Artificial Intelligence (AI)

Cambridge School of Bucharest

DISCLAIMER

The information contained in this document has been prepared exclusively for the internal use of Cambridge School of Bucharest. Since this document is confidential and proprietary to Cambridge School of Bucharest, third parties are not entitled to rely on it in any way. The present document has been issued in English.

This document is proprietary to Cambridge School of Bucharest. All external communication, reproduction or publication must be authorised by Cambridge School of Bucharest. This is an uncontrolled copy that must be considered obsolete after the print date.

	PREPARED BY	CHECKED BY	APPROVED BY
Name	Matthew Wemyss	BATTLEGROUP	School Director
Title	Assistant School Director	Data Protection Officer	Rita Maalouf

Next Review Date:	August 2027

STAFF ESSENTIALS

- Use only an approved AI tool for an approved school purpose. Personal-account AI and all connectors to School accounts or repositories are prohibited.
- Do not enter personal, confidential, safeguarding, health, SEND, HR, admissions or financial data unless the exact workflow has written approval.
- Keep professional judgement and human decision-making. AI must not make safeguarding, admissions, assessment, recruitment, disciplinary or other significant decisions.
- Check accuracy, bias, source quality, copyright, tone, age suitability and accessibility before an output is shared, acted on or added to a record.
- Follow the Article 50 disclosure and labelling rules in Appendix B for interactive AI and realistic synthetic media.
- Report a child-safeguarding concern to the DSL immediately. Report suspected data or cyber incidents to the DPO and IT. Do not wait for the AI Lead.

1. PURPOSE, STATUS AND SCOPE

This policy sets the rules for safe, lawful and responsible use of artificial intelligence on behalf of Cambridge School of Bucharest (CSB). It is written for staff and other adults who work for, at, or on behalf of the School.

For this policy, staff includes employees, leaders, governors and Proprietor representatives, agency and supply staff, contractors, consultants, coaches, volunteers, interns, trainees and visiting professionals. It applies whether work is carried out on site or elsewhere, using a school or personal device or account.

The policy applies to generative AI, predictive or analytical AI, computer vision, biometric systems, AI-enabled recording and transcription, AI agents, and AI features built into other products. It covers use in teaching, safeguarding and pastoral work, admissions, recruitment and HR, communications and social media, finance, IT, facilities and school operations.

Pupil use is governed by the School's Acceptable Use, Academic Integrity and Assessment, Behaviour and Online Safety arrangements. This policy tells staff when they may direct, approve or supervise pupil use.

2. PRINCIPLES

- Child first. The best interests, dignity, safety and rights of pupils take priority over convenience or efficiency.
- Human accountability. The staff member remains responsible for the work, decision and communication. AI is a support tool.
- Approved and purpose-limited. A tool may be used only for the approved feature, users, data and purpose recorded in the live AI Tool and Use-Case Register.
- Minimum data. Use no personal or confidential data where a task can be completed with public, synthetic or properly anonymised information.
- Fair and accessible. Check for bias, stereotypes, exclusion, accessibility barriers and disproportionate effects on pupils, applicants, families or staff.
- Transparent and contestable. People must receive the disclosures required by law and School policy, and consequential decisions must be explainable, reviewable and open to human challenge.

3. ROLES AND RESPONSIBILITIES

Role	Core Responsibility
Proprietor / Governing Board	Approves the policy, ensures material AI risks appear on the School risk register, and receives annual assurance on use, incidents, exceptions and training.
School Director	Is accountable for implementation; appoints the AI Lead and Editorial Responsibility Owner; approves any high-risk, sensitive or materially autonomous use.
AI Lead	Maintains the live register; coordinates classification, approval, training, monitoring and annual review; advises staff on ordinary AI use.
Data Protection Officer (DPO)	Advises on lawful basis, transparency, contracts, international transfers and data subject rights; advises on and monitors DPIAs and data-protection compliance. The authorised School decision-maker approves processing and accepts residual risk.
Designated Safeguarding Lead (DSL)	Assesses child-safeguarding implications of tools and incidents. The DSL route takes priority for any concern about a child.
IT Network Manager	Assesses security, identity, permissions, integrations, logging, filtering and monitoring; disables unapproved connectors and unsafe features.
Editorial Responsibility Owner	Holds effective editorial control and ultimate responsibility for public-facing AI-assisted text that may concern matters of public interest.
Managers and budget holders	Ensure local use is approved, staff are trained, outputs are reviewed and procurement does not bypass this policy.
All staff and persons in scope	Follow the policy, use professional judgement, protect information, make required disclosures and report concerns promptly.

