CAS AI INITIATIVE

Artificial Intelligence in Education

IB & UNESCO Research Synthesis for CAS

A working reference for the AI Working Team and Committee

2026–2027 | Casablanca American School

Detailed contents

Select any entry to move directly to that section. Section headings throughout the document include a return link to this contents page.

› Executive synthesis
› 1. Mandate, scope and evidence frame
  1.1 What this document is—and is not
  1.2 How sources were weighted
  1.3 CAS mandate crosswalk
› 2. The shared educational position
  2.1 The centre of gravity
  2.2 A practical CAS decision test
› 3. Governance, ethics, rights and safety
  3.1 From values to operating controls
  3.2 Data, privacy and procurement
  3.3 Age, well-being and human oversight
› 4. Academic integrity, authorship and assessment
  4.1 What the IB position requires
  4.2 A usable continuum of assistance
  4.3 Assessment design and authentication
› 5. Teacher competency and professional learning
  5.1 UNESCO teacher competency architecture
  5.2 Differentiated pathways for CAS
  5.3 Workload, agency and organizational learning
› 6. Student competency and AI literacy
  6.1 UNESCO student competency architecture
  6.2 The literacy CAS should protect
  6.3 Evidence of learning
› 7. A Nursery–Grade 12 learning continuum
  7.1 Design principles
  7.2 Proposed developmental bands
  7.3 Integration model
› 8. Pedagogy and practical use
  8.1 Purpose before tool
  8.2 High-value patterns
  8.3 Patterns that require restraint
› 9. Inclusion, equity, language and culture
  9.1 Inclusion as a design requirement
  9.2 Access is more than device availability
  9.3 CAS design responses
› 10. Platform, vendor and data governance
  10.1 Gemini as institutional priority, not curriculum
  10.2 Approval and review lifecycle
  10.3 Minimum procurement questions
› 11. Implementation, pilots and evidence
  11.1 A staged implementation logic
  11.2 What to measure

- 11.3 Evidence standards and review rhythm
- 12. Tensions, gaps and local decisions
  - 12.1 Where the research agrees
  - 12.2 Where CAS must choose
  - 12.3 Open questions for governance
- 13. Consolidated CAS recommendations
  - 13.1 Immediate foundations
  - 13.2 2026–2027 development priorities
  - 13.3 Decisions to defer or constrain
- 14. Blueprint for CAS deliverables
  - 14.1 AI policy and interim guidance
  - 14.2 AI literacy framework
  - 14.3 Nursery–Grade 12 learning continuum
  - 14.4 Professional learning, pilots and reporting
- Appendix A. CAS mandate–evidence crosswalk
- Appendix B. Complete source register
- Appendix C. Working glossary

Executive synthesis

The IB and UNESCO materials converge on a clear position: schools should neither prohibit AI categorically nor adopt it simply because it is available. AI belongs in education only when it advances a defensible learning purpose, preserves human agency and accountability, protects rights and well-being, and produces evidence that the intended benefit is real. For CAS, this means building an education program and governance system together—not treating policy, professional learning, curriculum and platform decisions as separate projects.

The strongest sources are unusually consistent. The IB requires students to own their thinking and to acknowledge AI-generated material used in assessed work. It expects schools and teachers to authenticate work through knowledge of the student, visibility into the process and the student's ability to explain decisions. UNESCO extends the frame beyond academic integrity: privacy, equity, safety, inclusion, cultural and linguistic diversity, environmental impact, human oversight, transparency, participation and the long-term development of human capacities all belong in the decision.

The evidence does not support a single skill level for everyone. UNESCO's student framework moves from Understand to Apply to Create across four aspects; its teacher framework moves from Acquire to Deepen to Create across five aspects. These progressions are non-linear. CAS should therefore define a universal foundation for all learners and staff, then offer differentiated pathways for application and advanced creation. The school should avoid making advanced tool production the default expectation for every student or teacher.

A coherent Nursery–Grade 12 continuum should be spiral, developmental and largely integrated into existing learning. All students need recurring experiences with human agency, safe and ethical use, critical evaluation of outputs, basic ideas about data and algorithms, transparency and appropriate authorship. Older or more ready students can move into system analysis, structured design and creation. Dedicated electives, clubs and enrichment can carry specialized technical depth without overloading the core curriculum.

