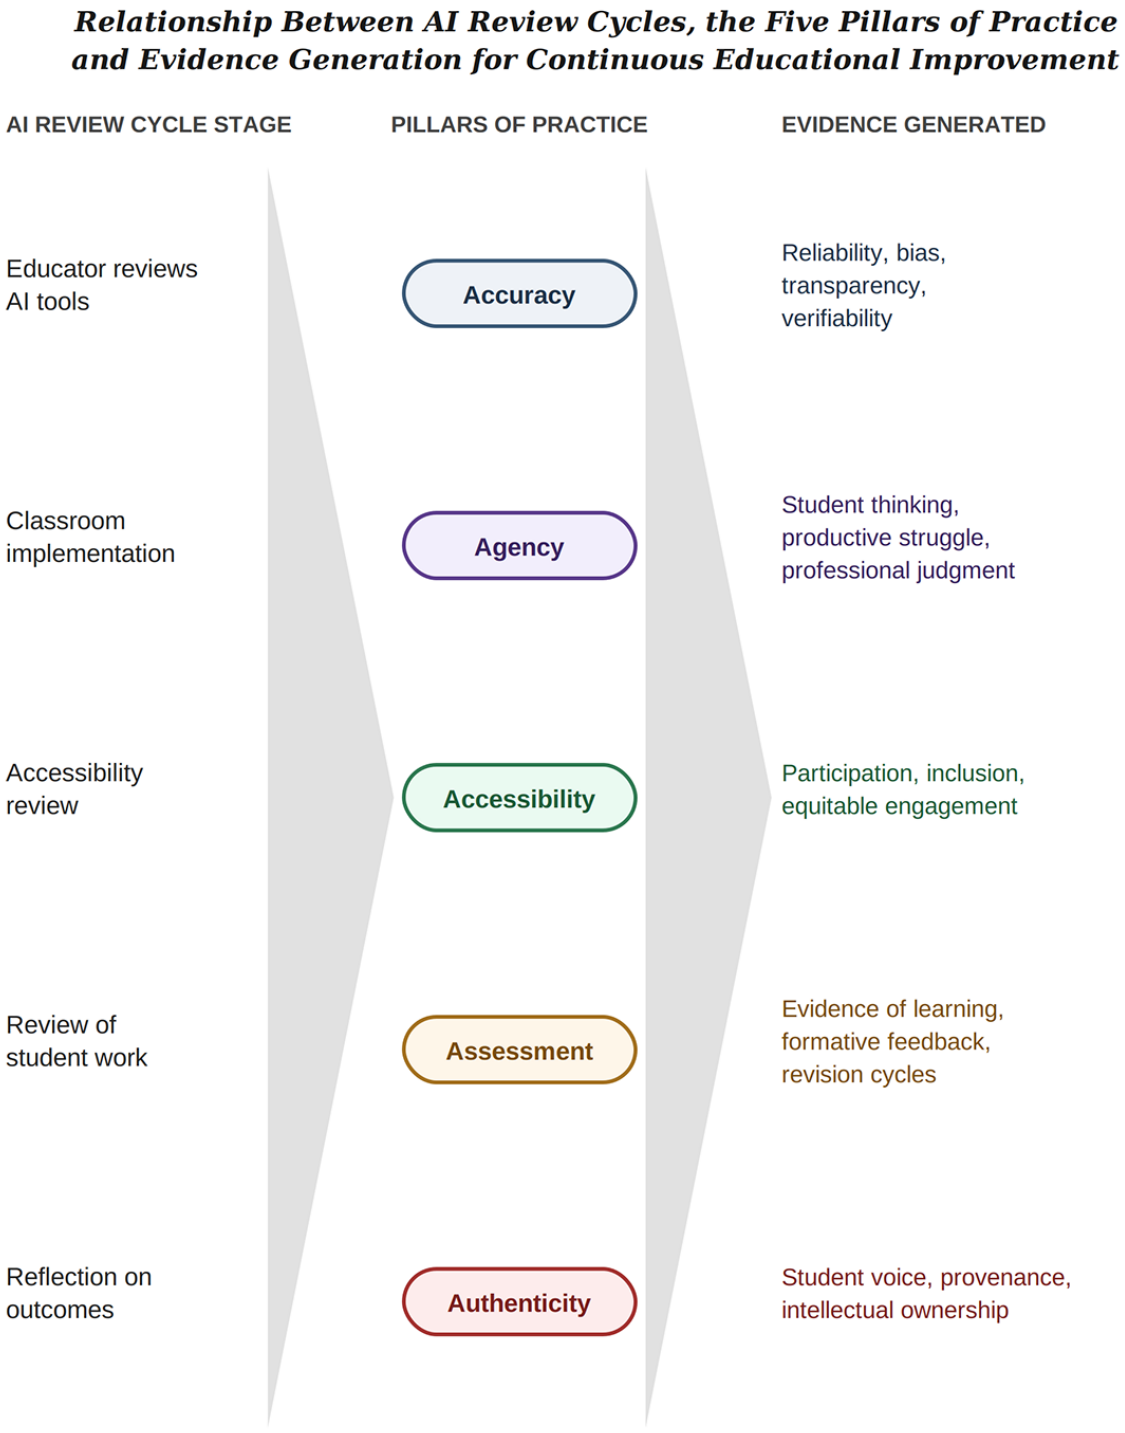


*Figure 4*. Relationship Between AI Review Cycles, the Five Pillars of Practice, and Evidence Generation for Evaluating AI Tools in K–12 Education (Source: IAEP Center 2026).

The review cycle organizes the stages through which educators examine AI tools within authentic instructional settings, while the Five Pillars of Practice framework provides the evaluative lens used to interpret evidence generated at each stage.

Initial reviews emphasize **Accuracy**, focusing on reliability, transparency, bias, and verifiability. Classroom implementation shifts attention to **Agency**, examining where cognitive work resides and whether AI preserves productive struggle, metacognition, and professional judgment. Reviews of learning environments then consider **Accessibility**, asking whether AI expands meaningful participation while maintaining high expectations for all learners. Examination of student work centers on **Assessment**, investigating whether AI strengthens formative assessment, feedback, revision, and evidence of learning. Finally, educator reflection emphasizes **Authenticity**, considering student voice, provenance, intellectual ownership, and the integrity of AI-supported learning.

