ChatGPT

AI for Student Social and Emotional Learning in K-12

Executive summary

Used well, AI can strengthen student social and emotional learning by helping schools do three things that are already hard at scale: give students timely reflection and coaching, help educators notice patterns earlier, and reduce the planning burden of high-quality SEL instruction. Used badly, the same technologies can turn SEL into surveillance, over-automate sensitive judgments, or create false confidence in tools that have not actually been shown to improve student well-being. The practical conclusion is that AI is most defensible in K-12 when it augments adult relationships, not when it substitutes for them. [1]

The evidence base is promising but uneven. Traditional schoolwide SEL programs have a strong research record. Digital SEL interventions also show positive effects, especially for social-emotional skills, but the strongest modern school-based evidence for generative AI is still thin and mixed. A 2026 meta-analysis of digital SEL interventions found positive effects on social-emotional skills and behavior, but effects on affect and attitudes weakened after publication-bias correction. A 2026 school-based randomized trial of generative AI for self-regulated learning found modest motivational benefits but no clear gains in effort or domain learning over control. Meanwhile, research on affective pedagogical agents and educational chatbots suggests small motivational benefits and improved satisfaction, but not a reliable basis for high-stakes SEL decisions. [2]

For U.S. districts, the highest-value near-term use cases are teacher-facing curriculum design, structured student reflection or coaching inside bounded classroom tasks, low-stakes student check-ins, and analytics that help adults coordinate support. The weakest and riskiest use cases are automated emotion recognition from faces or voices, always-on monitoring, opaque student "risk scores," and any student-facing agent that blurs the line between tutor, friend, and therapist. Those categories carry the greatest risks of bias, privacy intrusion, cultural mismatch, discrimination, and over-reliance. [3]

The market is ahead of the evidence. Most reviewed products are AI-enabled workflows, not rigorously validated SEL interventions. Vendor claims often emphasize safety, efficiency, or usage metrics, while public evidence on subgroup performance, false positives, and long-term student outcomes is limited. Districts should therefore treat procurement as a governance decision first and a software decision second, using U.S. Department of Education [4] guidance, the National Institute of Standards and Technology [5] AI risk framework, and clear human-in-the-loop protocols as the baseline. , , and the are the most useful starting documents. [6]

SEL foundations and what AI can realistically add

The most widely used U.S. school framework, from CASEL [7], organizes SEL around five competencies: self-awareness, self-management, social awareness, relationship skills, and responsible decision-making. A complementary lens from the OECD [8] frames social and emotional skills as teachable patterns of thoughts, feelings, and behaviors linked to academic success, health, and later life outcomes. Those definitions matter because they imply that SEL is not just “feeling better”; it is a set of capacities practiced across relationships, routines, and decisions. [9]

That framing also clarifies where AI fits. AI can scaffold reflection, practice, feedback, and pattern detection. It is relatively well suited to supporting self-awareness and self-management through check-ins, journaling prompts, coping strategy rehearsal, goal-setting, and metacognitive coaching. It can also help with responsible decision-making through scenario-based discussion and guided questioning. It is less well suited to replacing the human conditions that actually build social awareness and relationship skills: trust, attachment, co-regulation, classroom belonging, and culturally grounded understanding. In other words, AI can help students rehearse SEL, but schools still have to supply the relationships in which SEL becomes real. [10]

A good district test is simple: if the tool primarily expands adult capacity, improves the timeliness of support, or gives students structured opportunities to reflect and practice, it is more likely to be appropriate. If it claims to infer emotion, diagnose need, or form a quasi-relationship with students, scrutiny should increase sharply. [11]

AI use cases for K-12 SEL

Personalized coaching is the clearest use case. LLM-based systems can prompt students to identify goals, name emotions, break down problems, rehearse coping strategies, and reflect on choices. The strongest current application is not “AI therapy,” but bounded coaching linked to school tasks such as frustration tolerance during writing, planning during project work, or reflection after peer conflict. The most credible products keep teachers or trained staff in the loop and preserve a clear boundary between educational coaching and clinical care. [12]

Conversational agents and tutors can also facilitate SEL indirectly. Educational chatbot research shows that chatbots are used not only as teaching agents, but often as peer-like agents, and that personalized or collaborative designs can improve learning and satisfaction. In practice, this means AI can support reflection, perspective-taking, rehearsal of restorative conversations, and guided problem-solving. The limitation is that most studies evaluate learning or engagement, not durable SEL transfer, and few are rigorous U.S. K-12 SEL trials. [13]

