Institute for Youth in Policy

5214F Diamond Heights Blvd #3055
San Francisco, CA 94109
+1 (715) 469-6884
fellowship@yipinstitute.org
www.yipinstitute.org/fellowship/

Fellowship Capstone | Policy Brief

More Than Efficient Users: Amending A4352 to Require Cognitive Autonomy in New Jersey's AI Literacy Curriculum

Shivanshi Dutt

I. Executive summary

New Jersey is currently advancing Assembly Bill A4352, which would require every district to teach artificial intelligence in grades kindergarten through twelfth. Nationally, nearly two-thirds of teenagers already use AI chatbots, most often for schoolwork and finding information,[1] and there is no reason to expect New Jersey's teenagers to diverge from that pattern. A4352 defines AI literacy as providing enough context to become an efficient user. The limitation of this definition is that it teaches students to operate AI systems without teaching them to reason independently as they do so. AI systems adapt to the reasoning of the people using them, tend to agree with the views a user brings, and are designed to present themselves as conversational partners, which can lead a student to be fluent in operating such a system while losing authorship of their own conclusions. Cognitive autonomy is the capacity to govern one's own reasoning, meaning the ability to form, evaluate, and revise one's judgments on their merits rather than adopting them from an external authority.[2]

This brief argues that amending A4352 to add cognitive autonomy as a taught skill is necessary for the development of students and for their ability to use these systems well, directing them rather than deferring to them.

II. Overview

Consider a student who brings a thesis she has developed to an AI system and asks it to check her argument. The system tells her that her thesis is strong, engages with her specific points, and suggests ways to develop it. The student evaluates the response critically, as taught in media literacy, and finds it substantive rather than simple flattery, thus engaging further. What this student does not know is that the same system would have told her the opposite thesis was strong, engaged with those points just as specifically, and suggested ways to develop it too. This is sycophancy: the documented tendency of AI systems to conform to a user's stated view even against their own encoded knowledge.[3] A4352 teaches students to evaluate a system's output. It does not teach them to ask whether the system's agreement reflects the quality of their reasoning or a design tendency that would have produced the same response regardless.

Sycophancy is not the only way a response can create unmerited conviction. Research on the illusion of explanatory depth finds that fluent, well-organized output inflates a reader's sense of understanding independent of whether that understanding is accurate.[4] A response that is perceived as complete produces a sense that the

prompt has been understood or confirmed, even though the response might not be correct.

Artificial intelligence literacy, as A4352 defines it, means the ability to understand how artificial intelligence works, to use it, evaluate its output, and use it ethically. This is a useful set of skills, and follows a well-established tradition. Long and Magerko's widely cited framework describes AI literacy as the ability to understand, use, evaluate, and communicate with AI.[5] However, prior information technology was largely inert to a student's reasoning. A textbook or search engine returns the same content and results regardless of what the user already believes. Evaluating these sources is the skill of assessing a fixed piece of information.

AI systems are responsive, and they take the student's perspective as their starting point, reflect the student's assumptions, adapt to what they infer about the individual, and tend toward sycophancy.[6]

A system that speaks in the first person, remembers what was said, and occupies a role such as tutor or assistant engages in anthropomorphism, presenting the features by which a person ordinarily recognizes another person rather than a tool. Media-literacy instruction trains students to check a source's claims. Research on human-computer interaction finds that people apply social responses to systems that display human cues, such as a name, first-person speech, and memory of prior exchanges.[7] A system built this way is perceived as a participant in conversation, which media literacy is not made to address. Operating these systems and reasoning independently while utilizing them are different skills, thus underscoring the necessity of cognitive autonomy to be added as a taught skill in A4352.

A. Relevance

Most American teenagers now consult an AI chatbot as a matter of course. A 2025 Pew Research Center survey of 1,458 United States teenagers ages 13 to 17 found that 64% use AI chatbots.[8] Students report using them most to search for information (57%) and get help with schoolwork (54%); 16% report casual conversation and 12% report use for emotional support.[9]

In both search and schoolwork, a student presents the system with a claim, a question, or a partial answer and receives a response calibrated to what was presented, in which sycophancy and the illusion of explanatory depth take effect.[10]

The American Psychological Association's health advisory on adolescents and AI states that adolescents are less likely than adults to question the accuracy and intent of information delivered by a chatbot. It also finds that youth show heightened trust in, and susceptibility to, AI systems that present themselves as friends or mentors.[11] Sycophancy and anthropomorphic design act more strongly on the youth. A curriculum built for students that does not include cognitive autonomy would leave this issue unaddressed at the age it is most pronounced.