These review cycles transform evaluation from a one-time procurement activity or compliance checklist into an ongoing process of professional inquiry as shown in Figure 4. More importantly, they reposition educators from consumers of AI technologies to investigators of AI-mediated teaching and learning. Evaluation no longer concludes once a tool has been purchased or approved. Instead, each implementation generates new evidence about how AI influences classroom practice, student development, and instructional decision-making.

This **AI in K-12 education 2.0 policy** progression reflects a broader shift in educational improvement. Rather than evaluating digital technologies including generative AI in isolation, schools begin evaluating the relationships among technologies, pedagogy, assessment, and learning. In doing so, AI Review Cycles, when coupled with core learning science principles like the Five Pillars, create continuous feedback loops through which classroom evidence informs professional learning, procurement decisions, district policy, and future instructional design.

Evaluation thus becomes not simply a mechanism for selecting technologies, but a systematic process for improving educational practice itself.

# VII. Operationalizing Guidance Through the 5POP/AI Review Cycles Evaluation Framework

Translating broad policy principles into classroom practice requires more than aspirational guidance. Educators need practical tools that help them interpret policy, examine evidence, and make consistent instructional judgments across diverse educational settings.

To operationalize these principles, an integrated evaluation framework such as the Five Pillars of Practice ("5POP Evaluation framework") must provide educators, district leaders, researchers, policymakers, and educational organizations with a common architecture for evaluating AI-mediated

learning environments, instructional tools, district policies, professional learning initiatives, and emerging technologies.

Within this integrated architecture, AI Review Cycles serve as an example of the implementation methodology for the 5POP Evaluation Framework. Through recurring cycles of structured educator inquiry, collaborative evidence generation, reflection, and professional learning, AI Review Cycles transform the Five Pillars from a set of evaluative criteria into a knowledge-building process that supports continuous improvement across classrooms, schools, districts, and educational systems.

Unlike traditional technology evaluation rubrics, which often emphasize technical functionality, implementation fidelity, or regulatory compliance, the 5POP Evaluation Framework focuses explicitly on educational quality. Its central question is not simply whether an AI application functions effectively, but whether it advances the kinds of developmentally oriented learning, teaching, assessment, accessibility, and professional judgment that educational systems seek to promote.[ix]

Importantly, the framework is intentionally scalable. Individual teachers may use it to evaluate classroom applications of AI. Professional Learning Communities may employ it to examine instructional practices across grade levels or departments. Schools and districts may use it to review policies, procurement decisions, and professional learning initiatives. State agencies, accreditation organizations, and researchers may apply the same framework when evaluating broader systems of AI implementation. The use of common evaluative criteria across these multiple levels creates opportunities for meaningful comparison while respecting local context and professional judgment.

Potential applications include:

- evaluating AI tools, systems, and machine learning applications before and during implementation;
- guiding district and state AI policy development;
- designing instructional programs and professional learning;
- assessing organizational readiness for AI implementation;
- informing strategic planning and innovation initiatives;
- supporting responsible adoption of emerging technologies;
- shaping standards for accreditation and external program review; and
- promoting ambitious teaching and deeper learning across educational settings.

## Developing Criterion-Based Evaluation Systems: A Deeper Dive into AI Purposes in the K-12 Classroom

A central contribution of the 5POP Evaluation Framework is its use of criterion-based evaluation to translate broad educational principles into observable indicators of quality. Educational policies

frequently reference concepts such as equity, student agency, authentic assessment, accessibility, and responsible AI use.

However, without clearly articulated criteria or robust teacher participation, these concepts often remain aspirational rather than operational (Aljemely, 2024). Criterion-based evaluation bridges this gap by specifying what educators should look for when determining whether policies, instructional designs, or AI-mediated learning environments actually embody these principles.

Each of the Five Pillars of Practice is organized around a core evaluative question that captures the fundamental educational purpose of the domain. Within each pillar, multiple evaluation criteria operationalize that purpose by identifying specific dimensions of instructional design, policy language, organizational practice, or classroom implementation that reviewers can examine systematically.

Rather than relying on general impressions, reviewers gather evidence against clearly defined expectations that can be interpreted consistently across classrooms, schools, districts, and educational organizations.

Each criterion contains four integrated components:

- A **conceptual definition** describing the educational construct being examined;
- Observable **indicators** identifying the evidence reviewers should seek;
- A **developmental rating continuum**, ranging from *Emerging* to *Exemplary*, that recognizes varying levels of implementation; and
- A **pedagogical design principle** synthesizing the broader educational implications of the criterion and connecting individual judgments to the larger goals of ambitious teaching and deeper learning.

This developmental architecture distinguishes the 5POP Evaluation Framework from traditional compliance rubrics. Conventional evaluation instruments often determine whether required elements are present or absent, encouraging binary judgments about compliance. By contrast, criterion-based evaluation recognizes that educational quality develops over time. Policies, instructional systems, AI applications, and professional practices evolve through successive cycles of implementation, reflection, feedback, and refinement. Evaluation therefore becomes a mechanism for supporting organizational learning rather than merely documenting compliance.

This developmental orientation aligns closely with California's long-standing commitment to continuous improvement. Rather than producing static reports, criterion-based reviews generate actionable evidence that educators, school leaders, districts, developers, researchers, and policymakers can use to strengthen future practice. As review findings accumulate across classrooms and educational organizations, they contribute to an expanding body of professional knowledge regarding how AI supports—or constrains—ambitious teaching and deeper learning.

This developmental approach shifts the purpose of review itself. The goal is no longer simply to determine whether an AI policy or technology meets established requirements. The goal is to generate trustworthy evidence that helps educational systems learn from implementation, improve instructional practice, and continually refine how AI contributes to student learning and development.

## A Concrete Example: Evaluating AI and the Role of Student Agency in K-12 Classroom Settings

The **Agency** pillar illustrates how criterion-based evaluation operates within the 5POP Evaluation Framework. Rather than asking reviewers whether a policy merely references human oversight or student choice, the framework examines the extent to which AI-mediated learning environments preserve the intellectual, social, and professional conditions necessary for meaningful learning. In this way, broad educational values are translated into observable criteria that educators, school leaders, policymakers, and researchers can examine systematically.

The core evaluative question guiding the Agency pillar asks policymakers to focus on **key AI pedagogical practice priorities**:

> ***Does the policy or practice preserve meaningful human judgment, intellectual labor, and collaborative decision-making among students, educators, and other educational professionals in AI-mediated learning environments?***

Agency concerns the extent to which AI supports rather than supplants human thinking, decision-making, and learning. High-quality AI policies and practices recognize that learning requires active engagement, productive struggle, reflection, and opportunities for both students and educators to exercise professional judgment. Rather than maximizing efficiency alone, agency-centered policies seek to preserve the cognitive, social, and professional work that underpins ambitious teaching and deeper learning.