Ten conclusions for CAS

- **1. Education leads; technology follows.** Every use begins with a learning, professional or operational purpose and an explicit reason AI improves the approach.
- **2. Human agency is the non-negotiable centre.** Students and educators remain responsible for judgment, meaning, relationships, creative direction and high-stakes decisions.
- **3. Transparency must become routine.** CAS needs shared language for disclosure, acknowledgement, process evidence and acceptable assistance across divisions and subjects.
- **4. AI literacy is broader than prompting.** It includes ethics, rights, data, bias, systems, evaluation, disciplinary judgment, citizenship and knowing when not to use AI.
- **5. Teacher learning is infrastructure.** Policy without sustained, differentiated professional learning will be inconsistently interpreted and weakly implemented.
- **6. Age and development matter.** Independent use, adult mediation, data exposure and the cognitive work protected by a task must vary across Nursery–Grade 12.
- **7. Assessment needs redesign, not only detection.** Process visibility, dialogue, drafts, oral explanation and authentic performance are more durable than reliance on AI detectors.
- **8. Inclusion must be designed in.** AI can support language, accessibility and differentiation, but can also deepen unequal access, representation and surveillance.
- **9. Platform approval is a lifecycle.** School-managed Gemini can provide a coherent default, but every use still requires age, privacy, contractual, pedagogical and accessibility review.
- **10. CAS should learn through bounded pilots.** Start small, collect teacher and student evidence, review harms and workload, publish what was learned, then scale, revise or stop.

Recommended operating model

Layer	Purpose	Primary CAS output
Guardrails	Define rights, responsibilities, prohibited or restricted uses, accountability and review.	Interim guidance followed by the AI Policy.
Competencies	Define what teachers and students should understand, apply and create.	CAS AI Literacy Framework.
Progression	Sequence age-appropriate learning from Nursery through Grade 12.	CAS AI Learning Continuum.
Practice	Support useful, discipline-sensitive applications without replacing thinking.	Professional learning, exemplars and pilot protocols.
Evidence	Determine whether uses improve learning, access, workload and well-being without unacceptable harm.	Baseline, pilot evidence, annual report and review cycle.

The immediate priority is coherence, not comprehensiveness: agree on common decision rules, publish interim expectations, establish approval and incident pathways, diagnose staff and student readiness, and begin a small number of observable pilots. The policy, literacy framework and learning continuum can then be built from real CAS evidence while remaining anchored in the sources synthesized here.

1. Mandate, scope and evidence frame

1.1 What this document is—and is not

This is an internal working reference for the CAS AI Working Team and committee. It synthesizes the supplied IB and UNESCO corpus against the 2026–2027 CAS mandate: a coordinated schoolwide approach; responsible and ethical use; policies and procedures; differentiated staff and student learning; practical applications; student and staff leadership; and evidence-informed improvement. It is designed to inform the preliminary implementation outline before policy language and curriculum expectations are finalized.

It is not a substitute for current IB program standards, subject guides, assessment procedures, Moroccan law, CAS safeguarding requirements or contractual review. Where those authorities set a firmer rule, they prevail. UNESCO materials are global frameworks and guidance, not automatically binding school policy. Several IB items are classroom exemplars, and the 2026 IB AI Design Principles are draft consultation material. The document labels these distinctions to prevent accidental overclaiming.

1.2 How sources were weighted

Tier	Use in this synthesis	Typical examples
1 — Core	Defines the main ethical, competency, governance or assessment foundation.	UNESCO competency frameworks, rights report, ethics recommendation and GenAI guidance; current IB assessment position and usage guidance.
2 — Applied	Shows how principles may translate into curriculum, subjects or practice.	IB DP subject guidance and classroom series; UNESCO K–12 mapping, UAE case study, inclusion compendia and Beijing Consensus.
3 — Context	Adds historical, futures, system or implementation perspective.	UNESCO forum reports, ICT/OER policy guidance, futures publications and strategy; older IB technology-enhanced learning material.

Weight is not a judgment of overall quality. It describes applicability to the CAS mandate. A contextual source may contain a powerful idea, but it should not override a more current and directly relevant framework. Likewise, a practical activity may be useful without constituting IB endorsement of a particular platform or method.

