Same Tools, Different Rules: Generative AI Acceptable-Use Policies Across Fifteen Sectors, with a Complete Policy Matrix for K-12 Schools

Samo Plibersek
Alpha Quantum · Shadow AI Tools
samo@alpha-quantum.com

September 2026

Abstract

Guidance on generative AI is written sector by sector, yet the tools employees and students reach for are the same everywhere. We formalise fifteen sector profiles, from K-12 schools to defense, as layered decision rules over a functional taxonomy of AI tools, and apply every profile to the same register of 11,890 active AI tools. The share of tools a profile blocks ranges from 10.0% for marketing and creative agencies to 81.6% for defense and critical infrastructure. "Allow with controls" is the most common verdict in ten of fifteen sectors, so the practical question is rarely a choice between banning and permitting AI. Two sectors picked at random disagree on 41.3% of tools on average; only 13.4% of tools receive the same verdict everywhere, while 51.9% are blocked in one sector and allowed in another. Profiles cluster: healthcare and insurance differ on 3.2% of tools, K-12 schools and universities on 76.0%. We present the complete policy matrix for one sector, K-12 schools, with a rationale for each of the 18 categories and all 165 subcategories. Under this profile 71.5% of tools are blocked, and 35.4% of verdicts are set by a subcategory rule rather than a category default. We argue that sector policies should be published as explicit, testable rules at the subcategory level, and we release the aggregate statistics as open data.

Keywords: generative AI, acceptable use policy, AI governance, K-12 education, shadow AI, sector regulation

1 Introduction

Most organisations now have staff using generative AI, often without approval. In a 2024 survey of knowledge workers, three quarters reported using AI at work, and most of them brought their own tools (Microsoft and LinkedIn, 2024). Unapproved use is a data-protection problem: an analysis of 1.6 million workers found that 11% of what employees pasted into ChatGPT was confidential (Cyberhaven, 2023), and a large electronics manufacturer restricted generative AI after engineers uploaded source code (Gurman, 2023). The pattern is familiar from earlier research on shadow IT (Silic and Back, 2014).

The answer to "which AI tools should we allow?" depends on who is asking. A school must protect minors and academic integrity (United States Congress, 2000, 1998, 1974). A hospital must keep protected health information out of third-party systems (U.S. Department of Health and Human Services, 2013). A bank must preserve records and prevent leaks of client data (United States Congress, 1999; FINRA, 2024; European Parliament and Council of the European Union, 2022). A law firm must protect client confidences (American Bar Association Standing Committee on Ethics and Professional Responsibility, 2024), a federal agency must inventory and govern its AI use (Office of Management and Budget, 2024), and a defense contractor must

keep controlled technical data within the country (U.S. Department of State, 2024). Horizontal frameworks such as the NIST AI Risk Management Framework (National Institute of Standards and Technology, 2023) and the EU AI Act (European Parliament and Council of the European Union, 2024) add obligations on top, but leave the choice of individual tools to the organisation.

Sector guidance therefore exists, but it is written in prose. It says what must be protected, not which of thousands of tools are acceptable. This paper asks what happens when sector guidance is translated into explicit rules and applied to the same population of tools. We contribute:

- a layered decision model that turns a sector profile into a verdict (block, allow with controls, allow) for any tool described by a functional taxonomy;
- the verdict distributions of fifteen sector profiles over 11,890 AI tools, with an analysis of where sectors agree and where they differ;
- the complete matrix for one sector, K-12 schools, with a rationale for every category and subcategory, as an example of how such a policy can be documented and audited.

The verdicts are recommended defaults derived from the rules, not observations of what organisations actually do.

2 Background

Sector guidance on generative AI. Education has the most developed guidance. UNESCO recommends protecting data privacy and setting an age limit for independent conversations with generative AI (Miao and Holmes, 2023). The U.S. Department of Education asks schools to keep humans in the loop and to align AI use with educational goals (U.S. Department of Education, Office of Educational Technology, 2023). The EU AI Act lists AI systems that determine access to education or evaluate learning outcomes as high-risk (European Parliament and Council of the European Union, 2024). Design codes for children's services add data minimisation and high-privacy defaults (Information Commissioner's Office, 2020). In finance, regulators stress that existing rules on supervision, records and third-party risk apply to AI unchanged (FINRA, 2024; Financial Stability Board, 2024; European Parliament and Council of the European Union, 2022). Legal ethics guidance focuses on confidentiality and informed consent (American Bar Association Standing Committee on Ethics and Professional Responsibility, 2024).

AI in the classroom. Research on large language models in education sees real potential for tutoring and feedback, together with risks to assessment and to the development of writing skills (Kasneci et al., 2023; Cotton et al., 2024). Detection tools are an unreliable remedy: they misclassify both human and AI-written text (Weber-Wulff et al., 2023) and are biased against non-native writers (Liang et al., 2023). A school policy therefore cannot rely on detection alone and must decide which tools pupils can reach.

3 Data and method

Tools. The population is the register of 11,890 active AI tools compiled by AI Tools Blocklist on 26 September 2026 (Plibersek, 2026a); the taxonomy and the register's training-policy data are described in a companion paper (Plibersek, 2026b). Tools were classified by a machine-learning pipeline into a fixed taxonomy of 18 categories and 165 subcategories; labels were then reviewed tool by tool against each tool's own homepage in an LLM-assisted review with spot checks by the author. A tool has one primary category and usually several subcategory labels. Tools that cannot be placed would form a nineteenth bucket, *Uncategorized*; after the final review of the

register no active tool remains in it. The register also records, per tool, whether it is AI-native or an established product with added AI (AI-enabled), a risk rating, data-sovereignty exposure, whether the tool serves an abusive purpose, and its stated policy on training with customer inputs.

Profiles. A profile encodes one sector. It has a strictness level (2 for regulated sectors and minors, 1 standard, 0 permissive), a default verdict for each category, and optional verdicts for subcategories. The fifteen profiles were written from the sector sources cited above.

Decision layers. A tool's verdict is set in five steps.

1. *Hard blocks*, identical in every sector: abusive-purpose tools, tools with a high risk rating and tools with high data-sovereignty exposure are blocked.
2. *Baseline.* If any of the tool's subcategories has a rule in the profile, the strictest such rule applies. Otherwise a cross-sector subcategory rule applies if one exists (for example, voice cloning is blocked and deepfake detection is allowed everywhere). Otherwise the category default applies.
3. *Tightening* by attributes. Training on inputs by default moves the verdict one step, or two in strict profiles. Opt-out training moves it one step in standard and strict profiles. In strict profiles, an allowed tool with no stated training policy moves to controls, and medium risk or medium sovereignty exposure moves the verdict one step.
4. *AI-enabled relief.* An established product with added AI features is never blocked on category grounds alone; a block becomes controls.
5. *Enterprise relief.* If a block was caused by training on inputs and the vendor states that its business tier does not train, the verdict becomes controls.