To operationalize this construct, the framework organizes evaluation around **five complementary criteria** that together capture the major dimensions of agency within AI-mediated educational environments.

***Criterion 1:*** **Cognitive Labor** examines whether AI encourages learners to engage in the intellectual work required for deeper learning or instead bypasses that work through excessive automation. Reviewers examine expectations for analysis, reasoning, interpretation, synthesis, independent sense-

making, and the distinction between educational assistance and cognitive substitution and/or offloading.

***Criterion 2*: Productive Struggle** evaluates whether AI preserves opportunities for challenge, persistence, revision, and problem solving. Rather than eliminating difficulty, high-quality policies and practices recognize that learning often develops through effort, uncertainty, iterative improvement, and thoughtful feedback.

***Criterion 3*: Goal-Directed Learning** examines whether AI supports self-regulated learning by encouraging students to establish goals, monitor progress, engage in reflection, and participate actively in instructional decisions. Evaluation emphasizes learners' capacity to direct, monitor, and evaluate their own learning rather than passively responding to AI-generated recommendations.

***Criterion 4*: Collaborative Decision-Making** evaluates the extent to which AI preserves teacher professional judgment, student voice, and collaborative interpretation of AI-generated recommendations. Rather than defaulting to automated decisions, effective policies and practices position AI as an advisory partner that informs—but does not replace—human educational judgment.

***Criterion 5*: Reflective Delegation of Effort** examines whether learners and educators thoughtfully consider when AI should appropriately support cognitive work and when human effort remains educationally essential. Particular attention is given to documenting AI-assisted processes, recognizing the trade-offs associated with delegation, and distinguishing between interactive learning and interpassive reliance on AI systems.

Figure 5 shows how each criterion is evaluated using a developmental continuum consisting of four implementation levels:

- **Emerging** — The policy and practices emphasize efficiency, productivity, or automation with limited attention to learning processes or human judgment.
- **Developing** — The policy and practices recognize the importance of the construct but provides limited implementation guidance or supporting evidence.
- **Established** — The policy and practices explicitly support the construct through clear expectations, implementation strategies, and opportunities for educator judgment.
- **Exemplary** — The policy and practices function as a foundational design principle that consistently shapes instructional practice, professional learning, organizational decision-making, and continuous improvement.

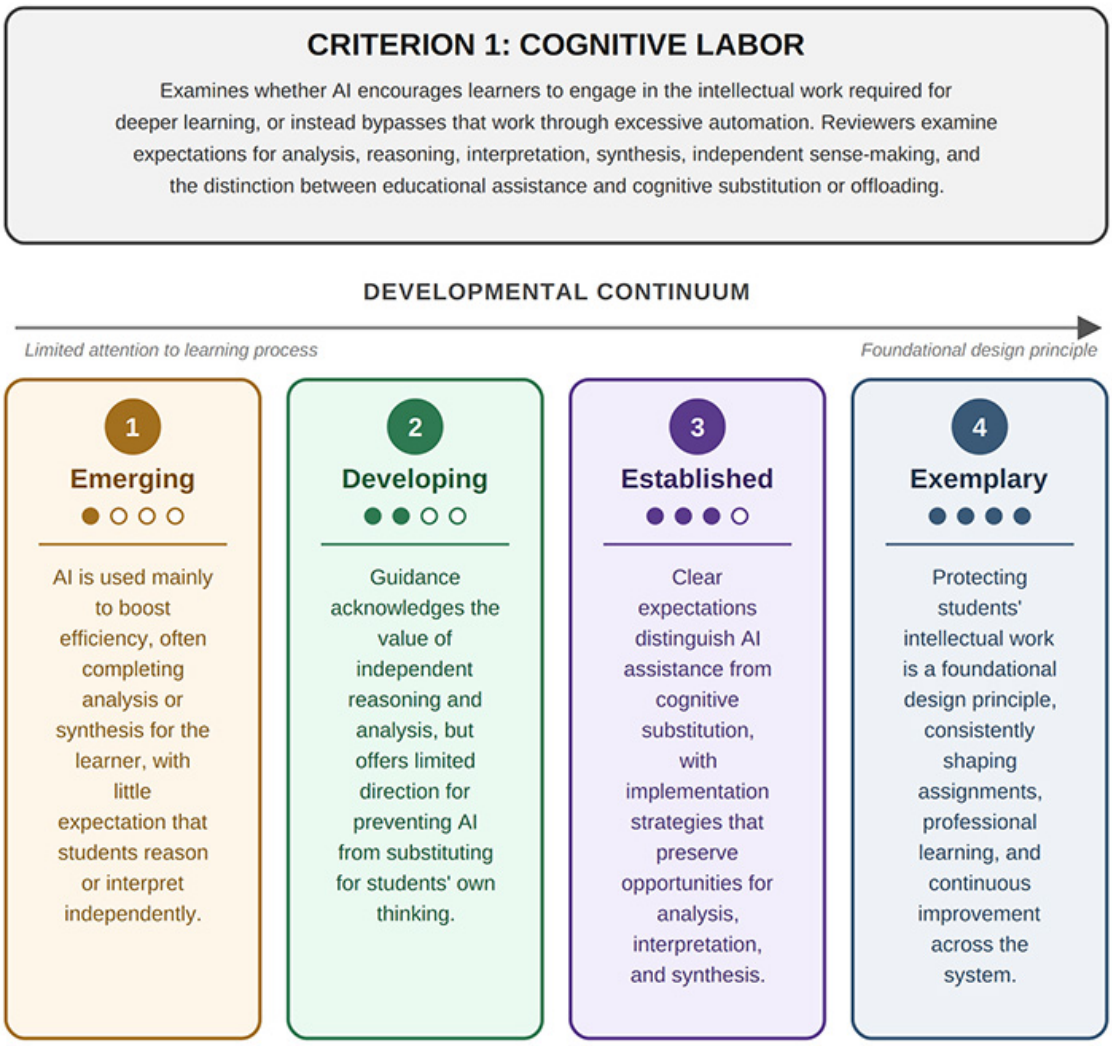


*Figure 5.* Example of Developmental Continuum of AI in Education Implementation (Source: IAEP Center, 2026).