1.3 CAS mandate crosswalk

CAS mandate element	Research contribution	Resulting design requirement
Teachers + AI	UNESCO teacher framework; IB subject and classroom guidance.	Differentiated, task-based professional learning with peer support and evidence of classroom transfer.
Students + AI	UNESCO student framework; IB inquiry, ATL, authorship and subject guidance.	A developmental continuum combining agency, ethics, critical evaluation, foundations and design.
Policy + documentation	UNESCO rights, ethics and governance guidance; IB assessment position.	Clear responsibilities, transparency, age/data controls, human oversight, academic integrity and review.
Coordinated approach	UNESCO whole-system, multi-stakeholder and readiness guidance.	Named governance roles, cross-divisional participation, consultation and an annual action cycle.
Differentiation	Non-linear competency levels and contextual localization.	Universal foundations plus supported and advanced pathways; Lower/Upper School distinctions.
Evidence and feedback	UNESCO pilot, monitoring and evaluation recommendations.	Baseline measures, bounded pilots, disaggregated evidence and revise/scale/stop decisions.

2. The shared educational position

2.1 The centre of gravity

Across the strongest IB and UNESCO sources, the educational purpose is human development. AI may extend access, feedback, expression, inquiry and professional capacity, but it should not displace the relationships, struggle, judgment, authorship and participation through which education develops people. UNESCO describes a human-centered approach; emerging IB principles connect use to caring, balance, inquiry, educator agency, safety, transparency and continuous evidence-led adaptation [UN04–UN05; IB34].

This leads to a balanced institutional stance. A ban would fail to prepare learners for an AI-shaped world and would drive use underground. Uncritical adoption would allow product capabilities and market pressure to define educational aims. CAS should instead cultivate informed agency: the ability to understand AI, choose whether and how to use it, evaluate its effects, disclose its role and remain accountable for the result.

The 2026 UNESCO anthology The algorithm in the room sharpens this stance through three horizons: see the algorithm, steer the algorithm and rewrite the algorithm [UN24]. Learners and educators first recognize that AI systems embody human choices and interests; then exercise pedagogical and ethical control over use; and ultimately participate in shaping better purposes, designs and governance. This is a stronger aim than producing efficient tool users.

2.2 A practical CAS decision test

Question	A defensible ‘yes’ requires

Question	A defensible 'yes' requires
Purpose	A clear learning, professional or operational goal—not novelty, pressure or time-saving alone.
Human work	Clarity about the thinking, relationship, practice or judgment that must remain human.
Development	Age, readiness and subject fit; adult mediation where needed.
Rights and safety	Acceptable privacy, data, bias, accessibility, well-being and safeguarding conditions.
Transparency	Users know when AI is involved and can disclose, explain and document its contribution.
Responsibility	A named person retains authority and accountability; high-impact decisions are reviewable.
Equity	A meaningful alternative exists and the design does not create avoidable access or language disadvantage.
Evidence	Success and harm indicators are defined; the use can be reviewed, changed or stopped.

CAS implications

- Adopt a schoolwide purpose-first decision protocol that applies to classroom, professional and operational uses.
- Frame AI literacy as informed human agency, not product proficiency.
- State explicitly that faster output is not automatically better learning or better work.
- Make refusal, limitation and non-AI alternatives legitimate outcomes of the decision process.

3. Governance, ethics, rights and safety

3.1 From values to operating controls

UNESCO's ethics architecture gives CAS a stable base: human rights and dignity; environmental and ecosystem flourishing; diversity and inclusion; peaceful and just societies; proportionality and do no harm; safety and security; fairness; privacy; human oversight; transparency; accountability; literacy; and adaptive, multi-stakeholder governance [UN22]. The operational lesson is that values must appear across the AI lifecycle—from selection and design through use, monitoring, incident response and retirement.

The 2025 rights report organizes system readiness through five connected dimensions: coordination and leadership; content and solutions; capacity and culture; connectivity and infrastructure; and cost and sustainability, with data cutting across all five [UN02]. For CAS, these dimensions prevent a narrow policy that tells individuals how to behave while ignoring institutional choices about procurement, access, training, infrastructure and accountability.

Participation matters. Students, teachers, leaders, families and specialists experience different benefits and risks. UNESCO repeatedly calls for meaningful involvement, especially from children, youth and marginalized groups. Consultation should not merely validate a finished policy; it should surface real practices, trade-offs, language needs, accessibility issues and unintended consequences before rules are fixed.

3.2 Data, privacy and procurement

Generative AI commonly requires users to share prompts, documents or other information with a provider. Institutional approval must therefore address what data is collected, where it is stored, whether it is used to train models, how long it is retained, who can access it, what happens when an account is deleted, and what contractual or technical controls protect children and staff. Users cannot carry this responsibility alone [UN02; UN14; UN22].