The rules are deterministic, so every verdict can be traced to the layer that produced it.

4 Fifteen sectors, one population

Table 1 and Figure 1 give the verdict shares for each profile. The strictest profile, defense, blocks 81.6% of tools; the most permissive, agencies, blocks 10.0% and allows more than half. Strictness explains the broad ordering but not all of it: government and banking share strictness level 2 with K-12, yet block about half of all tools (49.3% and 46.5%) against 71.5% in K-12, because their category defaults allow business software under controls.

Controls is the modal verdict in ten sectors. Only defense and K-12 block a majority of tools; government is split almost evenly (49.3% block, 50.1% controls). Software companies, agencies and, narrowly, universities allow more than they control. For most organisations the work of an AI policy is therefore in defining controls: which accounts, which data, which logging.

How much do sectors disagree? For each pair of profiles we count the tools that receive different verdicts. Across all 105 pairs the mean disagreement is 41.3%. Only 13.4% of tools receive the same verdict in all fifteen profiles, mostly hard blocks and cross-sector rules; 51.9% are blocked somewhere and allowed somewhere else. Figure 2 shows selected pairs. Regulated sectors that protect personal or client data form a tight cluster: healthcare and insurance differ on 3.2% of tools, banking and insurance on 4.0%, healthcare and pharma on 4.6%. At the other end, agencies and defense differ on 83.2%.

Two pairs are instructive. K-12 schools and universities both serve learners, yet disagree on 76.0% of tools: 3,337 tools blocked for pupils are allowed outright for adult students. K-12 and defense, in contrast, differ on only 11.7% of tools, although their reasons have nothing in common.

Table 1: Verdict shares per sector profile over the same 11,890 AI tools. Strictness: 2 = regulated or minors, 1 = standard, 0 = permissive.

Sector profile	Strictness	Block (%)	Controls (%)	Allow (%)
Defense & critical infrastructure	2	81.6	18.2	0.2
K-12 schools	2	71.5	28.3	0.2
Government & public sector	2	49.3	50.1	0.7
Banks & financial services	2	46.5	52.2	1.2
Insurers & brokers	2	45.6	53.2	1.2
Healthcare	2	45.5	53.4	1.2
Pharma & life sciences	2	44.4	54.8	0.8
Law firms	2	43.0	52.9	4.1
Accounting & tax firms	2	39.2	55.9	5.0
MSP default (managed clients)	1	34.3	58.1	7.6
Manufacturing & engineering	1	31.3	42.0	26.7
General business	1	14.7	59.0	26.3
Higher education	1	13.1	41.9	45.0
Software & technology companies	0	12.9	40.3	46.9
Marketing, media & creative agencies	0	10.0	35.5	54.5

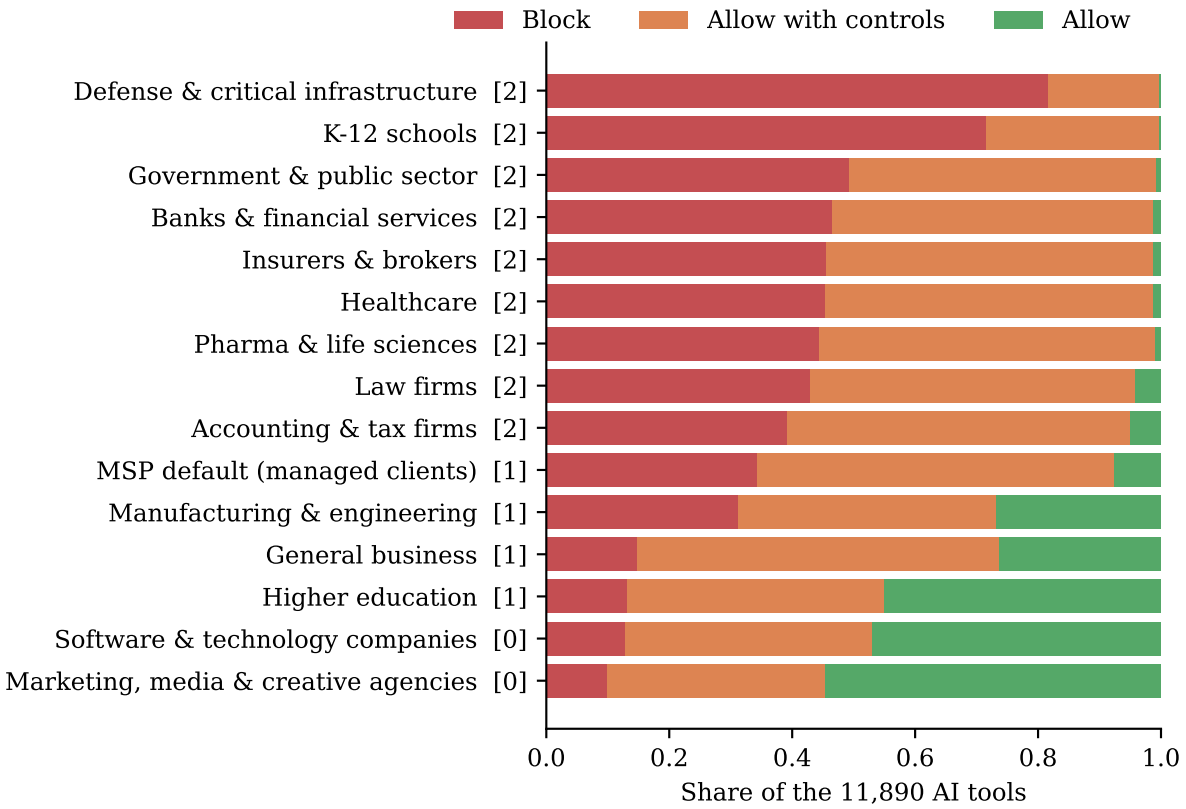


Figure 1: Verdict shares per sector profile. Strictness level in brackets.