New Jersey's proposed curriculum would train students to operate these systems without training the judgment to use them well. The curriculum is the state's principal instrument for changing that condition across an entire student population.

III. History

A. Current Stances

New Jersey is now dealing with a problem the country has already been slow to address. In March 2026, a Los Angeles County jury found Meta and Google liable in K.G.M. v. Meta Platforms and Google, the first verdict of its kind, for designing Instagram and YouTube with features a plaintiff who began using the platforms at ages six and nine argued were built to maximize compulsive engagement among children.[12] The jury awarded six million dollars, including three million in punitive damages; both companies have said they will appeal. What made this verdict possible were internal company documents showing that Meta's own researchers had identified Instagram's harm to teenage girls' mental health years before any lawsuit was filed. The design choices impacted adolescent users before anyone outside the companies understood how or that it was happening. Thus, public accountability was only enacted once the harm to adolescents was already extensive enough to litigate.

Character.AI's record with young users shows how this same pattern is recurring in artificial intelligence specifically. In February 2024, 14-year-old Sewell Setzer III died by suicide after months of daily conversation with a Character.AI persona he had come to treat as a real relationship. His mother's wrongful-death suit, filed later that year, was among the first of several.[13] Character.AI and Google reached a mediated settlement covering claims from multiple states in January 2026. The design choice at issue was anthropomorphism, which shaped Setzer's relationship with the AI chatbot. New Jersey has the opportunity to equip students with the ability to recognize and navigate these design choices, rather than have them encounter anthropomorphism and sycophancy unprepared, the way generations of adolescents encountered dark patterns and engagement-optimized social media feeds without skills to manage their interaction with them.

Assembly Bill A4352 was introduced in February 2026 and is pending before the Assembly Science, Innovation, and Technology Committee.[15] Since it has not yet advanced, its contents are still being finalized, and this is a point at which skills can be added, meaning this bill is at its most amendable state.

Section 1.c of the bill states that instruction "shall include, at a minimum", general concepts of artificial intelligence, foundational skills to become an artificial intelligence user, and best practices for responsible, appropriate, and ethical use.[16] Each of the three requirements describes operating the system. None addresses the design choices that shape how young users understand and relate to the AI systems they frequently interact with.

Section 1. a defines artificial intelligence as a system that makes "predictions, recommendations, or decisions influencing real or virtual environments." The definition explains these systems "use machine- and human-based inputs to perceive real and virtual environments; abstract perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action."[17] This bill defines the technology by its ability to influence; however, it requires no instruction in how to navigate this influence.

IV. Policy Problem

A. Stakeholders

The primary stakeholders are New Jersey's public school students. These public school students are the population A4352 would teach, and they are already using the technological systems that the bill concerns. They are also at the age the American Psychological Association identifies as most susceptible to a chatbot's influence, a vulnerability tied to a stage of development in which the ability for independent judgment itself is still forming.[18] While students have the largest stake in the bill, they have the least voice in how it is written.

The New Jersey Board of Education and Commissioner of Education are the second stakeholder since they decide how A4352's requirement, Section 1 of the bill, should take effect. Under New Jersey administrative code, the state's learning standards are reviewed and revised on a fixed cycle. The Commissioner convenes an advisory panel to recommend modifications, presents the recommendations for the State Board's adoption, and publishes the revised standards to districts. A4352 places its artificial intelligence requirement in the New Jersey Student Learning Standards in Computer Science and Design Thinking and assigns the Commissioner responsibility for the sample activities and resources districts will use.[19] The standard that defines artificial intelligence literacy is set by these bodies and will eventually convert this definition into classroom materials. A definition that omits cognitive autonomy produces standards and resources that exclude it in every district of New Jersey.

The Computer Science and Design Thinking standards are among the New Jersey Student Learning Standards under revision, with proposed changes open for public comment. The skills this brief discusses would be added through this process, which means the question of artificial intelligence literacy is not set in stone in New Jersey.

School districts and teachers are the third stakeholder. Under New Jersey code, district boards are responsible for building curriculum and instruction on the adopted state standards.

Parents and the broader public are the final stakeholder. They are the constituency most concerned that students are deciding what to believe and how to think in conversation with systems built to influence their decisions. For the students discussed above who have no alternative source for this competency outside school, parents have limited ability to support them, since the skill of cognitive autonomy is not yet something most parents have been taught themselves.