The **developmental continuum** (Figure 5) of the responsible use of AI in educational policy reinforces the notion that educational quality is not a binary condition but an evolving process. Policies, instructional systems, research-based pedagogies, and AI-mediated learning environments may occupy different points along the continuum, providing educators with opportunities to identify strengths, prioritize improvements, and monitor progress over time rather than simply determining whether a policy "passes" or "fails."

The Agency pillar concludes with a broader design principle (Duckor & Holmberg, 2026) that synthesizes the educational purpose underlying the individual criteria. High-quality AI policies do more than keep humans "in the loop." They preserve the intellectual, social, and professional

conditions necessary for learning. Educational AI should strengthen human growth, professional judgment, productive struggle, and deeper learning rather than simply increasing efficiency or automating cognitive work.

Although Agency serves as the illustrative example presented in this paper, the same criterion-based architecture is applied across each of the remaining Four Pillars—Accuracy, Accessibility, Assessment, and Authenticity. Together, these five domains provide educators and policymakers with a coherent, research-informed framework for evaluating whether AI implementation advances the educational purposes that schools ultimately seek to achieve.

## Criteria Design Principles Aimed at Deeper Learning: More than "Human in the Loop" AI Plays

High-quality AI policies do more than keep humans "in the loop." They nurture and preserve the intellectual, social, and professional conditions necessary for learning.

Within the Agency pillar, for example, AI policies at the highest levels of the developmental continuum (Figure 5) recognize that the purpose of educational AI is not merely to increase efficiency or automate routine tasks, but to strengthen human growth, professional judgment, productive struggle, collaborative decision-making, and deeper learning.

Human oversight alone is insufficient if opportunities for thinking, reflection, creativity, and intellectual engagement have already been delegated to AI. The central question is therefore not simply whether humans remain involved, but whether AI expands or diminishes the educational work through which learners and educators develop expertise over time.

More broadly, the Five Pillars of Practice provide a criterion-based architecture for translating California's emerging AI guidance into systematic educational inquiry. By articulating observable criteria for evaluating AI-mediated learning environments, instructional practices, policies, professional learning, and organizational decision-making, the framework enables educators, school leaders, district personnel, researchers, policymakers, and other educational partners to generate evidence that informs instructional improvement and continuous learning. Rather than relying exclusively on periodic policy revisions, procurement reviews, or vendor claims, educational organizations can examine how AI functions within authentic educational settings and whether it advances the learning goals they seek to promote.

Because the 5POP framework is organized around common evaluative criteria, it can be applied through a variety of review processes and organizational contexts. Schools, districts, county offices of education, researchers, accreditation organizations, and state agencies may incorporate the Five Pillars into professional learning, continuous improvement cycles, program evaluation, policy review, procurement protocols, instructional design, or research initiatives. AI Review Cycles represent one

example of how the framework can be operationalized through structured educator inquiry, but the underlying evaluative architecture is intentionally designed to support multiple implementation approaches.

The implications extend well beyond individual schools. As evidence accumulates across classrooms, schools, districts, counties, and states, it begins to form a distributed knowledge base capable of informing procurement decisions, professional learning priorities, policy refinement, accreditation processes, product development, and future educational research. What begins as local educator inquiry becomes organizational knowledge that supports system-level improvement. Thus, California's investment in educator expertise becomes inseparable from its investment in responsible AI use and implementation.

The Five Pillars of Practice are therefore more than a rubric for reviewing AI technologies. They provide a learning-centered evaluative architecture through which educator judgment becomes evidence, evidence becomes professional knowledge, and professional knowledge supports continuous educational improvement. Whether implemented through AI Review Cycles or other criterion-based review processes that promote new knowledge networks for emerging PLCs, the framework positions evaluation not as a compliance exercise but as an ongoing process of educator inquiry, evidence generation, reflection, and organizational learning.

Most importantly, this approach keeps questions of learning, teaching, assessment, and student development at the center of AI implementation rather than allowing technological innovation alone to drive educational decision-making. It reinforces the principle that the success of AI in education should ultimately be judged not by the sophistication of the technology, but by the quality of the learning environments it helps educators create and the developmental opportunities it affords students.

The final section of this paper considers the implications of this learning-centered approach for California's next generation of AI policy. It argues that the state's opportunity is no longer simply to provide guidance for responsible AI adoption, but to establish an integrated system in which policy, pedagogy, professional learning, educator inquiry, and continuous improvement work together to advance ambitious teaching and deeper learning for all students.

# VIII. Conclusion and Implications for California's Next Generation of AI Policy

California has emerged as one of the nation's leaders in developing thoughtful guidance for artificial intelligence in K–12 education. Across the 2023 *Artificial Intelligence Guidance*, the 2024 *AI Resource Kit*, and the 2025 *California Public Schools Artificial Intelligence Working Group* recommendations, the state has steadily expanded its attention from awareness and responsible use toward implementation,

governance, professional learning, data privacy, equity, procurement, and instructional practice. These developments represent a significant maturation of educational AI policy and provide a strong foundation for districts navigating a rapidly evolving technological landscape.

Yet the evolution of California's guidance also reveals the next policy challenge. As artificial intelligence becomes increasingly integrated into teaching and learning, the central questions are no longer primarily technological or merely regulatory.

They are educational and pedagogical:

- **What counts as learning when AI increasingly participates in thinking, writing, problem solving, assessment, and knowledge production?**

- **What constitutes trustworthy evidence of understanding when students can generate sophisticated academic products with varying degrees of AI assistance?**

- **How should educators distinguish between AI that extends learning and AI that substitutes for the cognitive work necessary for development?**

These questions cannot be answered through procurement policies, acceptable use guidelines, governance frameworks, or technical standards alone. They require an explicit theory of learning that explains how students develop understanding, how teachers design ambitious learning experiences, and how assessment data captures real growth over time. Without such a framework, educational systems risk confusing polished performance with learning, efficiency in outputs with development, and access to increasingly powerful AI technologies with meaningful and *educative* educational opportunities.

This white paper has argued that the Five Pillars of Practice—*Accuracy, Agency, Accessibility, Assessment, and Authenticity*—provide one such learning-centered policy architecture. Rather than functioning as an alternative policy framework, the Five Pillars of Practice serve as a developmental and interpretive lens through which existing policies and classroom-based practices can be examined, strengthened, and continuously improved.