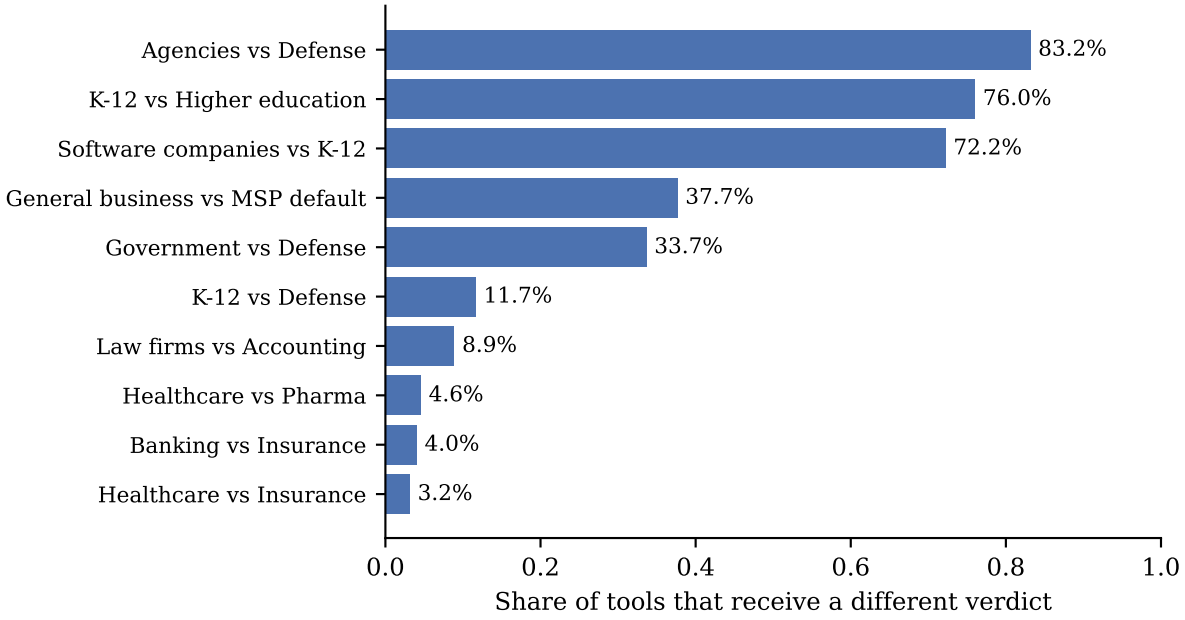


Figure 2: Share of tools whose verdict differs between two sector profiles, for selected pairs.

A school blocks an image generator to protect minors; a defense contractor blocks it to keep data within the perimeter. Similar verdicts do not imply similar policies, which is one reason to publish rationales alongside rules.

Specific contrasts. Figure 3 follows four categories across sectors.

- *Productivity & Collab*, which contains meeting assistants and office copilots, is blocked for 64.9% of tools in healthcare, 73.3% in banking and 79.5% in defense, but for 7.9% in general business. Meeting assistants record conversations that in these sectors contain patient, client or classified information.
- *Code & Development* is allowed for 81.8% of tools in higher education and blocked for 67.0% in defense. Software companies, perhaps surprisingly, allow only 4.7%: code assistants see proprietary source, so they are placed under controls.
- *Marketing, Sales & SEO* is allowed for 87.1% of tools in software companies and blocked for 73.6% in K-12 and 58.5% in government.
- *Education & Learning* is allowed for 79.4% of tools in higher education and for none in K-12, where the default is controls and 16.1% end up blocked.

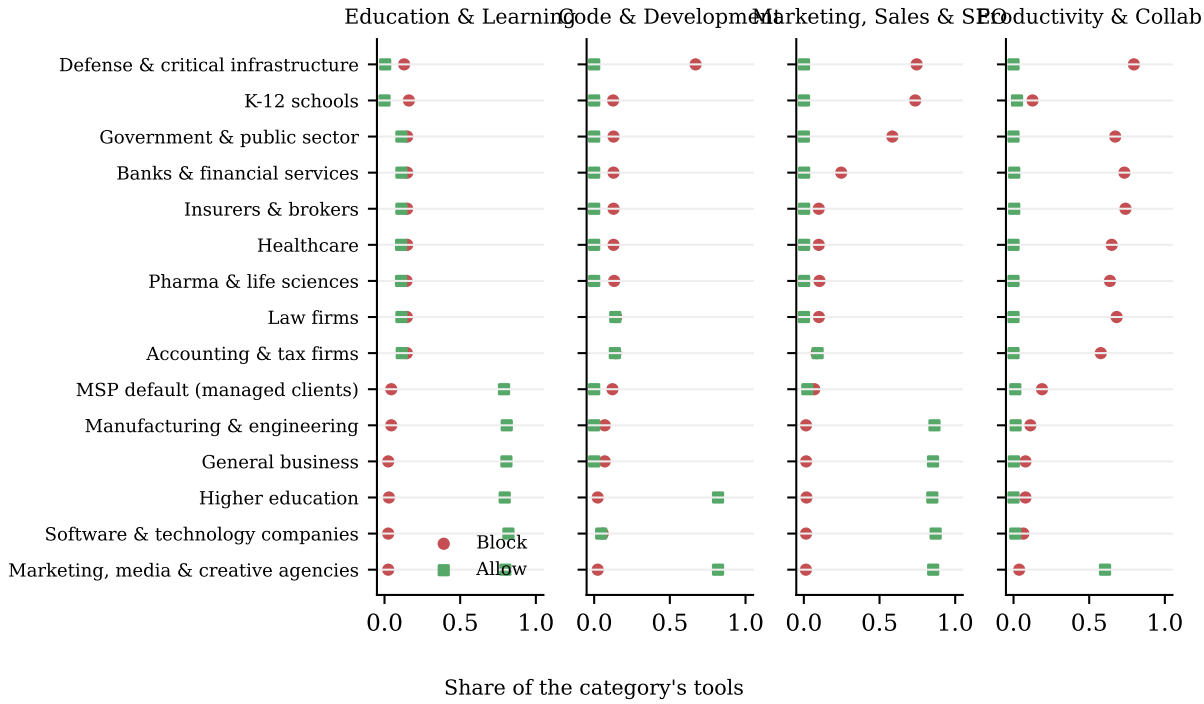


Figure 3: Block and allow shares for four categories across the fifteen sector profiles, ordered by overall strictness.

5 The K-12 policy in full

K-12 schools are a useful case for full documentation. The users are minors, the legal duties are specific (United States Congress, 2000, 1998, 1974), and the educational stakes of generative AI are well studied (Kasneci et al., 2023; Cotton et al., 2024). The profile states three goals in order: protect minors, protect academic integrity, and keep the vetted classroom tools.

5.1 What "controls" means in a school

In the K-12 profile, *allow with controls* means that a tool is reachable only through a district-managed account, that the vendor has signed a student data-privacy agreement, that users under 13 are covered by school-authorised consent (United States Congress, 1998), and that a teacher has approved the tool for a class. These conditions follow the recommendations of U.S. Department of Education, Office of Educational Technology (2023) and Miao and Holmes (2023).