Teacher analytics and early warning systems are another major category. AI can synthesize attendance, behavior, survey responses, writing, and help-seeking patterns so that teachers or student support teams can spot changes sooner. In a careful implementation, this can reduce missed signals and make supports faster and more consistent. In a careless implementation, it becomes indiscriminate monitoring. The difference lies in scope, proportionality, transparency, and whether flagged patterns trigger human review rather than automated action. [14]

Curriculum design is the most mature and lowest-risk category. Teacher-facing AI can draft discussion prompts, role-play scenarios, reflection journals, multilingual explanations, restorative conversation stems, and differentiated lesson plans. This is attractive because it directly addresses teacher time constraints. It also has the cleanest governance profile because schools can use AI without exposing large volumes of sensitive student data. The drawback is that time-saving is not the same as student impact; districts should not confuse teacher efficiency with evidence of SEL improvement. [15]

Peer collaboration facilitation is promising but underdeveloped. AI can form groups, suggest norms, structure peer feedback, or coach discussion quality. Systematic reviews show that a meaningful share of educational chatbots are designed as peer agents and that collaborative learning patterns do appear in the literature. But direct evidence that AI improves empathy, belonging, or conflict resolution in regular K-12 classrooms remains limited. [16]

Assessment of SEL skills should be approached conservatively. The most defensible measures remain triangulated student self-report, teacher observation, and work artifacts. AI can help summarize or score reflective writing, organize check-in data, and surface trends. It should not, at present, be treated as a reliable stand-alone assessor of a child's emotional state or social competence. That caution is especially important for emotion recognition systems that rely on cameras, voice, or behavioral traces; those systems are among the least mature and most ethically fraught uses schools could adopt. [17]

Emotion recognition deserves a separate judgment: it is the most controversial AI-for-SEL category and, for most districts, the least ready for routine use. The research case is weaker than the marketing case, and the legal and ethical risks are higher because inference errors can map directly onto discipline, surveillance, stigma, or differential treatment. In the current U.S. K-12 environment, camera-based or voice-based emotional inference should be treated as high risk and, if used at all, limited to low-stakes, aggregate, opt-in research or product pilots with explicit oversight. [18]

Evidence base

The central research fact is that conventional SEL is much better validated than AI-enabled SEL. A large recent meta-analysis of universal school-based SEL found significant benefits for academic achievement, school functioning, social and emotional skills, behaviors, attitudes, and school climate. That gives districts a strong "business as usual" benchmark. AI should be judged against that benchmark, not against the absence of SEL supports. [19]

The best current synthesis of digital SEL itself is a 2026 meta-analysis of 37 studies with 4,742 K-12 students. It found positive effects on social-emotional skills (SMD 0.631, reduced to 0.369 after trim-and-fill) and prosocial/antisocial behavior (SMD 0.285, trim-and-fill 0.353). Effects on affect and attitudes were initially positive (SMD 0.267) but became non-significant after publication-bias correction. The practical takeaway is that digital supports can help, but the effects are heterogeneous, and the most robust gains are in skills and behavior rather than broadly improved feelings or attitudes. [20]

The most directly relevant recent school-based GenAI trial is the 2026 randomized controlled trial of 371 students in grades 7-9 using theory-informed generative AI supports for self-regulated learning during physics or English lessons. Students in the "utility value" condition showed a better trajectory on perceived usefulness of learning content than a strategy-prompting condition, but there were no statistically significant advantages over the control condition for effort, knowledge, or elaboration strategy use. That is a useful caution: GenAI may help preserve or scaffold motivation, but its classroom effects are not automatically broad or deep. [21]

Evidence on emotionally expressive instructional agents is somewhat stronger than the evidence on open-ended general chatbots. A meta-analysis of affective pedagogical agents found small but positive effects on positive emotions (g = 0.26), intrinsic motivation (g = 0.26), retention (g = 0.26), and transfer (g = 0.34). That suggests design matters: emotionally supportive cues can help some students stay engaged. But these are still small effects, typically measured in controlled multimedia environments rather than whole-school SEL implementations. [22]

For special populations, socially assistive robots remain one of the clearer evidence pockets. A 2022 meta-analysis of robot-mediated interventions for children and youth on the autism spectrum found a significant effect on social functioning (g = 0.35) across randomized trials, but not on emotional or motor outcomes, in part because those outcome pools were small. This is informative for targeted SEL and social-skills work, especially in special education, but it does not support broad claims about robots improving SEL for all students. [23]