4. SAFE USE BY STAFF

4.1 Before Using AI

Before starting, answer all five questions. If an answer is no or unclear, stop and ask the AI Lead. A safeguarding or data incident must still be reported through the routes in section 11.

1. Is this exact tool, feature, account and use case on the live approved register?
2. Can the task be completed without personal, confidential or sensitive information?
3. Will a competent person review the output before it is shared, acted on or recorded?
4. Could the use affect a pupil's education or welfare, an applicant's admission, a candidate's recruitment, or a person's employment? If so, has the higher-risk route been completed?
5. Does the use involve direct AI interaction, synthetic media, public publication, recording, monitoring, a connector, an agent or an external action? If so, have the specific controls been approved?

4.2 Permitted Low-Risk Support

Subject to the approved register and the data rules, AI may assist with brainstorming, first drafts, templates, summarising public or approved material, translation, accessibility adaptations, lesson-resource ideas, formatting, coding support in a controlled environment, and analysis of public, synthetic or properly anonymised information.

4.3 Staff Review Standard

Before an output is shared, acted on or entered into a school record, the responsible staff member must:

- verify facts, quotations, references, calculations and any cited source against reliable material;
- check for bias, stereotypes, omissions and unfair or discriminatory effects;
- check age, curriculum, cultural, language and accessibility suitability;

- check copyright, permissions, confidentiality and data protection;
- edit for professional tone and make sure the final wording reflects the staff member's own judgement; and
- record material AI assistance where it influenced a consequential recommendation, formal record or regulated process.

Human review must be meaningful. The reviewer must understand the system's role, examine the relevant inputs and output, have enough time and authority, and be able to disregard or change the output rather than merely approve it.

AI output is never evidence merely because it appears confident, detailed or well written. Staff must not invent a source, citation, event, pupil observation, meeting contribution or professional judgement.

5. PROHIBITED AND HIGHER-RISK USE

5.1 Uses the School Does Not Permit

- Manipulative or deceptive systems that materially impair informed choice or exploit vulnerability linked to age, disability or social or economic circumstances.
- Social scoring, or using AI to infer or rank a person's character, trustworthiness, motivation, likelihood of offending, protected characteristics, mental state or emotions.
- Emotion-recognition AI in education or the workplace, except a formally approved medical or safety use that is lawful and necessary.
- Biometric categorisation that infers race, political opinions, trade-union membership, religion or belief, sex life or sexual orientation; or untargeted scraping of facial images to build a recognition database.
- Generation or manipulation of sexual or intimate material involving a child or identifiable person, including nudification, deep-nude and child sexual abuse material. CSB prohibits this from the policy effective date; the added EU AI Act prohibitions apply from 2 December 2026.
- Covert recording, monitoring or surveillance; impersonation; voice or image cloning of pupils, parents, staff or other identifiable people without the full approval and permissions required by the DPO, DSL and School Director.
- Solely automated decisions, or decisions materially dictated by AI, about safeguarding, admissions, grades, progression, access to learning, recruitment, pay, performance, discipline, promotion, contract renewal or dismissal.
- Using a personal-account AI tool or agent for School work or with School-related information, including copied, paraphrased or uploaded text, files, folders, emails, contacts, calendar entries, images or records.
- Enabling any connector, integration, plug-in, add-in, synced source, organisational grounding, browser extension, file picker, API token or similar permission that lets AI search, retrieve, sync or act on School email, address books, calendars, files, folders, cloud storage, Teams, SharePoint, Google Workspace, MIS/SIS, HR, admissions, safeguarding, finance or other School systems. This applies to personal and School-managed accounts.
- Allowing an AI agent to publish or share externally, send email or messages, purchase or book, alter or delete School data, change permissions, or operate another School system. An approved agent may draft or create only inside its designated School-managed workspace; a person must review and take the later action.
- Using or bringing into operation a general-purpose AI-enabled wearable device on School premises or during a School activity, including smart glasses, AI pins or pendants and worn devices capable of AI recording, transcription, identification, analysis or transmission.
- Bypassing filters, access controls, age restrictions, safety settings, licensing conditions or the approval process.