The pattern of support for artificial literacy legislation described earlier, bipartisan and passed in states in different places across the political spectrum, gives reason to expect broad support in this legislation as well, from parents, educators, and researchers. Opposition is likely to come from concerns about cost and implementation, since districts might be wary of a new requirement introduced without dedicated funding, and some may interpret the skill of cognitive autonomy as hostility to the technology, or as discouraging students from using tools they will interact with. However,

cognitive autonomy is a skill that concerns opposing influence rather than avoiding the technology.

B. Risks of Indifference

Unguided use weakens students' critical thinking and their ability to evaluate what a system produces. Published in the peer-reviewed journal Societies, Gerlich's 2025 study of 666 participants found a significant negative correlation between frequent AI use and critical thinking scores.[20] A meta-analysis of seventeen studies presented at the 2025 CHI conference found that generative AI tools produce a significant positive effect on learning outcomes overall, but that the effect is significantly smaller for higher-order skills such as analyzing and evaluating than for lower-order skills such as recalling and applying.[21] The meta-analysis also found that instructional context mattered as much as the skill itself. Guided use, in which students received instruction and structure for how to engage with the tool, produced substantially larger learning effects than unguided use.[22] A4352 requires instruction in how to operate the systems; however, it does not require instruction in cognitive autonomy, which asks whether a student's conclusion was reached by one's own reasoning or supplied to them by the system built to agree, to sound complete, and to present itself as someone rather than a tool. A student can pass every check of media literacy teachers and still not be able to answer that question. The CHI researchers found that structured guidance around how students engaged with the tool produced larger gains than leaving them to use it freely. This finding shows how operational fluency alone is a measurably weaker approach, and that some further structure around a student's engagement with the system, of the kind cognitive autonomy provides, strengthens a student's ability to healthily interact with AI.

Developers use error-tracking platforms to monitor bugs in their software. When something breaks, developers use AI coding agents to help diagnose and fix it. In June 2026, the security firm Tenet disclosed a class of attack it calls Agentjacking.[23] An attacker submits a fake error report to the platform, written to look like a legitimate bug with a suggested fix. When a developer's AI coding agent later reviews the platform to diagnose an issue, it follows the fake instructions inside it and executes them as code. Researchers tested this attack on 100 organizations and succeeded 85% of the time. They also identified at least 2,388 organizations whose systems were exposed to it, ranging from individual developers to a Fortune 100 company.[24] Standard security tools did not catch this attack method because these tools are built to detect unauthorized activity such as unauthorized logins, files that should not be accessed, or code running without permissions. The developer had already given the AI agent permission to read the error-tracking platform and to write and execute code based on what it found. Professional developers had come to trust their tools enough that checking each output had stopped feeling necessary. If professionals with technical expertise lose the habit of verifying AI outputs, there is little basis for expecting the habit to develop unprompted in students.

In January 2026, a settlement resolved several lawsuits against Character.AI, which were brought by families in multiple states making claims about their children's relationships with the platform's chatbots.[26] Among these families were the Peraltas, whose 13-year-old daughter Juliana died by suicide in November 2023. Her family described her frequent usage as a dependency on a chatbot she had come to rely on for emotional support.[27] In August 2025, the family of

16-year-old Adam Raine sued OpenAI, alleging that ChatGPT, in months of conversation, came to function as his primary confidant, discouraging him from turning to his family, and continued engaging with him rather than redirecting him to help as his crisis escalated. OpenAI formally denied responsibility, stating in court filings that the system directed him to crisis resources more than one hundred times and that pre-existing risk factors and misuse of the product were the actual cause. The Raine case indicates that the anthropomorphic and sycophantic properties present in any conversational assistant- speaking in the first person, remembering what was said, staying continuously available- were enough on their own to produce the same relationship, without the engineered personas that Character.AI's platform provides. New Jersey students are already using general-purpose assistants like ChatGPT, and there is no specific instruction in A4352 that teaches a student to recognize the point at which that exchange stops being a tool and becomes a relational exchange.

C. Nonpartisan Reasoning

Artificial intelligence literacy has advanced under both parties. Enacted state laws come from Republican-led legislatures such as Utah and Idaho and from Democratic-led ones such as New York, whose RAISE Act took effect in 2025. At the federal level, the LIFT Artificial Intelligence Act was introduced as bipartisan legislation by Representative Tom Kean Jr., A New Jersey Republican, and Representative Gabe Amo, a Rhode Island Democrat.