The weakest point in the evidence base is the area many vendors emphasize most: mental-health-adjacent chat and passive detection. A qualitative study of secondary students' views of AI therapists found strong concerns about inaccuracy, privacy, dependency, and the quality of emotional support, and the American Psychological Association [24] has called for safeguards against manipulation, unsafe advice, and over-reliance in adolescent AI use. This does not mean schools should never use chat-based wellbeing supports; it means they should deploy them as bounded, supervised supports rather than as substitutes for trusted adults. [25]

The research limitations are consistent across categories: many studies are small, short-term, outside the United States, focused on narrow populations, or evaluate user satisfaction rather than student outcomes. Vendor studies frequently lack comparison groups, subgroup error analysis, or transparency about how models behave across language, disability status, race, or crisis contexts. Districts should therefore demand evidence of effectiveness, precision, and equity by subgroup before scaling. [26]

K-12 platform landscape

The table below compares representative AI-enabled or algorithmically driven products that U.S. districts are likely to encounter under the broad banner of "AI for SEL." Inclusion is descriptive, not an endorsement. In several cases, the reviewed sources were strong on feature descriptions and privacy language but weak on independent efficacy evidence.

Platform	Primary SEL-related use	Age / grades	Main data collected	Privacy safeguards	Cost model	Evidence of effectiveness	Accessib equity conside
	Student-facing coaching/ tutoring spaces; teacher dashboard with insight summaries and alerts	K-12	Student interactions, account/ roster data, interaction-derived insights and sentiment-style summaries	School/ district control of information; family opt-out, access, correction, deletion; trust and privacy center	Free for teachers; paid team/ school/ district plans	Publicly claims ESSA Tier 4 logic-model validation; reviewed sources did not surface a peer-reviewed K-12 SEL outcome study	Purpose classroo guardra modera reviewe sources not surf public V ACR
	Teacher-facing curriculum design; student-facing guided AI rooms and writing/ learning support	K-12	Prompts, outputs, account data, usage/ engagement data	District-ready privacy docs and DPAs; states it does not sell PII and does not train on PII; FERPA/ COPPA/SOC 2-aligned claims; guardrails for student use	Free for teachers; enterprise/ district contracts	Strong evidence of teacher time savings in vendor reports; reviewed sources did not surface a peer-reviewed SEL impact study	WCAG-a with VP/ ACR, scr reader support equity p guidanc
	Personalized tutoring, writing feedback, teacher tools, progress summaries; can support reflection and self-regulation in instruction	Teacher access broadly; student classroom access through school/district implementations	Chat history, user prompts, moderation events, progress/ activity data	Adult visibility into student chat history; flagged-content alerts; guidance not to enter PII; official privacy and accessibility statements	Free for teachers; $4/ month for learner/ parent accounts; district pricing by request	Strong educational positioning and safety controls; reviewed sources did not surface a peer-reviewed K-12 SEL outcome study	Accessib settings speech- text/rea aloud feature languag settings oversig protect but redu student privacy

Platform	Primary SEL-related use	Age / grades	Main data collected	Privacy safeguards	Cost model	Evidence of effectiveness	Accessib equity conside
	Text-based wellbeing support for students, with human wellbeing coaches supported by AI tools	Adolescent / secondary-oriented in reviewed sources	Chat content and limited safety-related personal data	States chats are confidential, does not monitor social media, and only shares with school for serious safety concerns	District partnership / demo-led; public list pricing not surfaced	Reviewed sources show strong product positioning; no peer-reviewed school-based efficacy study was identified in reviewed sources	Text-bas access r broader in couns short di high importa human escalati clear cli bounda
	Student check-ins, surveys, MTSS coordination, early warning, AI-powered student summaries/ intervention support	K-12	SIS data, attendance, behavior, assessments, survey/ check-in responses, intervention records	Privacy and client information policies; district control over surveys/ questions	Demo/ quote-based; public list pricing not surfaced	Vendor reports show 60% of tracked interventions met stated goals on-platform, but this is not causal efficacy evidence	Accessib surveys languag stated accessib best pra help wit inclusio requires careful interpre of surve
	Daily wellness check-ins, short coping/skill activities, dashboards for staff	K-12	Emoji-based mental/ emotional/ energy/ physical/ social check-ins, custom comments, identifiers, staff/district data	FERPA-oriented design claims; U.S.-based servers; no sale of student data to marketers; role-based access for staff	Demo/ contract; public list pricing not surfaced	Claims ESSA Tier 3 evidence and broad adoption; reviewed sources did not surface recent peer-reviewed RCT results for the current product	Short, d light ch may imp particip self-rep remains sensitiv culture, and res honesty