5.2 High-Risk Uses

The EU AI Act identifies certain education, recruitment, employment and biometric uses as potentially high-risk. Classification depends on intended purpose, material influence, profiling and the Article 6(3) filter. Human involvement alone does not remove high-risk status if staff materially rely on the output.

Area	Examples Requiring High-Risk Triage
Admissions and placement	Evaluating or ranking applications; deciding access, admission, scholarship eligibility, placement, programme or educational level.
Learning and assessment	Evaluating learning outcomes where the output steers learning; determining educational level; monitoring and detecting prohibited behaviour during tests.
Recruitment and HR	Targeted job advertising through profiling; screening, filtering, ranking or evaluating candidates; AI-supported decisions on terms, promotion or termination; task allocation based on behaviour or traits; performance or behaviour monitoring.
Biometrics	Remote biometric identification, biometric categorisation or emotion recognition where a use is not already prohibited.

No staff member may procure, trial or deploy a high-risk use without written approval from the School Director after AI Lead, DPO, DSL, IT and, where relevant, HR review. The Annex III high-risk rules apply from 2 December 2027; product-embedded Annex I rules apply from 2 August 2028. Existing GDPR, employment, equality and safeguarding duties apply now.

6. SAFEGUARDING AND CHILD-FACING AI

AI use is part of the School's whole-school approach to safeguarding and online safety. This section must be read with the current CSB Safeguarding and Child Protection Policy and KCSIE 2026, particularly the four areas of online risk: content, contact, conduct and commerce.

- Content: AI can produce harmful, sexual, violent, hateful, extremist, false or developmentally inappropriate material.
- Contact: generative AI can simulate human relationships, groom, manipulate, elicit disclosures or encourage secrecy and dependence.
- Conduct: AI can be used for bullying, impersonation, harassment, sexual extortion and making or sharing explicit images, including deepfakes and deep nudes.
- Commerce: AI can amplify scams, phishing, deceptive advertising and financial harm.

6.1 Staff Response to a Concern

A safeguarding concern linked to AI must be reported immediately to the DSL or deputy under the Safeguarding and Child Protection Policy. Staff must not ask an AI system to assess credibility, risk, threshold or next steps, and must not delay a report while consulting the AI Lead.

- Do not investigate the child, alleged perpetrator, account or device beyond the actions required by the safeguarding procedure.
- Do not copy, download, print or forward sexual or intimate imagery. Preserve the device, account and context and follow the DSL's instructions.
- Do not promise confidentiality. Reassure the pupil, record the facts in the approved safeguarding system and follow the School's reporting route.
- Treat any alert produced by a child-facing AI system as information for human safeguarding review, never as a conclusion.

6.2 Approval Standard for Pupil-Facing AI

A pupil-facing system must have a clear educational purpose, suitable age terms, effective content filtering, staff supervision, proportionate logging and monitoring, an accessible disclosure that it is AI,

and an immediate route to a human. It must not imply feelings, consciousness or a personal relationship, encourage secrecy, or replace pastoral, mental-health or safeguarding support.

The SLT lead for filtering and monitoring, supported by the DSL and IT, will review the effectiveness of filtering and monitoring at least once each academic year and after a serious incident or material system change. The review will cover relevant devices, locations, accounts and AI features.

7. DATA PROTECTION, CONFIDENTIALITY AND SECURITY

Tool approval does not itself authorise the use of personal data. Approval is attached to a defined workflow, data category, account, settings and contractual arrangement.