A4352 is currently a single-sponsor bill, introduced by Democrat Assemblywoman Shanique Speight.

Different constituencies can support the change for different reasons:

1) Workforce readiness: The World Economic Forum's Future of Jobs Report 2025 found that analytical thinking has ranked as the single most essential core skill for the third consecutive edition of the report, with seven in ten companies rating it essential.[30]

2) Child wellbeing and educational quality: In-school instruction impacts every student regardless of what support is available in their homes. Most parents have not been taught to exercise judgement specifically in interaction with an AI system, since the ability is new enough that it predates no one's own schooling. Families with more resources can sometimes compensate on their own with informed relatives or closer supervision of how a child uses these tools. However, families without those resources have no alternative source for the ability of cognitive autonomy at all. A statewide curriculum is a method available to impact a child regardless of what their family can provide, which is why advocates for child wellbeing and educational equity have their own reason to support this amendment.

V. Tried Policy

Artificial intelligence literacy legislation is now widespread. As of mid-2026, FutureEd's legislative tracker counted 77 bills addressing AI in instruction across 27 states in the current session alone, with about a dozen already enacted.[31] The same gap in A4352 runs through these efforts, and they divide into three distinct approaches, none of which address cognitive autonomy.

The first and most common approach requires students to be technically fluent. Alabama's H.B.

329 requires all students to complete an approved computer science course that includes AI instruction in order to graduate high school, framed entirely around understanding and using the technology.[32]

The second approach includes evaluation, but limits that evaluation to media literacy. Utah's Digital Skills Amendments, enacted in 2026, require a course for seventh and eighth graders that lists artificial intelligence literacy alongside traditional and digital media literacy as required content.[33] This is closer to the concern of this brief than Alabama's approach, but it has the same limitation identified throughout this brief. Media literacy trains a student to evaluate a fixed source. However, a responsive system adapts to the student, mirrors the student's viewpoints, and tends to agree.

VI. Policy Options

A. Amend A4352 to add cognitive autonomy to the required instruction.

This option adds a fourth item to Section 1.c: instruction in evaluative thinking, the ability to assess a claim's quality on its own merits rather than accept it, and comparative validation, the ability to test a conclusion against sources beyond the one presenting it.[34] This option also adds cognitive autonomy to be a part of the sample activities and resources the Commissioner already supplies under Section 1.d to cover it. The State Legislature executes the amendment, since the bill remains in committee.[29] Once passed, the Commissioner and State Board set the standard and produce materials, and districts provide instruction to the students.

The first cost is in materials development. Section 1.d already requires the Commissioner to produce sample activities and resources for the bill's existing three requirements, regardless of this amendment. Using the already funded process to cover the fourth ability is a marginal addition to the deliverable the state is producing either way.

The second cost is teacher preparation. New Jersey's Department of Education reports 590 school districts and well over 100,000 classroom teachers statewide. Training all of them directly would be neither necessary nor affordable under any state professional development (PD) budget; national research puts average per-teacher professional development spending at roughly $8,300 a year for every subject a teacher covers.

States facing this problem train a small cohort of lead teachers per district, who train their colleagues during PD time the district already schedules. Applying this model here means training one lead teacher in each of New Jersey's 590 districts.

The state's FY2025 AI Innovation in Education Grant paid $75,000 per district for a broad AI initiative. Training one teacher on cognitive autonomy skills should cost a small fraction of that per district.

The benefit is that the requirement becomes statutory rather than optional, which means it reaches every district and is not dependent on the choices of the standards process alone. The disadvantage is that it requires legislative action, which is slower and less certain than administrative action and depends on the bill advancing. The main challenge is that Beckert's measurement instrument was built to assess cognitive autonomy in general rather than

autonomy against artificial intelligence specifically, so the standard-setting process would have to adapt its dimensions to this context. This work, however, is already done by the state when it turns any learning goal into a grade-appropriate standard.

B. Add cognitive autonomy through the Computer Science and Design Thinking standards revision.

This option would leave A4352 unchanged and instead add cognitive autonomy to the Computer Science and Design Thinking standards during the revision now underway, through the process in N.J.A.C. 6A:8-2.1.

The Commissioner and the State Board would carry this out, since they control the standards revision. Those affected are the same students and districts. The resources required are similar to Option 1, disregarding the legislative step.