Platform	Primary SEL-related use	Age / grades	Main data collected	Privacy safeguards	Cost model	Evidence of effectiveness	Accessib equity conside
	Early warning for self-harm, suicide, and violence through online-activity monitoring	Districtwide K-12 deployment; strongest fit is typically secondary	Web searches, browsing, documents, email, contextual alert data, escalation settings	Product privacy policy; customizable escalation rules; professional services; accessibility release notes	Quote-based with volume discounts and professional services	Vendor cites safety and prevention metrics; reviewed sources did not surface a peer-reviewed public efficacy study or subgroup error analysis	Accessib work is docume but surveilla burden false po can fall unequa margina student
	AI-supported student wellness monitoring using browsing data, “wellness levels,” and staff inputs	K-12	Browsing/activity data via extension, student identifiers, wellness levels, dashboard data	Student privacy policy, rights exercised through schools, enhanced privacy mode	Demo/quote-based; public list pricing not surfaced	Public product materials emphasize capability; reviewed sources did not surface peer-reviewed public efficacy evidence	Enhanc privacy is helpfu the over surveilla model r the sam proport and bias concern other monitor tools

A district should notice the pattern. The tools with the cleanest value proposition for SEL are usually the least invasive: teacher planning tools, structured classroom reflection tools, and bounded check-ins. The tools with the most intrusive data collection often have the weakest public efficacy evidence. [35]

Ethics, legal frameworks, and mitigation strategies

The most important ethical issue is not whether AI can sound empathic. It is whether schools can use it without turning vulnerable students into data exhaust. Research on school surveillance companies shows that many vendors monitor outside school hours, use AI to flag concerning activity, and in some cases create student risk scores, all with limited transparency about how the systems work or whom they misclassify. That is particularly concerning when the monitored signals are proxies for distress, identity, or family circumstance rather than direct evidence of need. [36]

Bias and discrimination are not hypothetical. The Department's Office for Civil Rights has made explicit that federal civil-rights protections apply when AI contributes to discrimination in educational programs and activities. That includes race, national origin, sex, and disability discrimination under Title VI, Title IX, Section 504, and the ADA. For SEL tools, that means districts should assume that any system used for flagging, screening, triage, communication, or discipline has civil-rights implications. If the district cannot explain the model's objectives, inputs, error modes, and override process, it should not be used for high-stakes decisions. [37]

The core U.S. privacy framework has three especially relevant pieces. FERPA protects education records and the personally identifiable information schools maintain or that vendors maintain as school officials. COPPA applies to operators of online services collecting personal information from children under 13; schools can consent on parents' behalf only for school-authorized educational purposes and not for unrelated commercial uses. PPRA governs certain surveys, analyses, or evaluations that touch protected topics such as mental or psychological problems, sex behavior, political beliefs, and family relationships. Any SEL screener or wellbeing check-in program should therefore be reviewed not just for pedagogical fit, but for lawful notice, consent, access, deletion, retention, and survey protections. [38]

Where EU data or international schools are involved, GDPR adds further requirements. The European Commission's guidance emphasizes parental consent thresholds that vary by member state between ages 13 and 16, reasonable efforts to verify consent, and child-facing notices in clear and plain language. It also highlights that preventive or counselling services offered directly to a child may be treated differently because they serve the child's best interests. These principles are useful even for U.S. districts, because they push schools toward comprehensible notices, minimal data collection, and proportionate age assurance. [39]

The safest mitigation strategies are straightforward. First, minimize data: do not collect facial video, full keystroke streams, location, or private-device browsing unless there is a compelling, documented reason. Second, narrow the purpose: tools bought for classroom reflection should not silently become discipline or threat-assessment tools. Third, preserve human review: no alert, score, or inferred emotion should be the sole basis for intervention, discipline, or special-education referral. Fourth, test equity explicitly: require vendors to report subgroup precision, recall, false positives, and content performance across languages and disability-related accommodations. Fifth, set retention and deletion rules up front. Sixth, make student and family notice understandable, plain-language, and age-appropriate. [40]