Data Category	Examples	Staff Rule
Public or synthetic	Published curriculum material, public website content, invented examples with no link to a real person.	May be used in an approved tool for an approved purpose.
Internal, non-personal	Internal templates, procedures and operational information that does not identify or reveal confidential information about a person.	Use an approved school account. Share only what is necessary.
Personal or confidential	Names, contact details, pupil work, reports, applications, CVs, staff records, unpublished school information.	Do not input unless the authorised School decision-maker has approved the exact workflow after documented DPO advice and the use is on the register.
Special, safeguarding or security	Child-protection information, allegations, health, SEND, biometrics, intimate images, passwords, access tokens, payment data.	Do not input into generative AI. A narrowly defined exception requires written School Director approval after DPO advice and relevant DSL or IT review.

- Removing a name does not necessarily anonymise data. Pseudonymised or indirectly identifiable data remains personal data; seek DPO advice.
- Use school-managed accounts, minimum access permissions and approved retention settings. Do not use school information to train or improve a vendor's general model unless expressly approved.
- Never enter passwords, API keys, authentication codes or security configurations into a generative AI tool.
- Do not upload a full document where a short, non-identifying extract is enough.
- Complete a DPIA before use where the processing is likely to result in high risk.
- AI-enabled transcription or recording requires an approved tool. Consent is not the default for pupils or staff and may be used only where it is freely given and refusal or withdrawal causes no detriment. Covert use is prohibited.
- General-purpose AI-enabled wearable devices are banned even if the user says their recording or AI function is disabled. Prescribed medical devices and assistive technology authorised through the School's reasonable-adjustment or accessibility process are not treated as general-purpose AI wearables.

8. TRANSPARENCY AND ARTICLE 50

CSB normally acts as a deployer of vendor systems. CSB will only deploy systems that meet applicable provider disclosure duties, and staff must preserve embedded provenance marks. Article 50(3) and (4) place duties directly on deployers in the situations below.

Trigger	Required School Action
A person interacts directly with AI	The provider must make the AI nature clear unless obvious. CSB must confirm the disclosure is present and add an in-context notice where needed. Provide it at the start of each session or first interaction.

Trigger	Required School Action
Realistic AI-generated or materially manipulated image, audio or video	If it could falsely appear authentic or truthful, label it clearly as AI generated or AI modified at first exposure. This can include an invented but plausible person, place, object or event.
Emotion recognition or biometric categorisation	Inform exposed people about the operation of the system and comply with data law. These uses are generally prohibited or not approved at CSB.
AI-generated or manipulated public-interest text	Either label it, or complete substantive human review or editorial control and ensure a named person or legal entity holds editorial responsibility. Review must cover substance and facts, not grammar alone.

All CSB public-facing AI-assisted text must receive substantive human review before publication. The School Director or designated Editorial Responsibility Owner has effective authority to approve, amend or reject the text and holds ultimate responsibility.

Routine emails, internal communications, teaching resources and ordinary promotional text do not require a label. Realistic synthetic media is assessed separately and may still require a label. Appendix B gives the operational test, wording and placement.

Disclosures must be clear, distinguishable and accessible at first interaction or exposure.

9. RULES FOR SCHOOL WORK AREAS

9.1 Teaching, Learning, Assessment and Reports

- Staff may use approved AI to develop ideas and drafts, but must verify curriculum accuracy, age suitability, representation and accessibility before pupil use.
- Pupils may use AI only where the teacher, approved tool terms, School policies and awarding-body rules permit it. The task must state what use is permitted and how it must be acknowledged.
- AI may support low-stakes formative feedback in an approved workflow after meaningful teacher review. A use that evaluates learning outcomes and materially steers feedback, a learning pathway, access, level, progression or assessment requires high-risk triage. AI must not set final marks, grades, progression, access arrangements or educational pathways.
- AI proctoring, or AI used to monitor or detect prohibited behaviour during a test, requires the high-risk approval route. An AI-detection score is not proof of misconduct and must not be the sole or decisive evidence. Concerns are investigated fairly under the Academic Integrity and Assessment Policy.
- Pupil reports and records must reflect verified evidence and the responsible professional's judgement. AI must not fabricate observations or infer needs, diagnoses, intent or risk.

9.2 Admissions

Approved AI may provide general admissions information, translate or organise documents, and draft communications that a staff member verifies. It must not evaluate, rank or score applicants; determine eligibility or scholarship access; recommend admission, rejection, placement or educational level; or infer family circumstances or protected characteristics.