Coupled with examples such as AI Review Cycles, the Five Pillars of Practice framework extends California's current guidance beyond questions of governance toward systematic inquiry into how AI influences teaching, learning, assessment, accessibility, and student development. Together, the Five Pillars of Practice with a focus on knowledge networking tools establish a practical bridge between policy and pedagogy by providing educators and educational leaders with a common language for evaluating AI-mediated learning environments and generating evidence that supports continuous improvement.

California now has an opportunity to move beyond responsible AI adoption toward a more comprehensive vision of responsible AI-enabled learning. Doing so would position the state not only as a leader in AI governance, but also as a leader in articulating what ambitious teaching and deeper learning should look like in an era of increasingly capable intelligent systems.

The state's existing infrastructure—including School Accountability Report Cards, Local Control and Accountability Plans, the Local Control Funding Formula, Learning Management Systems, Professional Learning Communities, and continuous improvement networks—provides an ideal foundation for embedding learning-centered AI evaluation within everyday educational practice.

## Implications for California's Next Generation of AI Policy

Building on the analysis and recommendations presented throughout this paper, **six complementary priorities** emerge for strengthening California's next generation of AI policy by more intentionally connecting governance to pedagogy, pedagogy to assessment, and assessment to continuous improvement.

The recommendations that follow are intentionally framed as statewide policy priorities rather than implementation directives. California's existing educational governance structure—including the California Department of Education, State Board of Education, county offices of education, local educational agencies, institutions of higher education, California Subject Matter Projects, professional learning organizations, and other educational partners—provides multiple pathways through which these priorities can be advanced.

Rather than prescribing a single implementation model, the recommendations identify opportunities for strengthening the connections among AI governance, pedagogy, assessment, and continuous improvement through California's existing educational infrastructure.

**First, future state guidance should more explicitly incorporate concepts from the learning sciences into California's AI policy framework.** As the California Department of Education, State Board of Education, county offices of education, California Subject Matter Projects, and other educational partners continue developing guidance and professional learning resources, concepts such as productive struggle, scaffolding, metacognition, developmental progression, internalization, transfer, self-regulated learning, and the Zone of Proximal Development can provide educators with a shared conceptual vocabulary for determining whether AI supports—or substitutes for—student learning. These concepts would complement existing guidance on ethics, governance, and responsible use by grounding AI policy within well-established theories of teaching, learning, and development.

**Second, California should encourage the development of practical implementation resources that operationalize the Five Pillars of Practice across multiple levels of the educational system.** Working through existing partnerships among the California Department of

Education, county offices of education, institutions of higher education, California Subject Matter Projects, and professional learning organizations, these resources could include district policy review protocols, school improvement rubrics, classroom observation instruments, AI Review Cycle facilitation guides, instructional design tools, student self-reflection protocols, and organizational dashboards that enable educators to evaluate AI-supported learning through consistent criterion-based processes.

**Third, California should continue integrating AI within its existing continuous improvement infrastructure rather than treating it as a separate educational initiative.** Existing policy and improvement systems—including School Accountability Report Cards (SARCs), Local Control and Accountability Plans (LCAPs), Local Control Funding Formula (LCFF) reporting, Learning Management Systems (LMSs), Professional Learning Communities (PLCs), and related statewide improvement processes—provide natural opportunities for documenting how AI influences teaching, assessment, educator professional learning, and student development. Building upon these established systems would strengthen transparency, expand local capacity, and avoid unnecessary policy fragmentation.

**Fourth, California should support the development of a statewide infrastructure for continuous AI learning through pilot implementation and refinement of the 5POP/AI Review Cycles Evaluation Framework described in this paper.** By encouraging distributed educator inquiry across Professional Learning Communities, schools, districts, county offices of education, institutions of higher education, and other educational partners, California can transform individual evaluations into a cumulative evidence base that informs procurement, professional learning, accreditation, policy refinement, educational research, and future innovation. In this way, educator judgment becomes evidence, evidence shared through digital Professional Learning Community (PLC) networks becomes professional knowledge, and locally vetted professional knowledge becomes the foundation for continuous educational improvement across California's educational system.

**Fifth, California should continue supporting broad stakeholder engagement to refine a shared vision for AI-mediated education.** Existing advisory bodies, statewide working groups, professional organizations, educational agencies, and cross-sector partnerships provide important venues for convening educators, researchers, policymakers, students, families, and technology developers as AI continues to evolve. Sustained dialogue across these communities will help ensure that future guidance remains adaptive while remaining anchored to enduring commitments to educator professional judgment, educational equity, democratic participation, and human development.

**Sixth, California should expand support for sustained statewide research-practice partnerships examining how AI influences teaching and learning over time.** Collaborations among the California Department of Education, county offices of education, school districts, institutions of higher education, nonprofit organizations, industry partners, and state agencies can

generate robust evidence regarding instructional effectiveness, student agency, classroom assessment, accessibility, and equitable implementation.

Future policy development should also examine how these research-practice partnerships can inform teacher educator preparation, pre-service programs, induction, credentialing, accreditation, and ongoing professional learning in collaboration with the California Commission on Teacher Credentialing, institutions of higher education, and other subject matter-focused educational partners. Such partnerships would enable policy to evolve in response to accumulating educational evidence rather than technological change alone while strengthening California's long-standing commitment to continuous improvement.

## Leaning into Guidance That Hits Home Where Students Live and Learn

The next generation of AI policy will not be defined solely by the sophistication of its technology guidance. It will be defined by the extent to which it strengthens the intellectual, social, relational, and professional conditions that make meaningful learning and human development possible across California's K–12 educational system.

California now stands at an important moment of statewide educational leadership. As new state leaders, educational agencies, county offices, school districts, and professional organizations continue shaping the future of AI in education, they have an opportunity to build upon an already strong foundation of governance by further integrating the learning sciences into statewide policy, professional learning, and continuous improvement.

The framework presented in this paper is intended not as a final blueprint, but as a resource for ongoing dialogue, adaptation, and evidence-informed policy development as educational technologies and instructional practices continue to evolve. The **Five Pillars of Practice** provide California with a learning-centered framework for evaluating the quality, equity, and developmental value of AI-mediated teaching, learning, and assessment. They help educators and policymakers examine how meaningful learning opportunities are created, experienced, and evidenced across diverse educational communities while establishing **AI Pedagogical Practice Priorities** that can strengthen statewide coherence.

As California continues strengthening its educational systems for the age of AI, sustained investments in educator preparation, credentialing, induction, and on-going professional learning will be essential for cultivating the professional judgment required to design equitable, evidence-informed, and deeply human learning environments.