Schools should also set a sharper boundary around relational design. Student-facing AI should support learning, not simulate replacement friendship or therapy. That is consistent with current educator guidance, adolescent-wellbeing guidance, and the stronger safety posture taken by some education vendors themselves. For SEL, this means avoiding anthropomorphic or dependency-inducing designs unless the use case is clinically supervised and explicitly justified. [41]

Implementation guidance for schools

The most effective implementation model is a bounded pilot with a sharply defined problem of practice. Start with one student population, one primary use case, one data-governance path, and one human response protocol. Examples include "use AI-supported reflection prompts to improve self-management in grade 8 writing," or "use low-stakes wellness check-ins plus counselor review to improve referral timeliness in one high school." Avoid combined pilots that mix tutoring, mental health chat, early warning, and discipline; they create too many variables to evaluate responsibly. [42]

```
flowchart TD
    A[Define SEL target problem and student group]
    B[Review law, privacy, accessibility, and civil-rights constraints]
    C[Screen vendors for evidence, bias, data minimization, and human oversight]
    D[Configure pilot with DPA, retention rules, notice/consent, and escalation
protocol]
    E[Train teachers, counselors, leaders, and family-facing staff]
    F[Run bounded pilot with baseline and comparison metrics]
    G[Review outcomes, incidents, subgroup equity, and user trust]
    H[Scale, revise, or stop]

    A --> B --> C --> D --> E --> F --> G --> H
    D --> I[Student and family communication]
    D --> J[Accessibility and accommodation check]
    F --> K[Human review of alerts and interventions]
```

This sequence reflects the common logic across the DOE toolkit, NIST risk guidance, FTC school-COPPA guidance, and early-warning implementation advice: governance first, staff capability second, classroom use third, scale last. [43]

Stakeholder roles should be explicit. The superintendent or cabinet sponsor sets the purpose and red lines. Academic leadership determines instructional fit. Student support leaders define triage and referral. Privacy/legal staff review FERPA, COPPA, PPRA, contracts, and retention. IT verifies SSO, data flow, logs, and access controls. Special education and multilingual learner teams review accessibility, accommodations, and language equity. Teachers co-design classroom use cases. Counselors and psychologists own escalation, not the vendor. Students and families should participate in design review before launch, not merely receive notice after procurement. [44]

Professional development should be practical rather than generic. Staff need to know what prompts to use, what data not to enter, how to interpret student outputs, when AI outputs are unreliable, how to document interventions, and what to do when a system flags self-harm, bullying, or identity-related distress. Districts should train to the failure modes, not just to the happy path. [45]

Metrics for success should span student outcomes, operations, and trust. For students, measure self-regulation or coping proxies, belonging/help-seeking indicators, referral uptake, and trends in absenteeism or behavior where relevant. For staff, measure time saved, response timeliness, and whether adults actually acted on the insights. For system quality, measure false positives, false negatives where detectable, subgroup disparities, accessibility issues, opt-out rates, and family concerns. If the district cannot observe both benefit and harm metrics, it is not really evaluating the pilot. [46]

A reasonable timeline is four to six weeks for governance and procurement review, two to four weeks for technical setup and staff training, eight to twelve weeks for classroom or school pilot use, and two to four weeks for evaluation and board-level decision. Early-warning or wellbeing-monitoring pilots should not go live until after-hours protocols, backup staffing, and responsible contacts are fully tested. [47]

The budget question is best framed as three scenarios rather than one number, because most district tools are quote-based and district size varies greatly.

Scenario	Likely scope	Typical stack	Illustrative planning range
Low	One school or one grade-band pilot	Teacher-facing AI planning tools, bounded student reflection use, light check-in or survey layer, internal PD	**$15,000-$50,000** annually
Medium	Three to five schools or one secondary feeder pattern	District licenses for student-facing support plus MTSS/check-in workflows, SSO/SIS setup, PD, pilot evaluation	**$75,000-$250,000** annually
High	Districtwide deployment with early warning and coordinated support	Enterprise monitoring or wellbeing platform, integrations, external privacy/accessibility review, formal evaluation, ongoing training	**$300,000-$1,000,000+** annually

These are planning estimates synthesized from public pricing signals and the fact that many K-12 enterprise vendors do not post standard district list prices. In reviewed sources, SchoolAI, MagicSchool, and Khanmigo expose free or low-cost teacher-facing entry points, while district platforms such as Panorama, GoGuardian, Sonar, Rhithm, and Securly are primarily demo-led or quote-based. [48]