5.2 Category rules and their rationale

Table 2 is the complete category matrix. Six categories default to controls and twelve to block. The profile also blocks tools the register cannot categorise, although after the final review no active tool is uncategorised. The rationale for each follows.

Image & Visual (block). Open image generation is the main route to explicit, violent or harassing imagery of real people, including classmates. Output filters are uneven across vendors. The whole category is blocked, including benign editing tools, because they share accounts and models with generators.

Text & Language (block). General text generators write the essay for the student. Open chat also exposes minors to unfiltered output and asks for personal data at sign-up. Blocked by

Table 2: The K-12 policy matrix. *Default*: baseline verdict for the category. The three right-hand columns give final verdicts for tools with that primary category, after subcategory rules and tool attributes.

Category	Default	Tools	Block	Controls	Allow
Image & Visual	Block	1,348	1,279	67	2
Text & Language	Block	1,135	958	177	0
Marketing, Sales & SEO	Block	1,045	769	275	1
Agents & Automation	Block	1,002	847	155	0
Video	Block	882	720	162	0
Business & Vertical	Block	842	722	120	0
Lifestyle & Entertainment	Block	694	660	34	0
Code & Development	Controls	688	86	602	0
Data, Analytics & Research	Block	635	464	171	0
Models & Infrastructure	Block	630	552	78	0
Audio, Voice & Music	Block	593	288	305	0
Customer Support	Block	592	559	33	0
Search, Knowledge & Docs	Block	477	390	87	0
Education & Learning	Controls	447	72	375	0
Productivity & Collab	Controls	430	54	366	10
Security & Detection	Controls	195	28	153	14
Design, 3D & Creative	Controls	193	43	150	0
Robotics & Embodied AI	Controls	62	13	49	0

default; only translation and proofreading are carved out, because they support learners rather than replace their work.

Marketing, Sales & SEO (block). No classroom use case. These tools collect contact data and send messages at scale, which has no place on a student network. Blocked.

Agents & Automation (block). Autonomous agents act on accounts, browse and send data without a human checking each step. In a school that means student records and email leaving the network unsupervised. Blocked; automation for staff can be approved individually.

Video (block). Text-to-video, avatars and dubbing produce synthetic media of real people, the same risk as image generation but more convincing. Blocked. Captioning and screen recording are carved out as accessibility and teaching tools.

Business & Vertical (block). Sector software for finance, HR, legal or healthcare professionals. Nothing here is built for pupils and many tools process sensitive personal data. Blocked on student networks; the district's own back office is out of scope for this profile.

Lifestyle & Entertainment (block). Companion chatbots, dating, astrology and games are the highest-risk area for minors: emotional dependence, sexual content and data collection. Wellness chatbots are not a substitute for a school counsellor. Blocked in full.

Code & Development (allow with controls). Computer science is taught in K-12, and code assistants are part of how software is now written. Allowed with controls: teacher-approved tools, managed accounts, and no personal data in prompts.

Data, Analytics & Research (block). Business analytics and scraping tools have no student use and often ingest uploaded files that may hold student records. Blocked. Research assistants are also blocked at K-12 level because they draft literature reviews wholesale.

Models & Infrastructure (block). Raw model APIs and hosting platforms bypass the filters that consumer apps apply. A student with an API key has an unfiltered model. Blocked; local runtimes are carved out for supervised computer-science teaching.

Audio, Voice & Music (block). Voice cloning and voice changers enable impersonation and bullying. Music and podcast generators have little classroom value. Blocked by default; speech-to-text and text-to-speech are carved out because they are assistive technology.

Customer Support (block). Helpdesk bots and call-centre agents are business software. No student use; blocked.

Search, Knowledge & Docs (block). AI answer engines return a finished answer instead of sources, and chat-with-PDF tools summarise set reading. Both undercut the skills the school is teaching. Blocked; ordinary filtered web search remains available.

Education & Learning (allow with controls). The category built for schools: tutors, quizzes, course creation, grading. Allowed with controls: the vendor must sign a student data-privacy agreement, and tools with under-13 users need school-authorised consent. Plagiarism detection is allowed for staff.

Productivity & Collab (allow with controls). These are the tools a school already runs for staff and older students: documents, slides, task boards. They are allowed with controls, meaning a district-managed account covered by a data-privacy agreement. Mind-mapping is allowed outright because it carries no free-text generation risk.

Security & Detection (allow with controls). Security and fraud tools are staff tools. Allowed with controls. Detection tools for AI content, deepfakes and moderation are allowed outright because they support safeguarding.

Design, 3D & Creative (allow with controls). Design, CAD, 3D and game development are taught subjects. Allowed with controls, so a teacher decides which tools a class uses. These tools rarely generate harmful content of real people.

Robotics & Embodied AI (allow with controls). Robotics kits and computer vision appear in STEM classes. Allowed with controls, as the risk is low and the teaching value is real.

Uncategorized (block). Tools the register could not place in a category. A school cannot vet what it cannot describe, so unknown AI is blocked until reviewed.

5.3 Subcategory rules

Of the 165 subcategories, 112 are blocked, 48 are allowed with controls and 5 are allowed. Twelve rules depart from their category default. Seven relax a blocked category for accessibility and teaching: translation, proofreading, speech-to-text, text-to-speech, captions, screen recording and local model runtimes move to controls. Five tools for safeguarding are allowed outright: mind-mapping, plagiarism detection, AI-content detectors, content moderation and deepfake forensics. Allowing detectors does not make their results reliable (Weber-Wulff et al., 2023; Liang et al., 2023); they are allowed because they do not expose pupils to generation. Appendix A lists all 165 rules with a rationale.

A rule can also tighten a permissive category. Because the strictest matching subcategory rule wins, an education tool that is also labelled as a writing or chatbot tool is blocked. As a result, 16.1% of tools whose primary category is Education & Learning are blocked in K-12, even though the category default is controls (Figure 4).