9.3 Recruitment, HR and Staff Management

Approved AI may help draft a job description, neutral advertisement, interview-question bank or administrative correspondence from criteria set by staff. It must not target job adverts using profiling; screen, filter, rank or score applications; analyse video or voice for personality or emotion; conduct automated background checks; or materially influence recruitment, pay, deployment, performance, promotion, discipline, contract renewal or termination.

9.4 Communications, Marketing and Social Media

AI may assist with drafts, translation, design concepts and low-risk content. A competent person must check every publication for facts, permissions, tone, accessibility and misleading impressions. Public-facing text follows the editorial process in section 8. Realistic synthetic media follows Appendix B. Native platform disclosure tools should be used where available, but they do not replace a clear on-content or audible label when one is required.

Staff must not auto-publish, impersonate a person, create a synthetic pupil or staff testimonial, clone a voice, or present a generated school event, building, classroom, result or product as real. Social-media and advertising labels and provenance metadata must be preserved through resizing, export and reposting.

SLT, Safeguarding Team and DPO must consider the risk that identifiable pupil images published online may be harvested for facial datasets, deepfakes or abusive content when setting image-consent, selection, access and publication arrangements.

9.5 Administration, Finance, IT and Operations

AI may assist with templates, public-information research, coding suggestions and operational analysis in approved environments. Staff must not permit autonomous payments, purchases, bookings, access changes, file deletion, security actions or external communication. Code, formulas and automations must be tested, peer-reviewed where material, and security-checked before deployment. Finance, payroll, credentials, contracts, CCTV and access-control data require the relevant specialist approval.

Any unusual or urgent request involving payment, credentials, personal data, access or confidential information must be verified with the purported sender through a known, separate channel. A convincing voice, image, video, message or caller identity is not sufficient verification.

9.6 Agents, Connectors and AI-Generated Code

- Personal-account agents must not be used for any School purpose or with School information. This includes personal ChatGPT, Copilot, Gemini or similar accounts, even where the staff member intends only to summarise, translate or improve a School file.
- A School-managed ChatGPT Business or Enterprise workspace agent may be used only by named staff after the exact agent and purpose are approved. Approval does not extend to other agents, tools, actions, memories, models or files.
- Connectors are not permitted. Agents must not search or act on School email, files, folders, calendars, contacts, Teams, SharePoint, cloud drives or other School systems. Administrators must keep connector access disabled where the platform permits.
- Staff may use native Canva AI, Canva Code and native Canva assistant or agent features in the School-managed Canva account to create draft resources for pupils.
- Any other tool that generates code, interactive resources, automations or websites may be used only by named, approved staff for the approved purpose.

10. TOOL APPROVAL, PROCUREMENT AND MONITORING

The live AI Tool and Use-Case Register is the authoritative record. Approval belongs to a use case, feature, version, account, settings, user group and data category, not to a brand name. A free trial, pilot, browser extension or newly enabled AI feature counts as use and may require approval.

1. Request. The sponsor records the intended purpose, users, data, outputs, decisions, integrations, external actions and expected benefit. For an agent, this is the Agent Card.
2. Triage. The AI Lead checks whether the system is AI, prohibited, subject to Article 50, potentially high-risk, or within an approved low-risk pattern.
3. Specialist review. IT reviews security and integrations; the DPO reviews roles, lawful basis, data flows, contracts, subprocessors, transfers, retention, model training, change notices, export and deletion; the DSL reviews child-facing or safeguarding implications; HR, Admissions,

Communications or Finance review relevant uses. The sponsor records a check against the DfE Generative AI Product Safety Standards as adopted School good practice.

4. Decision. The AI Lead may approve low-risk use. The School Director must approve high-risk, pupil-facing, biometric, monitoring, agentic, sensitive-data or materially autonomous use.
5. Implementation. The owner documents instructions, human oversight, disclosures, access, training, incident routes, logging, retention and review dates before go-live.
6. Monitoring and exit. Material vendor, model, feature, data-flow, terms, age-limit or intended-purpose changes trigger reassessment. Unsafe or non-compliant use is suspended. When approval ends, the owner exports required School records, removes access and obtains deletion confirmation where contractually available.