California's next golden opportunity is to lead with clarity and confidence in defining what counts as genuine learning and human development in an age of artificial intelligence. Realizing this opportunity will require sustained investment in the knowledge infrastructure through which

educators, policymakers, students, families, and communities continually improve the ways AI serves ambitious teaching, deeper learning, and the flourishing of every learner.

---

[i] See, e.g., Lee, V. R. (2026). *AI and education across California schools: Calls for more professional development and AI literacy*. Getting Down to Facts III, Stanford University.

[ii] See Cuban, L. (1986). *Teachers and machines: The classroom use of technology since 1920*. Teachers College Press and Cuban, L. (2001). *Oversold and underused: Computers in the classroom*. Harvard University Press.

[iii] AI Policy to practice gaps abound across the U.S. A recent study notes that although 35 states and territories had issued AI guidance for schools by the end of 2025, that guidance remains nonbinding and focused on broad principles rather than concrete requirements. Even among early-adopter districts, nearly half of surveyed principals nationally reported receiving no AI policy or guidance at all from their school or district (Doss et al., 2025; Dusseault et al., 2026).

[iv] Within the United States, these evolving knowledge networks extend well beyond state education agencies themselves. Educational policy organizations, research centers, professional associations, philanthropic initiatives, and intermediary organizations increasingly contribute to the development, interpretation, and dissemination of AI policy guidance. Organizations such as the Center for Democracy and Technology, Digital Promise, TeachAI, the Stanford Institute for Human-Centered Artificial Intelligence, the Education Commission of the States, Brookings, RAND, WestEd, and others regularly synthesize emerging state guidance, convene stakeholders across sectors, identify promising practices, monitor legislative developments, and translate research into practical recommendations for educators and policymakers. Collectively, these organizations function as knowledge intermediaries that connect research, policy, and practice while accelerating the exchange of ideas across states and educational systems.

[v] The emerging global ecosystem of educational AI policy is characterized less by convergence than by increasingly diverse national approaches to when, how, and why artificial intelligence should become part of children's educational experiences. Recent developments illustrate this divergence. In 2026, China adopted a comprehensive national AI + Education Action Plan that seeks to integrate artificial intelligence across all levels of education by 2030, beginning with primary school. The plan positions AI as a developmental resource for cultivating curiosity, problem solving, interdisciplinary learning, and whole-child development while simultaneously investing in national infrastructure, teacher preparation, intelligent learning companions, adaptive instructional systems, and large-scale educational data systems. AI is envisioned not simply as an instructional technology but as a foundational component of an intelligent education system extending from primary schooling through higher education and lifelong learning.

By contrast, Norway has adopted an age-tiered approach grounded, in part, in developmental psychology. Beginning with the 2026–2027 school year, students ages 6–13 are generally prohibited from using generative AI in classroom learning, students ages 14–16 may use AI only under direct teacher supervision, and older secondary students are expected to develop responsible, independent uses of AI as preparation for higher education and employment. Norwegian policymakers have linked these restrictions to concerns that premature reliance on AI may undermine the development of foundational literacy, numeracy, reasoning, creativity, and problem-solving capacities during early childhood. Rather than rejecting AI outright, the policy reflects a developmental judgment about when learners are cognitively prepared to benefit from AI-assisted learning.

These contrasting national strategies underscore a central argument of this paper. Although jurisdictions differ substantially in their governance approaches—from early, systemwide AI integration to delayed, developmentally staged adoption—they are increasingly participating in a shared international conversation about how artificial intelligence should support learning, human development, and public education. California's continuing influence will depend less on aligning with any single international model than on contributing theoretically informed, evidence-based frameworks that help connect AI governance with what we know about ambitious teaching, deeper learning, and the development of expertise.

[vi] The current landscape of K–12 AI policy across the United States reflects rapid growth in state-level guidance but comparatively limited attention to theories of teaching and learning. As of mid-2026, more than

thirty-five states and Puerto Rico had issued official guidance, policy frameworks, templates, or recommendations addressing artificial intelligence in schools. Although these documents vary considerably in scope, intended audience, and implementation strategy, comparative reviews identify broad convergence around responsible AI use, educator and student AI literacy, data privacy and cybersecurity, equity and accessibility, procurement and governance, workforce preparation, and the importance of maintaining meaningful human oversight over AI-supported decision making.

Within this broader landscape, several bellwether states—including California, Texas, Florida, North Carolina, Ohio, and Virginia—have assumed outsized importance because of their large student populations, distinctive approaches to educational governance, and influence on national policy conversations. Collectively, these states illustrate that there is no single American model for AI policy in education. Instead, they represent a range of governance strategies that differ in their emphasis on state versus local authority, instructional innovation, educator preparation, procurement, privacy protections, and implementation. As these large and influential educational systems continue to refine their approaches, their experiences are likely to shape policy development well beyond their own borders.

At the same time, guidance remains uneven in its treatment of pedagogy. Comparative analyses by the Center for Democracy and Technology (Dwyer, 2025) and Digital Promise (Roschelle at al., 2024) conclude that most state guidance prioritizes governance, implementation, and risk mitigation while providing comparatively limited attention to how AI should support ambitious teaching, classroom assessment, or deeper learning.

[vii] Policy experts realize the potential for policy overload and the system of support providers are struggling to absorb the pace and competing demands for their attention brought on by these new technologies. As Gallagher (2026) observes, artificial intelligence is entering educational systems at a time when district leaders are already managing an exceptionally crowded policy environment. Interviews with California superintendents suggest that leadership attention remains focused on foundational priorities such as early literacy, student well-being, fiscal pressures associated with declining enrollment, local governance challenges, and accountability under the California School Dashboard.

Within this context, AI often competes with numerous established initiatives rather than emerging as a primary strategic priority. At the same time, growing public concern about student technology use, smartphones, and declining educational outcomes may create pressure for simplified policy responses—including restrictions or bans—when educational systems lack the capacity to engage more comprehensive questions about teaching, learning, and implementation.

This observation reinforces the importance of designing AI policy that strengthens existing systems of continuous improvement rather than adding another competing initiative to an already crowded educational agenda.

[viii] These trends illuminate a broader challenge emerging across national and international policy conversations (See e.g., Duckor & Holmberg, May 2026). References to personalization, accessibility, agency, student engagement and "human-centered" AI appear throughout California's guidance—as they do within many international frameworks—but these concepts remain only loosely connected to explicit theories of cognitive development, sociocultural learning, classroom assessment, ambitious teaching, and the development of expertise.