CAS should distinguish public consumer tools from school-managed services. A managed account may improve contractual control, identity management and monitoring, but it does not make every use safe. Sensitive personal data, safeguarding records, health information, confidential student work, unreleased assessment materials and high-stakes personnel information require explicit prohibitions or tightly controlled workflows. Data minimization should be the default: use the least information necessary, prefer de-identified or synthetic material, and avoid uploading data simply because a tool makes it convenient.

3.3 Age, well-being and human oversight

UNESCO's 2023 guidance recommends age restrictions for independent use of general-purpose generative AI and states that the minimum threshold should be 13, while recognizing parental permission, provider responsibilities and national decisions [UN14, p. 21]. This should not be misread as a recommendation that independent use is educationally appropriate for every 13-year-old or every task. CAS must combine legal and provider requirements with developmental readiness, adult mediation, safeguarding and the specific data and interaction risks of each service.

The rights literature adds concerns about anthropomorphism, manipulation, dependence, cyberviolence, cognitive offloading, surveillance and effects on creativity and critical thinking. Younger learners may not reliably distinguish generated fluency from expertise or machine simulation from relationship. Teachers should therefore explicitly teach that AI systems do not understand, care or take responsibility in human ways, even when their language feels socially responsive.

High-impact decisions require human determination, explanation and appeal. This includes decisions affecting grades, progression, discipline, admission, safeguarding, well-being, accommodations and employment. AI may support pattern finding or administrative preparation only within approved conditions; it should not be the final authority. The draft IB design principles offer a compatible emerging red line: grading, progression and well-being decisions should never be fully automated [IB34].

CAS implications

- Create an AI governance charter naming decision owners, consultation groups, approval authority, incident pathways and the review calendar.
- Publish a data classification and AI handling rule that identifies information that may never be entered into unapproved systems.
- Set division-specific rules for independent, mediated and prohibited use; do not rely on a single age threshold.
- Guarantee meaningful human review and appeal for high-impact decisions, and disclose material AI involvement.
- Include environmental cost, accessibility and cultural/linguistic fit in tool review—not only privacy and price.

4. Academic integrity, authorship and assessment

4.1 What the IB position requires

The IB's current position is not a ban. Students may use AI, but material produced by an AI tool—text, images, graphs or other outputs—is not the student's own work and must be credited and referenced when included in assessed work [IB01; IB02; IB33; IB35]. Schools remain responsible for authentication. Teachers should know the student's process and capabilities, review development over time and ask students to explain choices and content.

The IB AI Usage Guide offers a particularly useful ethical test: use is acceptable when AI helps a student learn; it is unacceptable when it helps the student pretend to have done work they did not do [IB33]. The distinction is not reducible to the amount of generated text. A small intervention can replace the key

intellectual work of a task, while extensive support may be appropriate in a non-assessed learning process if it is transparent and preserves the intended learning.

4.2 A usable continuum of assistance

Zone	Typical use	CAS expectation
Green — generally supportable	Brainstorming after initial thinking; alternative explanations; practice questions; feedback on a draft; accessibility support; teacher planning.	Teacher/task permission, safe platform, verify output, disclose when material to the work.
Amber — conditional	Outlining, summarizing sources, translation, substantial editing, generated images/code/data, simulated dialogue, feedback close to assessment.	Explicit task-specific rules, process evidence, acknowledgement, student explanation and protection of the assessed construct.
Red — prohibited or tightly restricted	Submitting generated work as one's own; invented citations/data; hidden paraphrase; bypassing required reading/thinking; AI-written reflection; unauthorized help in controlled assessment; fully automated high-stakes judgment.	Do not use. Apply academic-integrity, assessment or governance procedures as appropriate.

A color model is useful only as a communication aid. The same activity may change zones depending on the learning objective, age, assessment conditions and subject. Translation, for example, can provide equitable access in one task and invalidate evidence of language proficiency in another. Teachers need authority to specify the permitted zone for a task and the responsibility to explain why.

4.3 Assessment design and authentication

AI detection should not be the centre of the CAS response. Detection systems can be unreliable and do not reveal whether a student learned. More durable strategies are process-rich: staged proposals, source notes, annotated prompts or interaction records where appropriate, version history, drafts, conferences, oral defence, in-class components, reflective explanation and comparison between supported and independent work.

Assessment redesign should protect the construct. If the aim is close reading, initial interpretation may need to precede AI interaction. If the aim is mathematical reasoning, students should verify with independent reasoning and another method rather than accept an answer. If the aim is reflection, generated prose defeats the purpose. If the aim is critique of AI, interaction with a model may be central. The DP subject guidance repeatedly supports this discipline-sensitive approach [IB37–IB44].