Before any high-risk deployment, CSB will require provider classification and compliance evidence; an AI impact assessment and, where GDPR requires, a DPIA; a legal determination on Articles 27, 49 and 86; and a documented Article 26 plan covering instructions, competent human oversight, controllable input data, monitoring, suspension and incident reporting, notices, logs and challenge. EU database checks will be made where applicable; private-school deployer registration is not assumed.

11. INCIDENTS, CONCERNS, COMPLAINTS AND RECORDS

Concern	First Route
Child safeguarding	Stop the activity if safe; report immediately to the DSL or deputy under the Safeguarding and Child Protection Policy.
Personal data or confidentiality	Report immediately to the DPO and follow the data-breach procedure.
Cyber security, unsafe connector or unauthorised action	Report immediately to IT and the DPO where data may be affected.
Harmful, biased, false or missing disclosure	Stop sharing or using the output; report to the AI Lead and relevant manager or Communications owner.
Staff conduct or protected disclosure	Use the line-management, Staff Code of Conduct, Disciplinary or Whistleblowing route as appropriate.

Do not delete or alter relevant records unless instructed through the incident process. Preserve prompts, outputs, timestamps, screenshots, system notices and audit information where lawful and safe. Safeguarding records follow the Safeguarding and Child Protection Policy. High-risk system logs, once those duties apply, will be retained for at least six months where under School control, unless another lawful retention period applies.

Pupils, parents, applicants and staff may challenge an AI-influenced outcome through the relevant School procedure and must receive meaningful human review. Complaints follow the Complaints Policy. Data-rights requests go to the DPO. Where applicable, CSB will provide the explanation required by Article 86 of the EU AI Act and information about the Article 85 right to complain to the market-surveillance authority. Breaches by staff are managed fairly and proportionately under existing conduct and disciplinary procedures; this policy does not set an automatic minimum sanction.

12. AI LITERACY AND TRAINING

All staff, including staff who do not work directly with pupils, receive this policy and basic AI, data-protection and safeguarding guidance at induction. Refresher training is provided at least annually and when a role, system or risk changes. Completion is recorded.

Training is proportionate to role and use. Additional training is required for leaders, teachers using pupil-facing systems, the DSL team, DPO and IT, Admissions, HR and recruitment, Communications and

social media, procurement owners, and anyone overseeing a higher-risk system. Access to a named agent, AI coding tool, governed exception or higher-risk feature is issued only after the required training is recorded. Article 4 requires measures supporting AI literacy; it does not require CSB to guarantee a particular level or issue a universal certificate.

13. GOVERNANCE, REVIEW AND RELATED POLICIES

The AI Lead reviews the register termly and reports material use, incidents, complaints, training and higher-risk proposals to SLT. The Proprietor or Governing Board receives annual assurance. The policy is reviewed by August 2027, or earlier after a material legal change, final high-risk classification guidance, serious incident, new high-risk use, or significant change in School systems.

This policy must be read with the current versions of:

- CSB Safeguarding and Child Protection Policy; Online Safety Policy; and filtering and monitoring arrangements.
- Data Protection Policy, privacy notices and data-breach procedure; Records Retention Schedule; and Acceptable Use and Information Security policies.
- Staff Code of Conduct; Behaviour; Recruitment and Selection; Disciplinary; Whistleblowing; and Complaints policies.
- Academic Integrity and Assessment Policy and awarding-body rules; and Social Media, Communications, Copyright and Intellectual Property requirements.