5.4 Which layer decides

Figure 5 shows where K-12 verdicts come from. The category default decides 59.1% of tools and a K-12 subcategory rule 29.0%; cross-sector subcategory rules decide 6.3% and hard blocks 5.5%. Tool attributes change 1,419 baseline verdicts. The largest single effect is AI-enabled relief, which

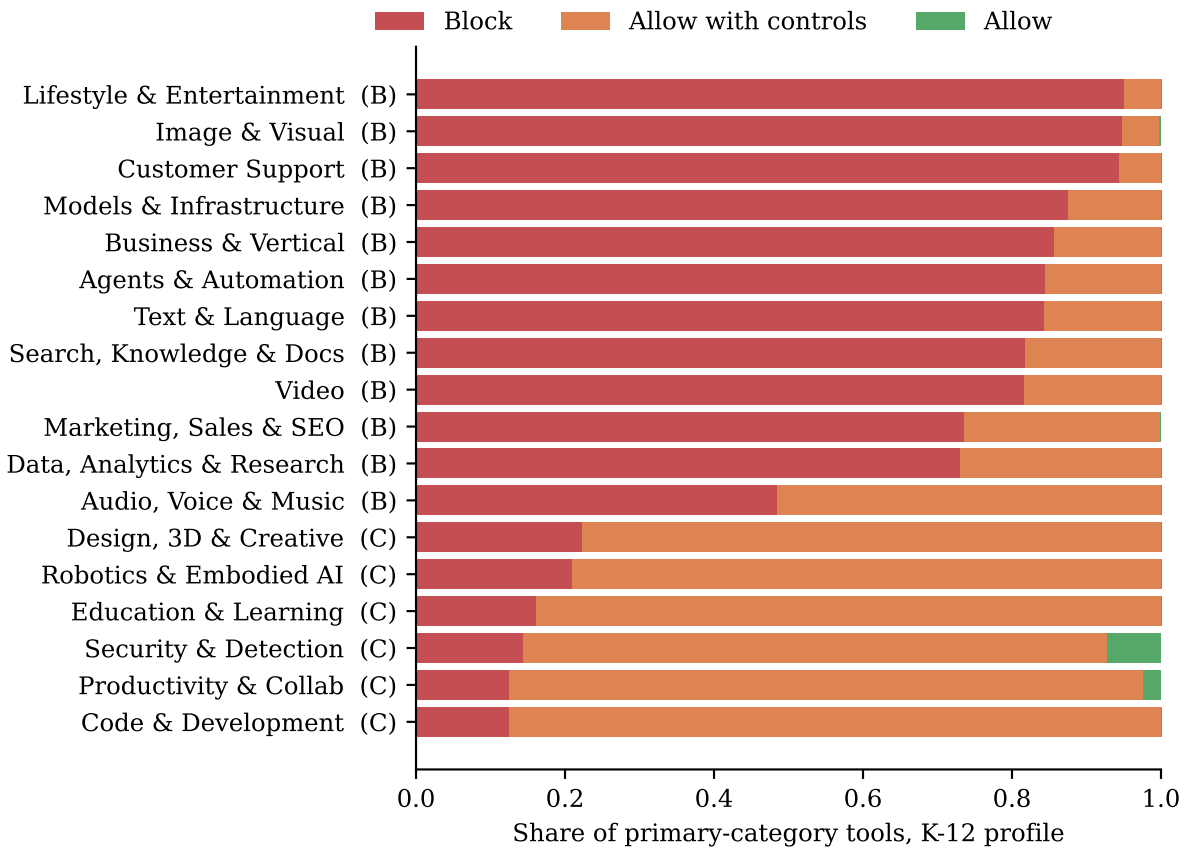


Figure 4: Final K-12 verdicts by primary category. Category default in brackets: B block, C controls, A allow.

turns 1,023 blocks into controls: in a school, an established learning platform that adds an AI feature is managed, not removed. Training on inputs by default changes 150 verdicts, 109 of them to block. A missing training policy moves 82 otherwise allowed tools out of the allow verdict, and enterprise tiers without training relax 53 blocks.

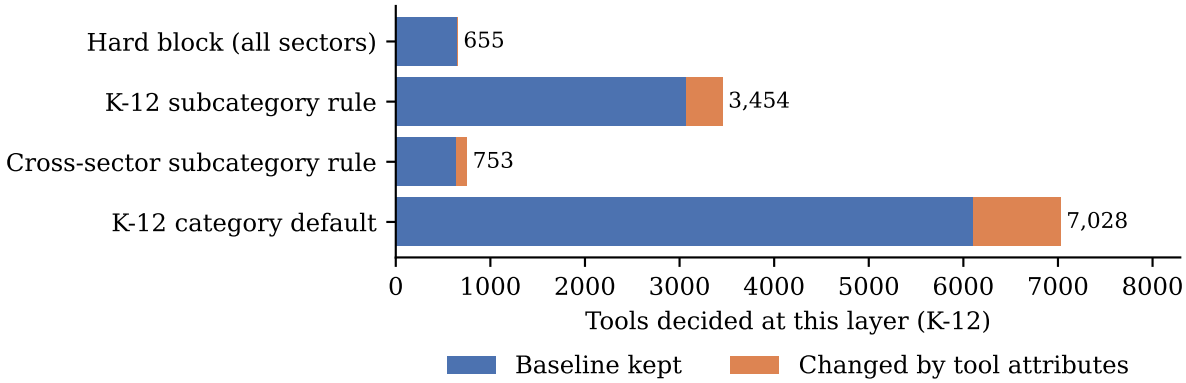


Figure 5: Number of K-12 verdicts decided at each layer, and how many of these were changed by tool attributes.

With these rules, K-12 blocks 8,504 tools (71.5%), places 3,359 (28.3%) under controls and allows 27 (0.2%). The allowed set is deliberately small. It is not a list of recommended products but the residue of tools that pose neither a content nor a data risk.

6 Discussion

Policy as rules, not prose. Sector guidance tells organisations what to protect. Turning it into rules exposes choices that prose leaves implicit: whether a school's education tools that also write essays are education tools or writing tools, or whether a bank's meeting assistant is productivity software or a recording device. Written as rules, these choices can be reviewed, audited and changed.

Subcategories carry the policy. In K-12, 35.4% of verdicts come from subcategory rules rather than category defaults. A policy expressed only at category level would either block translation tools with chatbots or allow chatbots with translation tools. The subcategory level is where sector differences become concrete.

Clusters suggest shared baselines. Healthcare, banking, insurance and pharma disagree on 3 to 5% of tools in the pairs we measured. A shared baseline for data-regulated sectors, with sector-specific exceptions, would remove most duplicated effort. Education is different: schools and universities need separate baselines, because the age of the user changes the answer for more than three quarters of tools.

Attributes matter at the margin. Most verdicts are set by what a tool does. Stated training policies and AI-enabled status change about one verdict in ten. Because most tools state no training policy, strict profiles treat silence as a reason for controls rather than as a neutral default.

7 Limitations

The profiles are normative: they encode one reading of sector guidance and are not an empirical survey of organisations' policies. Different readings would produce different numbers, which is the reason for publishing the rules. Tool attributes were assigned with automated, LLM-assisted methods, so individual tools can be misclassified. The register is a snapshot, and the verdict counts will change as tools appear and disappear.