CAS should use a simple disclosure convention appropriate to age: name the tool, describe what it contributed, identify what the student changed or verified, and provide the required citation or acknowledgement. For younger students, teachers may document class-level use; older students should increasingly maintain their own records. The rule should be proportionate and should not turn every low-risk classroom interaction into paperwork.

CAS implications

- Publish common authorship and disclosure principles, then allow programs and subjects to specify task-level conditions.
- Build authentication into assessment design through process evidence and dialogue rather than treating it as a forensic step after submission.
- Define restricted categories clearly: hidden substitution of thinking, fabricated evidence, unauthorized controlled-assessment support and AI-generated personal reflection.
- Provide models of age-appropriate acknowledgement and examples from multiple subjects and languages.
- Make academic-integrity responses educational and proportionate while preserving existing IB requirements.

5. Teacher competency and professional learning

5.1 UNESCO teacher competency architecture

UNESCO organizes teacher competency across five aspects and three levels [UN05]. The aspects are a human-centered mindset, ethics of AI, AI foundations and applications, AI pedagogy, and AI for professional development. The levels—Acquire, Deepen and Create—describe increasing capacity, not job rank. A teacher may be advanced in pedagogy and still need foundational knowledge about data or ethics.

Aspect	Acquire	Deepen	Create
Human-centered mindset	Human agency	Human accountability	Social responsibility
Ethics	Ethical principles	Safe and responsible use	Co-creating ethical rules
Foundations/applications	Basic techniques and applications	Application skills	Creating with AI
AI pedagogy	AI-assisted teaching	AI–pedagogy integration	AI-enhanced pedagogical transformation
Professional development	AI supporting professional development	AI enabling lifelong professional learning	AI enhancing organizational learning

UNESCO expects foundational acquisition for all teachers while deeper and creative levels depend on role, experience, interest and context. Progression is non-linear and should follow diagnosis. Training disconnected from real work is unlikely to transfer; teachers need opportunities to apply learning to planning, feedback, assessment, communication and subject-specific pedagogy, then examine student evidence and unintended effects.

5.2 Differentiated pathways for CAS

Pathway	Who it serves	Core experiences	Evidence of progress
Building Foundations	All staff, especially hesitant or beginning users.	What AI is/is not; approved access; privacy; bias; verification; authorship; one useful task; one deliberate non-use decision.	Safe completion of a real task, reflection on limits and correct response to a scenario.
Advancing Practice	Educators ready to integrate AI into learning design.	Subject-specific pedagogy; assessment redesign; differentiation; prompt and output evaluation; pilot design.	Student work/process evidence, teacher analysis, peer review and adjustment.
Leading and Creating	Coaches, early adopters, coordinators and interested specialists.	Tool evaluation, coaching, resource design, curriculum mapping, ethical-rule co-creation, structured experimentation.	Reusable resources, documented pilot results, mentorship and contribution to governance.

The first pathway should be universal but not remedial in tone. It establishes a common professional baseline. The second and third should remain flexible: not every excellent teacher needs to become an advanced AI creator. CAS should recognize disciplined non-use, careful limitation and pedagogical judgment as evidence of competence alongside effective use.

5.3 Workload, agency and organizational learning

Teacher workload is a legitimate design consideration, but automation can remove activities through which teachers learn about students. Marking, feedback and planning are not merely clerical tasks. Before automating, CAS should ask what professional knowledge the activity produces and how that knowledge will

be preserved. Teachers should retain final responsibility for feedback, evaluation, relationship-based communication and instructional decisions.

Professional learning should become organizational learning. Monthly Working Team meetings can examine a small number of cases, compare evidence, update shared guidance and identify needs. Peer coaching and communities of practice are better suited to a rapidly changing field than occasional one-off workshops. A curated internal repository should include approved use patterns, subject examples, disclosure models, pilot findings and decisions that were reversed as well as successes.

CAS implications

- Use a diagnostic organized by the five UNESCO aspects rather than a single confidence rating.
- Guarantee a common foundation for all staff and offer differentiated pathways aligned with CAS's two engagement levels.
- Make professional learning task-based, subject-sensitive and connected to student evidence.
- Protect educator agency and final judgment; do not evaluate staff by frequency of AI use.
- Build peer coaching, Working Team inquiry and a living resource repository into the annual action plan.