APPENDIX A: STAFF QUICK REFERENCE BY WORK AREA

Area	May Use AI For	Must Not
Teaching and resources	Ideas, drafts, quizzes, examples, translation and accessibility adaptations using approved information; staff-created Canva AI resources.	Unverified content; personal pupil data; fabricated sources or observations; directing pupils to unapproved AI features.
Pupil-facing use	Approved tool, clear learning purpose, age suitability, supervision, filters, disclosure and human route.	Unsupervised consumer tools; AI companion or counselling role; secrecy, dependency or replacing a teacher/DSL.
Assessment and reports	Teacher-reviewed formative support and drafting from verified evidence where the workflow is approved.	Final grading, progression or access decisions; invented observations; AI detector as sole or decisive evidence.
Safeguarding and pastoral	Approved administrative support without case data, where it does not assess the concern or determine action.	Entering case details; credibility or risk assessment; referral decisions; responding to a disclosure through AI.
Admissions	General information, translation, document organisation and human-reviewed correspondence.	Applicant scoring, ranking, eligibility, admission, scholarship, placement or educational-level decisions.
Recruitment and HR	Drafting neutral materials and questions from human-set criteria; routine correspondence.	Profiling, targeted ads, CV screening/ranking, video or emotion analysis, performance scoring or material employment decisions.
Communications and social media	Human-reviewed drafts, translations and design concepts; labelled synthetic media.	Auto-publishing, impersonation, cloned voices, fake testimonials or presenting a generated event/place/person as real.
Meetings and administration	Approved templates and an approved transcription workflow with lawful basis and notice.	Covert recording; external meeting bots; personal or confidential data in an unapproved tool.
Finance, IT and operations	Public-information research, controlled analysis, approved coding assistance and tested drafts.	Autonomous transactions/actions, credentials, payroll/payment data, unapproved or untested code, connectors, School-system grounding, biometric or CCTV analysis without approval.

APPENDIX B: ARTICLE 50 DISCLOSURE AND LABELLING GUIDE

All public-facing text receives substantive human review and a named editorial owner. All realistic synthetic media that could appear authentic or truthful receives a visible or audible label. If uncertain, label it and consult the AI Lead before publication.

Situation	Placement and Action	Standard Wording
Direct interaction with AI	At the start of each session or first interaction. The notice must be in context, clear and accessible.	You are interacting with an AI system operated by Cambridge School of Bucharest. / Interacționați cu un sistem de inteligență artificială operat de Cambridge School of Bucharest.
AI agent contacts a person for CSB	At the beginning of the contact, disclose that it is AI and that it acts for CSB. Human confirmation is required before any consequential action.	This is an AI system acting on behalf of Cambridge School of Bucharest. / Acesta este un sistem de IA care acționează în numele Cambridge School of Bucharest.
Realistic content created mainly by AI	Use an AI icon or label on the image/video and in accessible text. For audio, give an audible notice at the beginning.	AI GENERATED / GENERAT DE IA
Realistic content materially altered by AI	Use an AI icon or label. State what was changed where this helps avoid misunderstanding.	AI MODIFIED / MODIFICAT CU IA
Evidently artistic, fictional, satirical or creative deepfake	Disclose in a way that does not hamper the work, such as adjacent text, programme notes, opening or end credits. It must still be clear at first exposure.	Includes AI-generated or AI-modified content. / Include conținut generat sau modificat cu IA.
Public-interest text without qualifying human review	Place a clear label at or near the headline or beginning. Automated publication is not permitted without prior approval.	This publication includes text generated or modified by AI. / Această publicație include text generat sau modificat cu ajutorul IA.
Routine text substantively reviewed by a responsible person	No Article 50 text label is required. Keep the human author/editor responsible and reassess any realistic media separately.	No label required by this policy.

B1. Realistic Synthetic Media Test

Apply these four questions to every AI-generated or AI-modified image, audio or video before it is shared:

1. Was AI used to generate or materially manipulate the content? Ordinary technical editing, such as basic cropping or colour correction, is not enough on its own.
2. Does it depict or represent a real or plausible person, object, place, entity or event? A wholly invented but believable scene can qualify.
3. Could a reasonable person, in the actual context, falsely take it to be authentic or truthful? Consider captions, audience, realism and how it may circulate after reposting.
4. Is it evidently artistic, fictional, satirical or creative? It still needs disclosure, but the label may be placed in a way that does not hamper the work.

B2. Public-Interest Text Test

Treat AI-generated or manipulated text as Article 50 public-interest text only where all three answers are yes:

1. Is it published to an indeterminate and fairly large public audience, rather than sent privately or within the organisation?
2. Is its purpose to communicate knowledge, opinions or facts?
3. Does it concern a matter relevant to society and meriting public debate or scrutiny, such as public services, rights, safety, health, the environment, consumer safety, or significant economic, scientific, political or cultural developments?