8 Conclusion

Applied to the same 11,890 AI tools, fifteen sector profiles block between 10.0% and 81.6% of them, and two sectors disagree on 41.3% of tools on average. The differences are systematic: data-regulated sectors cluster, while schools and universities diverge sharply. Documenting one sector completely shows that a policy can be written as explicit rules with a stated reason for every category and subcategory. We encourage sector bodies to publish their guidance in that form.

Data availability

Aggregate sector statistics (verdict counts per sector and category) are documented at `https://www.shadowaitools.com/shadow-ai-sector-policy-dataset.php`. The underlying taxonomy statistics (version 1.3.0, the snapshot used here) are available under CC BY 4.0 on IEEE DataPort, doi: `10.21227/hrxe-fa21`, and documented at `https://www.aitoolsblocklist.com/ai-tools-taxonomy-dataset.php`. The full K-12 rule set is given in Appendix A.

Declaration on the use of AI

AI-assisted tools were used to review the classification of the tools in the register, to help draft and edit the text and to prepare figures. The author reviewed all content and takes full responsibility for it.

Competing interests

The author is affiliated with Alpha Quantum, which operates AI Tools Blocklist (`https://www.aitoolsblocklist.com`) and Shadow AI Tools (`https://www.shadowaitools.com`), the sources of the register and the sector profiles analysed here.

References

American Bar Association Standing Committee on Ethics and Professional Responsibility. Formal opinion 512: Generative artificial intelligence tools, July 2024.

Debby R. E. Cotton, Peter A. Cotton, and J. Reuben Shipway. Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2):228–239, 2024. doi: 10.1080/14703297.2023.2190148.

Cyberhaven. 11% of data employees paste into ChatGPT is confidential. Cyberhaven research blog, 2023.

European Parliament and Council of the European Union. Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). Official Journal of the European Union, L 333, 2022.

European Parliament and Council of the European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L series, 2024.

Financial Stability Board. The financial stability implications of artificial intelligence. Technical report, FSB, Basel, 2024.

FINRA. Regulatory notice 24-09: FINRA reminds members of regulatory obligations when using generative artificial intelligence and large language models, June 2024.

Mark Gurman. Samsung bans staff's AI use after spotting ChatGPT data leak. Bloomberg News, 2 May 2023, May 2023.

Information Commissioner's Office. Age appropriate design: A code of practice for online services. ICO, United Kingdom, 2020.

Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günnemann, Eyke Hüllermeier, et al. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103:102274, 2023. doi: 10.1016/j.lindif.2023.102274.

Weixin Liang, Mert Yuksekgonul, Yining Mao, Eric Wu, and James Zou. GPT detectors are biased against non-native English writers. *Patterns*, 4(7):100779, 2023. doi: 10.1016/j.patter.2023.100779.

Fengchun Miao and Wayne Holmes. *Guidance for Generative AI in Education and Research*. UNESCO, Paris, 2023. doi: 10.54675/EWZM9535.

Microsoft and LinkedIn. AI at work is here. now comes the hard part: 2024 work trend index annual report. Microsoft WorkLab, 2024.

National Institute of Standards and Technology. Artificial intelligence risk management framework (AI RMF 1.0). Technical Report NIST AI 100-1, NIST, 2023.

Office of Management and Budget. Memorandum m-24-10: Advancing governance, innovation, and risk management for agency use of artificial intelligence. Executive Office of the President, March 2024.

Samo Plibersek. AI tools taxonomy and training-policy statistics: 18 categories, 165 subcategories, 11,890 AI tools (version 1.3.0). IEEE DataPort, 2026a.

Samo Plibersek. A functional taxonomy for AI governance: 18 categories, 165 subcategories and the training policies of 11,890 AI tools. SSRN working paper, 2026b.

Mario Silic and Andrea Back. Shadow IT – a view from behind the curtain. *Computers & Security*, 45:274–283, 2014. doi: 10.1016/j.cose.2014.06.007.

United States Congress. Family educational rights and privacy act. 20 U.S.C. §1232g; 34 C.F.R. Part 99, 1974.

United States Congress. Children's online privacy protection act of 1998. 15 U.S.C. §§6501–6506; implemented by the FTC COPPA Rule, 16 C.F.R. Part 312, 1998.

United States Congress. Gramm-leach-bliley act, title v. 15 U.S.C. §§6801–6809, 1999.

United States Congress. Children's internet protection act. Pub. L. 106-554, codified at 47 U.S.C. §254(h) and 20 U.S.C. §9134, 2000.

U.S. Department of Education, Office of Educational Technology. Artificial intelligence and the future of teaching and learning: Insights and recommendations. Technical report, U.S. Department of Education, Washington, DC, 2023.

U.S. Department of Health and Human Services. HIPAA privacy and security rules. 45 C.F.R. Parts 160 and 164, 2013.

U.S. Department of State. International traffic in arms regulations. 22 C.F.R. Parts 120–130, 2024.

Debora Weber-Wulff, Alla Anohina-Naumeca, Sonja Bjelobaba, Tomáš Foltýnek, Jean Guerrero-Dib, Olumide Popoola, Petr Šigut, and Lorna Waddington. Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19:26, 2023. doi: 10.1007/s40979-023-00146-z.

A K-12 rules for all 165 subcategories

Rules are listed by category, in descending order of tool count. The *Blocked* column reports the final K-12 outcome for tools carrying the label, after all layers. It can exceed zero for a subcategory under controls because the strictest label of a tool wins, and it can fall below 100 for a blocked subcategory because AI-enabled relief applies.

Table 3: K-12 baseline for all 165 subcategories. *Rule*: baseline position (block, ctrl = allow with controls, allow); * marks a rule that departs from its category default. *Tools*: tools labelled with the subcategory. *Blocked*: share of those tools whose final K-12 verdict is block.