[ix] Each of the pillars that build the foundation of responsible AI use in K-12 education contains clearly articulated evaluation criteria describing observable characteristics of responsible AI practice. Reviewers (including teachers, instructional coaches, librarians, paraprofessionals, school leaders, district personnel, researchers, and other educational partners) examine evidence against these criteria using developmentally structured reflection protocols, criteria specific rating scales, and collaborative discussion processes that generate both quantitative ratings and qualitative evidence. The resulting information supports educator dialogue, organizational learning, policy refinement, procurement decisions, instructional design decisions, and continuous improvement. Asking the right questions at the right level of decision-making is key. (See **Appendix A** as an example of systemic approached to inquiry on AI uses in K-12 education).

# References


Aljemely, Y. (2024). Challenges and best practices in training teachers to utilize artificial intelligence: A systematic review. *Frontiers in Education*, 9, Article 1470853.

Alexander, R. (2008). Towards dialogic teaching: Rethinking classroom talk (4th ed.). *Dialogos*.

Bransford, J. D., Brown, A. L., & Cocking, R. R. (Eds.). (2000). *How people learn: Brain, mind, experience, and school* (Expanded ed.). National Academy Press.

Bruner, J. S. (1960). *The process of education*. Harvard University Press.

Bruner, J. S. (1990). *Acts of meaning*. Harvard University Press.

California Department of Education [CDE]. (2023). *Learning with AI, learning about AI: 2023 guidance regarding the role of artificial intelligence (AI) in California kindergarten through grade 12 (K-12) education.* https://www.cde.ca.gov/ci/pl/aiincalifornia2023.asp

California Department of Education. (2024). *Artificial intelligence resource kit*. https://www.cde.ca.gov/ci/pl/documents/cdeairesourcekit.pdf

California Public Schools Artificial Intelligence Working Group. (2025). *Artificial intelligence in California public schools*. California Department of Education. https://www.cde.ca.gov/ci/pl/aiincalifornia.asp

Cuban, L. (1986). *Teachers and machines: The classroom use of technology since 1920*. Teachers College Press.

Cuban, L. (2001). *Oversold and underused: Computers in the classroom*. Harvard University Press

Darling-Hammond (1997) *The right to learn: A blueprint for creating schools that work*. San Francisco: Jossey-Bass.

Darling-Hammond, L. (2006). Securing the right to learn: Policy and practice for powerful teaching and learning. *Educational Researcher*, *35*(7), 13–24.

Dewey, J. (1933). How we think: A restatement of the relation of reflective thinking to the educative process. In J. Boydston (Ed.), *The later works of John Dewey (Vol. 8)* (pp. 105-352). Carbondale, IL: Southern Illinois University Press.

Dewey, J. (1938). *Experience and education*. Macmillan.

Doss, C. J., Diliberti, M. K., Robbins, M. W., Ramchand, R., & Grant, D. (2025). *AI use in schools is quickly increasing but guidance lags behind: Findings from the RAND survey panels* (Research Report No. RRA4180-1). RAND Corporation.

Duckor, B., & Holmberg, C. (2026, May, 6). *Designing for agency in AI-mediated classrooms: From cognitive offloading to productive struggle aimed at deeper learning* [Conference presentation]. 2nd Annual AI Futures Symposium: On the Role of Agency in the Age of Innovation in Public Education, Budapest, Hungary.

Duckor, B., & Holmberg, C. (2025). Five pillars of ethical AI use for teaching and learning. *Kappan*, *107*(3-4), 14-18.

Duckor, B., & Holmberg, C. (2026). *AI and deeper learning: Promises, paradoxes, and evolving practices*. Harvard Education Press.

Duckor, B., & Holmberg, C. (2023). *Feedback for continuous improvement in the classroom: New perspectives, practices, and possibilities*. Corwin.

Duckor, B., & Holmberg, C. (2017). *Mastering formative assessment moves: 7 high-leverage practices to advance student learning*. ASCD.

Dusseault, B., Haderlein, S., Prymula, E., Waite, C., Fall, M., Berardino, M., & Harrison, D. (2026). How state policy can help teachers use AI well. *The State Education Standard, 26*(2). National Association of State Boards of Education.

Dwyer, M. (2025, April 15). *Looking back at AI guidance across state education agencies and looking forward*. Center for Democracy and Technology.

Fetterman, D.M., and Ravitz, J. (2018). A Google enhanced empowerment evaluation approach in a graduate school program. In Fetterman, D.M., Rodriguez-Campos, L., and Zukoski, A. (eds.) *Collaborative, Participatory, and Empowerment Evaluation: Stakeholder Involvement Approaches*. New York: Guilford Press.

Fullan, M., & Quinn, J. (2015). *Coherence: The right drivers in action for schools, districts, and systems*. Corwin.

Fullan, M., Quinn, J., & McEachen, J. (2018). *Deep learning: Engage the world, change the world*. Corwin.

Fullan, M. & Rincón-Gallardo, S. (2017). *California's golden opportunity: Taking stock: Leadership from the middle*.

Furger, R.C., Hernández, L.E., & Darling-Hammond, L. (2019). *The California way: The Golden State's quest to build an equitable and excellent education system*. Learning Policy Institute.

Gallagher, A. (2026, July 29). *I'm going to look at a level below the state, because districts are close enough to the classroom to drive* [Comment on the post "I wish I could get researchers, policy experts, education leaders, teachers, and others together and ask, 'How did current edtech backlash come to be?'"]. *LinkedIn*.

Grossman, P., Compton, C., Igra, D., Ronfeldt, M., Shahan, E., & Williamson, P. (2009). Teaching practice: A cross-professional perspective. *Teachers College Record, 111*(9), 2055–2100.

Hardwig, J. (1985). Epistemic dependence. *The Journal of Philosophy*, *82*(7), 335–349.

Kissinger, H. A., Schmidt, E., & Mundie, C. (2024). *Genesis: Artificial intelligence, hope, and the human spirit*. Little, Brown and Company.

Lampert, M. (2001). *Teaching problems and the problems of teaching*. Yale University Press.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.

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

Mehta, J., & Fine, S. M. (2019). *In search of deeper learning: The quest to remake the American high school*. Harvard University Press.

National Academies of Sciences, Engineering, and Medicine. (2018). *How people learn II: Learners, contexts, and cultures*. The National Academies Press.

National Research Council (NRC). (1999). *How people learn: Mind, brain, experience, and school* (J. D. Bransford, A. L. Brown, & R. R. Cocking, Eds.). National Academies Press.