6. Student competency and AI literacy

6.1 UNESCO student competency architecture

UNESCO's student framework combines four aspects with three levels [UN04]. The aspects are a human-centered mindset, ethics of AI, AI techniques and applications, and AI system design. The levels—Understand, Apply and Create—move from recognition and explanation to responsible use and, where appropriate, the design or creation of AI-enabled solutions.

Aspect	Understand	Apply	Create
Human-centered mindset	Human agency	Human accountability	Citizenship in the AI era
Ethics	Embodied ethics	Safe and responsible use	Ethics by design
Techniques/applications	AI foundations	Application skills	Creating AI tools
System design	Problem scoping	Architecture design	Iteration and feedback

UNESCO treats Understand as a foundation for all learners and Apply as broadly relevant. Create may be specialized or elective because it requires time, resources and interest. CAS can therefore promise universal literacy without promising identical technical depth. The four aspects should remain connected: ethics without enough technical understanding can become superficial, while technical skill without ethics, agency and design responsibility is incomplete.

6.2 The literacy CAS should protect

- **Agency and purpose:** decide when AI is useful, when it is not, and what human work must be protected.
- **Critical evaluation:** check claims, evidence, bias, omissions, uncertainty and fit for context rather than equating fluency with truth.
- **Data and systems:** understand that outputs arise from data, models, optimization choices, interfaces and human institutions.
- **Ethics and rights:** apply privacy, consent, fairness, safety, dignity, sustainability, transparency and accountability in concrete cases.
- **Authorship and participation:** disclose AI contributions, preserve one's own voice, explain decisions and contribute to shared rules.

- **Disciplinary judgment:** recognize that appropriate evidence, method, expression and assistance differ across subjects.
- **Design and creation:** scope problems, identify stakeholders, test assumptions, iterate and evaluate effects when building AI-enabled solutions.
- **Citizenship:** examine power, labour, culture, language, environmental costs, misinformation and who benefits or bears risk.

6.3 Evidence of learning

AI literacy cannot be measured reliably through a multiple-choice definition test alone. UNESCO recommends contextual, authentic and performance-based approaches. Students might compare generated and authoritative accounts, explain a recommendation system, conduct a privacy review, identify bias in a training set, debate a stakeholder dilemma, document changes made to an AI-assisted product, design class rules or defend why AI was excluded from a task.

Assessment should separate assisted performance from independent capability when that distinction matters. Students should sometimes work without AI to show foundational fluency and sometimes work with AI to show evaluation, orchestration and ethical judgment. Results should be examined by age, language background, learning needs and access so apparent gains do not conceal unequal participation.

CAS implications

- Base the CAS AI Literacy Framework on all four UNESCO student aspects, adapted to the IB Learner Profile and approaches to learning.
- Define a universal Understand foundation, broad Apply opportunities and optional/specialized Create pathways.
- Use authentic demonstrations, explanations and portfolios of process as evidence—not tool-use frequency.
- Assess both supported and independent performance when foundational competence could otherwise be hidden.
- Include citizenship, power, culture, language and sustainability rather than limiting literacy to safe prompting.

7. A Nursery–Grade 12 learning continuum

7.1 Design principles

The learning continuum should be developmental, spiral and local. Concepts return with increasing complexity; students encounter them through age-appropriate examples and authentic subject contexts. The curriculum should not depend on one vendor or assume continuous direct chatbot use. Many foundational ideas—patterns, classification, data, fairness, authorship, source evaluation, human choice and problem scoping—can be learned unplugged or through teacher-mediated experiences.

The UNESCO K–12 mapping found that government-endorsed curricula commonly integrate AI into existing subjects or offer it as an elective, and that teacher training and resources are essential. It recommends integration maps across subjects and grade levels, teacher participation, pilot testing and rigorous evaluation [UN21]. The UAE case study reinforces the importance of planned progression, infrastructure and teacher capacity while reminding CAS not to import a national curriculum unchanged [UN06].