Subcategory	Rule	Tools	Blocked (%)	Rationale
Text & Language default: block				
Writing & long-form	block	488	83	Produces complete essays; the core academic-integrity risk.
General assistants & chatbots	block	451	92	Open chat writes any assignment and has no age gate; a school-provisioned assistant can be approved separately.
Copywriting & marketing copy	block	360	71	Commercial copy tools; no classroom purpose.
Summarization	block	220	57	Summarises set reading so it need not be read.
Storytelling & fiction	block	180	98	Writes creative assignments for the student and can produce unsuitable content.
Translation & localization*	ctrl	176	23	Supports pupils learning in a second language; controls keep personal data out of prompts.
Prompt engineering & marketplaces	block	167	84	Prompt marketplaces include jailbreaks and adult prompts.
Resume & bio writers	block	118	86	Adult job-seeking tools; careers staff can request an exception.
AI-text detection & humanizers	block	117	91	Humanisers exist to disguise AI-written work; blocked in every sector.
Paraphrasing	block	109	80	Rewrites copied text to evade plagiarism checks.
Grammar, style & proofreading*	ctrl	103	50	Feedback on the student's own text; widely accepted in schools under a managed account.
Email writing	block	101	64	Business correspondence; no student need.
Productivity & Collab default: controls				
Meeting assistants	ctrl	197	27	Staff meetings only; recordings of minors need consent, so managed accounts.
Task & project management AI	ctrl	165	26	Useful for project work; managed accounts only.
Document drafting & office copilots	ctrl	161	40	Built into the office suites schools use; controlled through the district tenant.
Presentation & slide generation	ctrl	112	21	Classroom presentations; controls keep the student doing the content.
Email management & triage	ctrl	86	44	Staff inbox tools; must stay inside the district tenant.
Calendar & scheduling	ctrl	83	48	Staff scheduling; low risk with managed accounts.
Mind-mapping & diagramming*	allow	66	18	Structures the student's own ideas; low data and content risk.
Spreadsheet AI	ctrl	59	29	Data lessons and administration; managed accounts.
Image & Visual default: block				
Text-to-image	block	871	97	Main source of explicit or harassing imagery of real people.
Image editing & inpainting	block	496	93	Inpainting can alter photos of classmates.
Illustration & character	block	197	98	Character generators frequently host adult styles.
Graphic design & layout	block	180	84	Shares models and accounts with generators; design classes use approved design tools instead.
Avatars & headshots	block	170	98	Requires uploading a face photo of a minor.
Face swap & effects	block	162	99	Direct deepfake risk; blocked in every sector.

Subcategory	Rule	Tools	Blocked (%)	Rationale
Product photography & mockups	block	142	85	E-commerce tooling; no class use.
Background removal	block	128	91	Low risk alone, but blocked with the category; can be allowed per tool.
Upscaling & restoration	block	124	93	Enhances footage of real people; little class value.
Interior & architecture viz	block	111	81	Professional visualisation; design courses use CAD tools under controls.
OCR & image-to-text	block	93	67	Photographs worksheets to get answers from linked solvers.
Logo & brand assets	block	72	90	Business branding; no class use.
Pattern & texture	block	16	81	Minimal classroom need; blocked with the category.
Marketing, Sales & SEO	default: block			
SEO & GEO optimization	block	403	70	Commercial tool; no student use.
Sales prospecting & SDR agents	block	321	86	Harvests contact data and sends outreach at scale.
Social media management	block	305	70	Posting on behalf of accounts; minors' social media is outside school use.
Ad creative & optimization	block	275	77	Advertising tools; no student use.
CRM intelligence & lead scoring	block	180	71	Processes personal data of prospects; no student use.
Product marketing	block	139	79	Commercial tool; no student use.
Brand monitoring & sentiment	block	117	74	Commercial monitoring; no student use.
Email marketing & personalization	block	104	68	Bulk messaging; no student use.
CRO & A/B testing	block	74	64	Website optimisation; no student use.
Influencer & UGC	block	62	87	Creator marketing, often involving minors' images; blocked.
Agents & Automation	default: block			
AI workflow automation	block	616	76	Moves data between accounts without supervision.
MCP servers & integrations	block	387	68	Connects models to local files and accounts; data-exposure risk.
Agent-building frameworks	block	369	88	Lets users build unfiltered agents.
General autonomous agents	block	349	87	Acts on the web and accounts with no teacher in the loop.
Personal AI assistants	block	267	79	Reads mail and calendars; student accounts must not be connected.
Multi-agent orchestration	block	158	77	Developer infrastructure for autonomous agents.
Browser agents & web automation	block	109	82	Can complete online tests and forms for the student.
RPA with AI	block	43	74	Back-office automation; not for student networks.
Data, Analytics & Research	default: block			
Market & competitive intelligence	block	263	73	Business research; no student use.
BI & natural-language analytics	block	205	61	District analysts can request it; not for students.
Predictive analytics & forecasting	block	195	75	Business analytics; uploads may contain student data.
Document data extraction	block	181	85	Processes uploaded records, which in schools means student files.
Survey & feedback analysis	block	117	74	Survey data about pupils is FERPA-protected; staff exception only.
Data-viz generation	block	95	51	Generates charts from uploads; statistics classes can request a tool.
AI web scraping	block	95	85	Harvests personal data; blocked in every sector.
Academic research assistants	block	91	76	Drafts literature reviews; suited to universities, not K-12.
Deep-research agents	block	73	79	Produce finished reports; an academic-integrity risk.
Data cleaning & prep	block	61	66	Business data tooling; no student use.
Video	default: block			

Subcategory	Rule	Tools	Blocked (%)	Rationale
Text-to-video	block	773	93	Synthetic video of real people; highest-impact deepfake route.
AI video editing	block	323	75	Shares generation features with text-to-video.
Subtitles & captions*	ctrl	167	19	Accessibility tool for deaf and hard-of-hearing pupils.
Animation & motion graphics	block	144	88	Generative animation with weak content filters.
Talking avatars & presenters	block	131	92	Synthetic presenters built from real faces.
Dubbing, lip-sync & translation	block	128	67	Lip-sync makes anyone say anything.
Short-form clip repurposing	block	115	62	Social media production; no class use.
Video summarization & search	block	94	70	Summarises set videos so they need not be watched.
Screen recording*	ctrl	39	8	Teachers record lessons and walkthroughs.
Upscaling & restoration	block	15	93	Enhances footage of real people; little class value.
Business & Vertical	default: block			
Finance & accounting	block	275	77	Professional software; no student use.
HR & recruiting	block	207	84	Processes applicant data; no student use.
E-commerce	block	165	72	Commercial selling tools.
Healthcare & med-tech	block	160	85	Clinical tools; school nurses use approved systems.
Legal	block	155	84	Professional legal tools.
Real estate	block	86	87	Commercial listings tools.
Manufacturing & industrial	block	64	81	Industrial software.
Procurement & supply chain	block	48	79	Business purchasing tools.
Government & public sector	block	43	72	Administrative tools for agencies.
Insurance & risk	block	41	90	Professional underwriting tools.
Agriculture	block	13	85	Professional farm software; vocational programmes can request exceptions.
Energy & climate	block	10	90	Professional energy software.
Audio, Voice & Music	default: block			
Speech-to-text & transcription*	ctrl	336	16	Assistive technology for pupils with dyslexia or motor impairments.
Text-to-speech*	ctrl	230	35	Reads text aloud for pupils with reading difficulties.
Music generation	block	190	92	Generates complete songs; music classes can approve specific tools.
Real-time voice agents & phone AI	block	115	83	Places calls autonomously.
Voice cloning	block	109	61	Impersonation of teachers and pupils; blocked in every sector.
Podcast tools	block	82	39	Production tooling with voice synthesis.
Stem separation & remixing	block	63	97	Low classroom need.
Audio enhancement & mastering	block	46	74	Low risk but low need; blocked with the category.
Voice changers & real-time conversion	block	28	89	Disguises identity in calls and games; blocked in every sector.
Sound-effects generation	block	23	91	Low classroom need.
Code & Development	default: controls			
App builders & vibe coding	ctrl	248	16	Project-based computing lessons; teacher-approved accounts.
Autonomous coding agents	ctrl	234	31	Can complete assignments; teacher decides when.
Code assistants & autocomplete	ctrl	155	15	Standard in modern programming courses.
Testing & QA	ctrl	110	15	Teaches testing practice.
UI generation & design-to-code	ctrl	88	6	Web design lessons.
Code review	ctrl	84	17	Feedback on the student's own code.
DevOps, IaC & SRE	ctrl	66	21	Advanced courses and IT staff.
Documentation	ctrl	65	8	Supports coding projects.