Open questions and bottom line

The biggest unresolved questions are empirical, not technical. Which tools measurably improve SEL outcomes beyond existing school-based SEL programs? Which tools reduce staff burden without increasing surveillance or false positives? Which vendors can demonstrate equitable performance across disability status, language background, race, gender, and socioeconomic status? And how often do students start to substitute AI interactions for human help, especially when they are lonely or distressed? Public answers to those questions are still scarce. [49]

The bottom line is that AI can facilitate SEL in K-12, but only under a narrow and disciplined theory of action. The best uses are structured reflection, coaching, differentiated teacher support, and adult-facing coordination tools that make human care more timely. The worst uses are opaque emotional inference, broad surveillance, and any design that encourages students to treat AI as a confidant rather than as a bounded educational tool. Districts that lead with governance, evidence, privacy, accessibility, and human judgment can gain real value; districts that lead with novelty or crisis-driven procurement are much more likely to create new risks than to solve old problems. [50]

---

[1] [20] [26] [49] https://link.springer.com/article/10.1186/s40359-026-04434-4
https://link.springer.com/article/10.1186/s40359-026-04434-4

[2] [19] https://www.sciencedirect.com/science/article/pii/S2773233924000032
https://www.sciencedirect.com/science/article/pii/S2773233924000032

[3] [37] https://www.ed.gov/media/document/avoiding-discriminatory-use-of-artificial-intelligence-112274.pdf
https://www.ed.gov/media/document/avoiding-discriminatory-use-of-artificial-intelligence-112274.pdf

[4] [34] https://www.securly.com/aware
https://www.securly.com/aware

[5] [29] https://www.khanmigo.ai/pricing
https://www.khanmigo.ai/pricing

[6] [11] [17] [24] [35] [40] [42] [43] [44] [45] [46] [50] https://files.eric.ed.gov/fulltext/ED661924.pdf
https://files.eric.ed.gov/fulltext/ED661924.pdf

[7] [25] https://pmc.ncbi.nlm.nih.gov/articles/PMC11939552/
https://pmc.ncbi.nlm.nih.gov/articles/PMC11939552/

[8] [10] [12] [21] https://link.springer.com/article/10.1007/s10648-026-10133-8
https://link.springer.com/article/10.1007/s10648-026-10133-8

[9] https://casel.org/fundamentals-of-sel/what-is-the-casel-framework/
https://casel.org/fundamentals-of-sel/what-is-the-casel-framework/

[13] [16] https://link.springer.com/article/10.1007/s10639-022-11177-3
https://link.springer.com/article/10.1007/s10639-022-11177-3

[14] [31] https://www.panoramaed.com/
https://www.panoramaed.com/

[15] https://www.magicschool.ai/magicschool
https://www.magicschool.ai/magicschool

[18] [36] https://www.jmir.org/2025/1/e71998
https://www.jmir.org/2025/1/e71998

[22] https://www.sciencedirect.com/science/article/abs/pii/S1747938X22000756
https://www.sciencedirect.com/science/article/abs/pii/S1747938X22000756

[23] https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0269800
https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0269800

[27] https://schoolai.com/
https://schoolai.com/

[28] https://www.magicschool.ai/
https://www.magicschool.ai/

[30] https://www.sonarmentalhealth.com/
https://www.sonarmentalhealth.com/

[32] https://rhithm.app/faq/
https://rhithm.app/faq/

[33] https://www.goguardian.com/beacon
https://www.goguardian.com/beacon

[38] https://studentprivacy.ed.gov/ferpa
https://studentprivacy.ed.gov/ferpa

[39] https://commission.europa.eu/law/law-topic/data-protection/rules-business-and-organisations/legal-grounds-processing-data/are-there-any-specific-safeguards-data-about-children_en
https://commission.europa.eu/law/law-topic/data-protection/rules-business-and-organisations/legal-grounds-processing-data/are-there-any-specific-safeguards-data-about-children_en

[41] https://www.magicschool.ai/blog-posts/student-safety-companionship
https://www.magicschool.ai/blog-posts/student-safety-companionship

[47] https://www.panoramaed.com/blog/k-12-early-warning-system-implementation
https://www.panoramaed.com/blog/k-12-early-warning-system-implementation

[48] https://schoolai.com/pricing/
https://schoolai.com/pricing/