7.2 Proposed developmental bands

Band	Developmental emphasis	Illustrative learner evidence

Band	Developmental emphasis	Illustrative learner evidence
Nursery–Grade 2 | Encounter and care	People make tools; machines follow patterns; private information; asking an adult; fair/unfair examples; created versus found media.	Names a trusted adult, distinguishes person/tool, sorts examples, explains why some information stays private.
Grades 3–5 | Notice, question and verify	Inputs/outputs, simple data patterns, mistakes and bias, authorship, source checking, teacher-mediated AI comparison.	Finds an error or omission, records help received, compares two outputs, proposes a fair-use class rule.
Grades 6–8 | Use responsibly and explain	Model limits, privacy/consent, prompting as communication, verification, recommendations, deepfakes, discipline-specific rules, problem scoping.	Explains an AI-supported process, verifies claims, identifies stakeholders and consequences, produces an acknowledgement.
Grades 9–10 | Analyze systems and impacts	Data/model/interface relationships, bias testing, automation and labour, environmental cost, assessment design, iterative projects.	Conducts an impact or bias analysis, defends tool/non-tool choice, documents iteration, distinguishes evidence from generated assertion.
Grades 11–12 | Exercise judgment and contribute	Advanced disciplinary use, research and citation, governance, ethics by design, human oversight, civic participation, creation where appropriate.	Defends authorship, audits a system or proposal, co-creates policy, builds/tests a bounded solution, evaluates benefits and harms.

The bands are a design proposal, not a finished scope and sequence. The final continuum should identify prerequisite concepts, division-specific access rules, subject ownership, protected independent competencies and a small set of common demonstrations. It should remain adaptable for students who enter CAS at different grades and for learners with different language, accessibility and prior-experience profiles.

7.3 Integration model

Curriculum home	Best use	Risk to manage
Transdisciplinary / advisory	Shared norms, citizenship, well-being, media literacy, privacy and reflection.	Everyone assumes someone else teaches the technical or disciplinary detail.
Subject integration	Authentic methods, evidence, authorship and appropriate assistance in each discipline.	Fragmented terminology and inconsistent expectations.
Dedicated modules	Common foundations in systems, data, ethics and evaluation.	Crowding the timetable or becoming detached from real subject use.
Electives / enrichment	Advanced creation, coding, design, competitions and mentorship.	Advanced opportunity becomes available only to already-advantaged students.
Student leadership	Peer education, policy consultation, showcases and community inquiry.	Token participation without real influence or adequate support.

CAS implications

- Use a hybrid model: common foundations, subject integration, advisory/citizenship work and optional advanced pathways.
- Map existing curriculum before adding content; identify what can be integrated, condensed or protected from displacement.
- Write the continuum around durable concepts and performances, with Gemini-specific instructions kept in replaceable resources.
- Pilot proposed demonstrations with students at different grade levels and revise based on comprehension, access and workload.
- Build bridge experiences for students entering CAS without earlier stages of the continuum.

8. Pedagogy and practical use

8.1 Purpose before tool

The IB materials are most useful when read as patterns rather than recipes. Their recurring logic is to begin with pedagogy, make the intended student thinking visible, use AI for a bounded function, and require teacher or student evaluation. Current DP subject guidance repeatedly recommends independent interpretation or effort before AI support, preservation of disciplinary voice, attention to bias and source quality, and explicit responsibility for final judgment [IB37–IB44].

Examples include using AI to generate alternative perspectives, hints rather than full solutions, counterarguments, questions, feedback, simulations or representations. In language and literature, initial close reading should precede AI comparison. In mathematics, an answer should be checked through mathematical reasoning and an independent method. In philosophy and TOK, AI is useful as an object of inquiry or a fallible interlocutor, not a source of final answers. In arts subjects, it can support ideation or production while human creative direction, consent and process evidence remain central.

8.2 High-value patterns

Pattern	Why it can add value	Required safeguard
Multiple representations	Offers another explanation, reading level, language support, example or visualization.	Verify accuracy; preserve access to authoritative content and independent sense-making.
Feedback for revision	Creates rapid formative input and opportunities to evaluate advice.	Student makes and explains decisions; teacher retains assessment judgment.
Inquiry and perspective	Generates questions, counterarguments, stakeholders or hypotheses.	Begin with student thinking; investigate sources; identify stereotypes and omissions.
Practice and simulation	Enables low-stakes rehearsal, dialogue and adaptive examples.	Do not anthropomorphize; monitor quality, emotional risk and data exposure.
Teacher preparation	Drafts materials, examples, questions, communications or differentiation options.	Teacher verifies, contextualizes, protects data and owns the final product.
Design and prototyping	Supports problem scoping, rapid iteration, code or media prototypes.	Document contributions, test outputs, consider stakeholders, rights and environmental cost.