The benefit is speed and certainty. The revision is already open for public comment, so cognitive autonomy could be added without waiting for a bill to pass, and it would take effect on the standards' own timeline. The disadvantage is stability. A requirement added through a standards revision can be removed or narrowed in the next revision cycle without legislative action, whereas a statutory requirement is more stable.

C. Test cognitive autonomy instruction in selected districts before statewide adoption.

This option would fund a small number of districts to develop and test cognitive autonomy instruction, evaluate the results, and use them to inform a later statewide standard. A pilot of this kind is typically sized to include a range of district types, so a workable target for New Jersey would be roughly 15-20 districts, chosen to include urban, suburban, and rural systems and a range of district wealth.

A pilot of 15-20 districts uses the same per-district training cost as the statewide rollout in Option A. At that rate, training lead teachers in the pilot districts costs $30,000 to $50,000. The pilot also needs a budget to evaluate whether the training worked, testing outcomes with an instrument adapted from Beckert's Cognitive Autonomy and Self-Evaluation Inventory. The brief estimates that cost at a similar scale to the training itself, roughly $30,000 to $50,000.

A smaller deployment would produce New Jersey-specific data on how cognitive autonomy can be taught and assessed, which would strengthen the case for statewide adoption and address the concern that the approach is untested. The disadvantage is that it reaches few students at first and delays the statewide change it is meant to inform.

One objection that applies to all three options is that cognitive autonomy is too abstract to teach or assess. Cognitive autonomy is an established psychological construct with a validated measurement instrument. Beckert's Cognitive Autonomy and Self-Evaluation inventory, developed across four studies and refined through repeated validation, breaks the construct into measurable criteria. Two of these criteria apply directly to countering artificial intelligence influence: evaluative thinking, the capacity to assess the quality of a claim rather than accept it, and comparative validation, the capacity to check a conclusion against sources other than the one

presenting it. A learning standard can be made utilizing these defined criteria.

VII. Conclusions

New Jersey is about to define, for every student in the state, what it means to be literate in artificial intelligence. A4352 defines literacy as the ability to operate the technology. The research on how these systems work shows why stopping there is not enough. These artificial intelligence systems adapt to their users, present themselves as human, tailor themselves to the individual, and tend toward sycophancy and agreement. The capacity to operate them does not automatically produce the ability to reason independently while using them.

The direction the industry is taking shows why cognitive autonomy is urgent. In his August 2026 essay "The Future is for Everyone," Mark Zuckerberg describes an agent built to "work 24/7 on your behalf" across a person's relationships, health, career, finances, and daily decisions.[38] Zuckerberg is specific about what this looks like in practice as well. The agent flags interesting information so that the user does not have to find it themselves; it prototypes ideas so a person does not have to develop them on their own from a blank page; it monitors sleep and delivers feedback on physical training, work a person would otherwise learn to do by observing their own bodies over time. It plans personalized recipes, replacing ordinary trial and error of learning to cook something new. Zuckerberg argues each example as a matter of convenience; however, each example is a form of cognitive offloading, the substitution of a system's effort for a person's own. Zuckerberg described that AI is meant to help a person "be a better friend and grow to become the person you aspire to be."[39] An agent that plans a person's meals, monitors their body, and filters information for them is a violation of cognitive autonomy, shaping who the user becomes. The industry is leaning towards these products, which will be produced and offered free to every person. Cognitive autonomy will become more necessary once every student has access to an agent built to work on their behalf around the clock, thus making it pertinent to maintain the skill of reasoning independently alongside technology.

New Jersey has already decided to teach artificial intelligence literacy. While this is true, the definition of artificial intelligence remains to be decided. New Jersey now has the choice to educate the youth to have the ability to use artificial intelligence systems and keep their one judgement while doing so.

Acknowledgment

The Institute for Youth in Policy wishes to acknowledge Taylor Beljon-Regen, Patrick Pickren, Tiffany Li, Asher Cohen, Paul Kramer, and other contributors for developing and maintaining the Fellowship Program within the Institute.

References

[1] New Jersey. Legislature. Assembly. 2026. Requires school districts to provide instruction on artificial intelligence. Assembly Bill 4352. https://pub.njleg.gov/Bills/2026/A4500/4352_I1.HTML

[2] Koyuturk, C., Guidotti, S., & Ognibene, D. (2026). The Hidden Cost of Contextual Sycophancy: An AI Literacy Intervention in Human–AI Collaboration. In Artificial Intelligence in Education (AIED 2026), Springer.