Subcategory	Rule	Tools	Blocked (%)	Rationale
Debugging & error monitoring	ctrl	56	14	Helps learners understand errors.
SQL & database queries	ctrl	42	21	Database lessons.
Migration & refactoring	ctrl	32	3	IT staff and advanced courses.
API dev assistants	ctrl	13	31	Advanced computing courses.
Lifestyle & Entertainment	default: block			
AI companions & character chat	block	319	98	Emotional dependence and sexual role-play with minors; blocked in every sector.
Dating & relationships	block	107	97	Adult services; blocked in every sector.
Astrology & fun generators	block	106	98	No educational value; blocked in every sector.
Mental wellness & journaling	block	97	93	Not a substitute for the school counsellor; sensitive disclosures leave the school.
Health, fitness & nutrition	block	91	84	Diet advice to minors carries eating-disorder risk.
Gaming & AI-native games	block	71	99	Distraction and open chat; blocked on school networks.
Fashion & styling	block	58	95	Uploads body photos; no class use.
Travel planning	block	58	84	Consumer tool; no class use.
Gift & shopping recommendation	block	36	92	Commercial recommendations aimed at purchases.
Personal finance	block	35	86	Connects bank accounts; not for minors.
Cooking & recipes	block	29	90	Consumer tool; food technology classes can request one.
Customer Support	default: block			
Support chatbots & AI helpdesk	block	457	93	Business software.
Voice & call-center agents	block	287	94	Business software.
Conversational commerce & messaging bots	block	219	93	Sales messaging.
Ticket routing & agent-assist	block	92	82	District IT can approve its own helpdesk.
Review management	block	30	80	Business reputation software.
Models & Infrastructure	default: block			
Observability & evals	block	227	85	Developer tooling.
Hosting & inference platforms	block	170	89	Run any open model, including uncensored ones.
Foundation models & model APIs	block	167	89	Direct model access bypasses consumer safety layers.
LLM gateways & routing	block	158	91	Provides unfiltered access to many models behind one key.
Guardrails, safety & red-teaming	block	138	74	Developer tooling; red-team kits include jailbreaks.
Fine-tuning & training	block	78	83	Can remove safety training from models.
Local runtimes*	ctrl	60	30	Supervised computer-science teaching; the model runs on school hardware.
GPU cloud & compute	block	43	84	Paid compute; no student use.
Prompt management	block	39	87	Developer tooling.
Data labeling & annotation	block	38	89	Crowdwork platforms; not for minors.
Vector DBs & embeddings	block	33	91	Developer infrastructure.
Synthetic data	block	18	100	Developer tooling.
Search, Knowledge & Docs	default: block			
Chat-with-documents & PDF	block	212	81	Answers questions on set texts.
Note-taking & second brain	block	178	67	Consumer note apps with AI writing; school suites cover the need.
RAG platforms & retrieval infra	block	176	90	Developer infrastructure.
AI search & answer engines	block	167	75	Replaces source evaluation with a finished answer.
Enterprise search & KM	block	165	72	Business knowledge tools.
Browser copilots & reading assistants	block	77	70	Sit on every page and answer online tests.

Subcategory	Rule	Tools	Blocked (%)	Rationale
Education & Learning	default: controls			
AI tutors & personalized learning	ctrl	337	19	Core classroom use; requires a data-privacy agreement and age-appropriate design.
Quiz, flashcard & study	ctrl	175	15	Revision tools; managed accounts.
Course & curriculum creation	ctrl	103	13	Teacher planning tools.
Grading & feedback	ctrl	87	13	Teacher tools handling student work, a FERPA education record.
Language learning	ctrl	72	25	Established classroom use.
Kids' learning & STEM	ctrl	44	41	Built for the age group; still needs consent for under-13s.
Plagiarism & AI-content detection*	allow	25	60	Staff safeguarding tool; results are evidence, not proof.
Design, 3D & Creative	default: controls			
3D model & asset generation	ctrl	85	38	3D printing and design lessons.
Game-dev AI	ctrl	71	37	Game design courses.
UI/UX design assistants	ctrl	66	14	Design and computing lessons.
Architecture & CAD	ctrl	60	38	Technical drawing and engineering courses.
Fashion design & virtual try-on	ctrl	59	81	Textiles courses; try-on uses body photos, so controls.
AR/VR content	ctrl	13	54	Immersive learning content.
Font & typography	ctrl	8	38	Design lessons.
Security & Detection	default: controls			
AI for cybersecurity	ctrl	105	18	IT staff tools.
AI-content detectors*	allow	40	28	Safeguarding; allowed in every sector.
Fraud detection	ctrl	34	26	IT and finance staff.
Identity verification	ctrl	32	28	Collects ID documents; district-approved only.
Privacy & PII redaction	ctrl	32	38	Helps staff strip student data before sharing.
Content moderation*	allow	24	12	Safeguarding; allowed in every sector.
Deepfake detection & forensics*	allow	24	21	Investigating image-based abuse; allowed in every sector.
Robotics & Embodied AI	default: controls			
Computer-vision & edge vision	ctrl	65	46	STEM projects; cameras in class need controls.
Robotics platforms & foundation models	ctrl	23	22	Robotics clubs and courses.
Wearables & AI hardware	ctrl	14	43	Collect biometric data; controls.
IoT & smart-home AI	ctrl	12	25	Maker projects.
Autonomous vehicles & drones	ctrl	7	14	STEM projects under supervision.