National Research Council (NRC). (2001). *Knowing what students know: The science and design of educational assessment* (J. W. Pellegrino, N. Chudowsky, & R. Glaser, Eds.). National Academies Press.

New Mexico Public Education Department. (2025). *New Mexico AI guidance for K-12 education* (Version 2025.1). https://web.ped.nm.gov/wp-content/uploads/2025/05/NM-AI-Guidance-Signed-4-29-2025.pdf

OECD. (2025). *Education policy outlook 2025: Empowering lifelong learning*. OECD Publishing. https://doi.org/10.1787/c3f402ba-en

Ravitz, J. (2004). A doorway to new tools and practices: Supporting teacher education, research, and development with an online netcourse. In N. Sabelli and R. Pea. (Eds.). *Uniting people, technology and powerful ideas for learning. Six years of knowledge networking in learning sciences and technologies*, 94-101.

Ravitz, J. & Hoadley, C. (2005) Supporting change and scholarship through review of online resources in professional development settings. *British Journal of Educational Technology, 36*(6), 967-974. https://researchgate.net/publication/227633312

Ravitz, J. (2026, June). AI review cycles for educators: Professional learning as the engine for AI evidence, capacity, and better decision-making. [Post]. *LinkedIn* https://linkedin.com/pulse/ai-review-cycles-educators-jason-ravitz-ph-d--w6s2c/

Resnick, L. B., & Resnick, D. P. (1992). Assessing the thinking curriculum: New tools for educational reform. In B. R. Gifford & M. C. Connor (Eds.), *Changing assessments: Alternative views of aptitude, achievement, and instruction* (pp. 37-75). Kluwer Academic.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, *20*(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Rogoff, B. (1990). *Apprenticeship in thinking: Cognitive development in social context*. Oxford University Press.

Roschelle, J., Fusco, J., & Ruiz, P. (2024, February 26). *Review of guidance from seven states on AI in education*. Digital Promise.

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

Sadler, R. (1989). Formative assessment and the design of instructional systems. *Instructional Science*, 18, 119-144.

Shepard, L. A. (2000). The role of assessment in a learning culture. *Educational Researcher*, *29*(7), 4-14.

Shepard, L. A. (2021). Ambitious teaching and equitable assessment: A vision for prioritizing learning, not testing. *American Educator*, *45*(3), 28.

Shulman, L. S. (1987). Knowledge and teaching: Foundations of the new reform. *Harvard Educational Review, 57*(1), 1–23.

Tyack, D. B., & Cuban, L. (1995). *Tinkering toward utopia*. Harvard University Press.

UNESCO. (2024, October 24). *Artificial intelligence and the futures of learning.* https://www.unesco.org/en/digital-education/ai-future-learning

Vermont Agency of Education. (2026, January 23). *AI guidance for education* (Version 1.0). https://static1.squarespace.com/static/64398599b0c21f1705fb8fb3/t/69866f7bf56f56652beb60e8/1770418043581/Vermont+edu-aoe-ai-guidance-for-education.pdf

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.

Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice*, *41*(2), 64–70.

# About the Authors

**Brent Duckor, Ph.D.** is Professor in the Department of Teacher Education and Executive Director of the Center for Innovation in Applied Education Policy (IAEP) at San José State University. His scholarship focuses on classroom assessment, formative feedback systems, applied educational measurement, and teacher professional learning. For more than two decades, his research has examined how educators develop assessment expertise and how evidence of student learning can support ambitious teaching, instructional improvement, and more equitable learning opportunities. Dr. Duckor's internationally recognized work on *Mastering Formative Assessment Moves: 7 High-Leverage Practices to Advance Student Learning* (ASCD, 2017) and *Feedback for Continuous Improvement in the Classroom: New Perspectives, Practices, and Possibilities* (Corwin, 2023), with Dr. Carrie Holmberg, provides research-based frameworks for strengthening classroom assessment practice, formative feedback, and educator decision-making.

His current work focuses on the intersection of artificial intelligence, classroom assessment, and deeper learning. He is co-author with Dr. Holmberg of *AI and Deeper Learning: Promises, Paradoxes, and Evolving Practices* (Harvard Education Press, 2026), which introduces the Five Pillars of Practice framework© for guiding ethical and learning-centered uses of AI in education. Dr. Duckor holds a Ph.D. in Quantitative Methods and Evaluation from the University of California, Berkeley, a Master's in International Affairs from Columbia University, and a B.A. in Politics with Honors from the University of California, Santa Cruz.

**Carrie Holmberg, Ed.D.** is Co-Director of the Center for Innovation in Applied Education Policy (IAEP) at San José State University and an educational researcher specializing in formative assessment, feedback practices, teacher learning, instructional improvement, and deeper learning. Her scholarship examines how educators design classroom environments that strengthen student agency, equitable participation, and meaningful engagement with complex ideas.

Her current work focuses on supporting educators and educational systems as they navigate responsible uses of artificial intelligence for teaching, learning, and assessment. As co-author of *AI and Deeper Learning: Promises, Paradoxes, and Evolving Practices* (Harvard Education Press, 2026), she helps translate research from the learning sciences into practical frameworks for AI integration, professional learning, and instructional design in pre- and in-service preparation settings. Dr. Holmberg holds an Ed.D. in Educational Leadership from San José State University, an M.A in

Education from the Stanford Teacher Education Program, and B.A. in English from Stanford University.

**Jason Ravitz, Ph.D.** is President and Founder of Evaluation by Design, where he helps education, technology, workforce, and social impact organizations connect strategy, measurement, analytics, and evidence generation for learning and continuous improvement. He has more than 25 years of experience designing evaluation systems, measures, dashboards, mixed-methods studies, and evidence-building strategies to support educational innovation, including leadership roles at Google, Digital Promise, and the Buck Institute for Education. His work has supported initiatives funded by the National Science Foundation (NSF), National Institutes of Health (NIH), U.S. Department of Health and Human Services (HHS), and the U.S. Department of Education.

His current work through Evaluation by Design includes the development of AI Review Cycles, a professional learning-centered approach for generating practical evidence about artificial intelligence and educational technology use within authentic instructional settings. This work supports educators and organizations in moving beyond technology adoption toward systematic inquiry, evaluation, and improvement. Dr. Ravitz holds a Ph.D. in Instructional Design, Development and Evaluation from Syracuse University and a B.A. from Harvard College.