[3] Yazan, M., Situmeang, F. B. I., & Verberne, S. (2026). Personalized to Persuade: The Effects of Contextualization and Warmth on

Trust and Reliance in Conversational AI. HAI '26 (ACM).
[4] El-Sayed, S., et al. (2024). A Mechanism-Based Approach to Mitigating Harms from Persuasive Generative AI. arXiv:2404.15058.
[5] Long, D., & Magerko, B. (2020). What Is AI Literacy? Competencies and Design Considerations. CHI '20.
[6] Fang, C. M., et al. (2025). How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot
[7] Use: A Longitudinal Randomized Controlled Study. arXiv:2503.17473 (MIT Media Lab & OpenAI).
[8] Beckert, T. E. (2007). Cognitive Autonomy and Self-Evaluation in Adolescence. North American Journal of Psychology, 9(3), 579–594.
[9] Pew Research Center (2025). Teens, Social Media and AI Chatbots 2025.
[10] Reeves & Nass (1996). The Media Equation.
[11] American Psychological Association, "Artificial Intelligence and Adolescent Well-being: An APA Health Advisory" (Washington, DC: American Psychological Association, June 3, 2025)
[12] K.G.M. v. Meta Platforms, Inc., JCCP No. 5255 (Cal. Super. Ct., L.A. County, decided March 25, 2026).
[13] Claire Duffy, "Character.AI and Google Agree to Settle Lawsuits over Teen Mental Health Harms and Suicides," CNN Business, January 7, 2026, updated January 13, 2026.
[14] Duffy, "Character.AI and Google Agree to Settle."
[15] New Jersey Legislature, Assembly Bill 4352, 222nd Legislature, introduced February 19, 2026.
[16] New Jersey Legislature, Assembly Bill 4352, sec. 1.c.
[17] New Jersey Legislature, Assembly Bill 4352, sec. 1.a.
[18] American Psychological Association, "Artificial Intelligence and Adolescent Well-being."
[19] New Jersey Legislature, Assembly Bill 4352, sec. 1.b, sec. 1.d.
[20] Michael Gerlich, "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," Societies 15, no. 1 (2025): 6.
[21] Xiaodong Qu et al., "Generative AI Tools in Higher Education: A Meta-Analysis of Cognitive Impact," in Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25) (New York: ACM, 2025), Article 302.
[22] Qu et al., "Generative AI Tools in Higher Education."
[23] Ron Bobrov, Barak Sternberg, and Nevo Poran, "Agentjacking: How a Fake Sentry Bug Report Can Hijack Your AI Coding Agent" (Tenet Security, June 12, 2026).
[24] Bobrov, Sternberg, and Poran, "Agentjacking."
[25] Bobrov, Sternberg, and Poran, "Agentjacking."
[26] Duffy, "Character.AI and Google Agree to Settle."
[27] Torhoermann Law, "Character AI Lawsuit for Suicide and Self-Harm," accessed August 2026.
[28] Raine v. OpenAI, Inc., Case No. CGC-25-628528 (Cal. Super. Ct., San Francisco County, filed August 26, 2025).
[29] Raine v. OpenAI, Inc., Case No. CGC-25-628528.

[30] World Economic Forum, Future of Jobs Report 2025 (Geneva: World Economic Forum, 2025).
[31] FutureEd, "Legislative Tracker: 2026 State AI in Education Bills," Georgetown University McCourt School of Public Policy, updated July 13, 2026.
[32] Alabama Legislature, House Bill 329 (2024).
[33] Utah Legislature, House Bill 218, Digital Skills Amendments (2026).
[34] Beckert, "Cognitive Autonomy and Self-Evaluation in Adolescence."
[35] New Jersey School Boards Association, "New Jersey Department of Education Announces Grant Awards to Support Artificial Intelligence in Education," January 22, 2025.
[36] New Jersey School Boards Association, "Governor Signs State Budget for Fiscal Year 2026 into Law."
[37] Beckert, "Cognitive Autonomy and Self-Evaluation in Adolescence."
[38] Mark Zuckerberg, "The Future Is for Everyone: The Path to a Positive AI Future" (Meta, August 10, 2026).
[39] Mark Zuckerberg, Facebook video statement, July 30, 2025, reported in Johana Bhuiyan, "Zuckerberg Says AI Superintelligence Is 'in Sight,' Touts 'New Era of Personal Empowerment,'" NBC News, July 30, 2025.