Routine advertisements and product descriptions are usually outside this text rule unless they make public-interest claims, for example about health, safety or sustainability. A public school social-media post can still be in scope depending on its content. The channel does not decide the question.

B3. Placement and Accessibility

- Give the disclosure at first interaction or exposure. Do not rely only on terms, a link, hidden metadata or a general website notice.
- For an image, place the label on or immediately beside it with strong contrast and descriptive alternative text.
- For video, show the label at the beginning and after interruptions; keep it visible throughout where feasible, especially if clips or screenshots may circulate.
- For audio, state the disclosure audibly at the beginning and repeat it after interruptions or during long-form content where needed.
- Preserve machine-readable marks, platform labels and content credentials during editing, export and reposting.
- Use plain, child-friendly wording and formats accessible to people with disabilities. For audiences in Romania, provide Romanian; English may be added alongside it.

APPENDIX C: APPROVED TOOL AND USE-CASE SNAPSHOT

This snapshot records the baseline carried into the August 2026 policy. The live AI Tool and Use-Case Register is authoritative and may be more restrictive. Product names and features change; approval never extends automatically to a new account, connector, model, agent or feature.

Tool / Category	Staff Status	Pupil Status	Key Conditions
Microsoft Copilot Chat, School account	Only for registered low-risk uses where organisational-data access is disabled.	No.	Microsoft 365 Copilot or any feature that searches School mail, files or folders is not permitted.
ChatGPT Business / Enterprise School workspace and named agents	Permitted only for registered staff workflows and, for agents, an approved Agent Card.	No.	No connectors or external actions. Use approved data only; review every output. Personal/confidential files need workflow-specific approval.
Personal or consumer AI accounts and agents	Not permitted for School work.	No.	Do not copy, paraphrase, upload or discuss School information or use the account to produce School work.
Canva AI, Canva Code and native Canva assistants/agents	Staff-only resource creation in the School-managed Canva account.	Limited, supervised and task-bounded	No personal data, auto-publishing.
Other AI code, website or interactive-resource generators	Named staff and purpose require prior approval.	No.	Test with synthetic data; check accessibility, privacy and security. Runtime AI/API calls, logins, analytics and publication need separate approval.
Connectors, synced sources, organisational grounding and AI add-ins	Prohibited.	Prohibited.	No AI access to School email, contacts, calendars, files, folders, drives, Teams, SharePoint or other School systems, regardless of account type.
Synthetic audio, image, video or music tools	Feature-specific approval required.	No, unless separately approved.	No voice cloning or impersonation. Apply Appendix B and preserve provenance marks.
Mindjoy or other approved educational chatbot	Permitted for approved learning/activities.	Limited, supervised and task-bounded.	Clear AI disclosure, age suitability, filters, monitoring, staff oversight and a human escalation route.
Microsoft Teams transcription / recording	Limited to approved meeting types and settings.	No, unless separately approved for a defined educational purpose.	Advance notice, lawful basis, restricted access, human-checked minutes and approved retention.

APPENDIX D: APPROVAL AND REVIEW CHECKLIST

The tool sponsor must provide enough information for the reviewers to answer the following questions:

Review Area	Required Evidence
Purpose and benefit	What problem is being solved? What safer or non-AI alternative exists? What evidence supports the intended benefit?
Users and affected people	Which staff, pupils, applicants, families or third parties will use or be affected? What are the ages and accessibility needs?
Data, contract and intellectual property	What enters and leaves? What are the provider's data role, subprocessors, transfers, retention, model-training terms, change notices, export and deletion arrangements? What lawful basis, contract, rights and privacy information apply?
Risk classification	Is the use prohibited, subject to Article 50, potentially high-risk, profiling, biometric, monitoring, agentic or capable of materially influencing a decision?
Safeguarding and security	What filters, alerts, logging, access controls, testing, incident routes and human support exist? Confirm that connectors and external actions are disabled.
Human oversight and redress	Who checks outputs, makes the decision, explains it, corrects error and handles a challenge? Do they have authority, competence and time?
Implementation and review	What training, disclosure, record, retention, owner, expiry date and change triggers will be placed on the register?