8.3 Patterns that require restraint

- Generating the central interpretation, argument, research question or reflection that the task is meant to develop.
- Providing full solutions before a learner has attempted the reasoning that creates durable understanding.
- Replacing reading of primary or assigned sources with summaries that the learner cannot verify.
- Using simulated care, counselling or authority where a qualified human relationship is required.
- Automating grading, progression, discipline, safeguarding or well-being decisions.
- Uploading student information, protected materials or sensitive work to unapproved services.
- Scaling an attractive demonstration before testing learning, access, workload and unintended effects.

The IB classroom series supplies many concrete activities across PYP, MYP, DP and CP—unit planning, culturally relevant lesson ideas, lines of inquiry, expository writing, complex questioning, moral dilemmas, response refinement, simple model training, perspective shifts, historical simulations, service learning and career preparation [IB04–IB32]. These are a useful idea bank. They are not blanket endorsement of a named product, and several require adaptation for current privacy, age, accessibility and academic-integrity conditions.

CAS implications

- Curate exemplars by learning function and subject, not by product name.
- Require every exemplar to identify the protected human work, disclosure expectation, data boundary and evidence of learning.
- Use ‘think first’, ‘verify by two’ and ‘explain the process’ as memorable cross-curricular routines where appropriate.
- Treat the IB classroom activities as starting points for local design and review, not ready-made CAS policy.
- Include disciplined non-use examples in professional learning so judgment is visible.

9. Inclusion, equity, language and culture

9.1 Inclusion as a design requirement

AI can expand access through language support, alternative representations, assistive features, adaptive practice and reduced barriers to production. The UNESCO inclusion compendia illustrate this potential across varied contexts [UN07–UN08]. Yet the same systems can reproduce historical bias, marginalize local knowledge and languages, misread disability or cultural difference, and allocate opportunity unevenly. Inclusion is therefore not a later accommodation; it is part of tool selection, curriculum design, evidence collection and governance.

CAS’s multilingual international context makes linguistic and cultural representation especially important. Students should learn that dominant-language fluency is not neutral evidence of quality and that models may flatten traditions or omit locally meaningful perspectives. Educators should compare outputs across languages where possible, protect students’ developing voices and avoid making standardized machine prose the implicit norm of academic quality.

9.2 Access is more than device availability

Access includes reliable connectivity, approved accounts, compatible devices, accessible interfaces, language support, time, adult guidance, prior knowledge and the freedom to choose an alternative. Home access may differ sharply even when school access appears equal. Assignments that assume unrestricted personal AI use can transfer costs and risks to families and create hidden advantage for students with paid tools or expert support.

Evidence should be disaggregated. If a pilot improves average performance, CAS still needs to ask who participated, who opted out, whose workload increased, whether independent competence changed and whether specific language or learning profiles were helped or disadvantaged. Student voice is essential, but participation must be accessible and consequential rather than symbolic.

9.3 CAS design responses

- Provide a functionally equivalent non-AI route when AI is not essential to the assessed construct.
- Use school-managed access for required experiences; do not require students to create unapproved consumer accounts.
- Test interfaces and activities with multilingual learners and students with accessibility needs before scaling.
- Invite local and underrepresented perspectives into examples, datasets, evaluation criteria and policy consultation.
- Distinguish translation for access from translation that replaces a language-learning objective.
- Monitor participation, outcomes, confidence, well-being and independent capability by relevant student groups.

CAS implications

- Add an inclusion and accessibility review to every approved tool and pilot proposal.
- Guarantee equitable school access or an equivalent alternative for required work.
- Include cultural and linguistic bias explicitly in the AI Literacy Framework and source-evaluation routines.
- Use disaggregated evidence and meaningful student consultation before declaring an intervention successful.
- Avoid framing standardized, fluent output as the benchmark for voice, intelligence or learning.

10. Platform, vendor and data governance

10.1 Gemini as institutional priority, not curriculum

A school-managed Gemini environment can give CAS a coherent default for training, support, identity management and common practice. That operational choice should not define the learning continuum. Durable competencies—purpose, verification, data judgment, authorship, ethics, systems thinking and evaluation—must transfer across tools and remain useful when products, interfaces and contractual terms change.

Institutional priority should mean ‘supported and governed default’, not ‘mandatory for every task’. Some uses will require a different approved specialist tool; other tasks should remain non-AI. Teacher and student agency includes the ability to justify those choices within shared guardrails.

10.2 Approval and review lifecycle

Stage	CAS action	Decision output
Propose	State educational/operational purpose, users, data, age, alternatives and owner.	Complete intake; no use implied.
Screen	Check terms, privacy, security, safeguarding, accessibility, bias, cost and interoperability.	Reject, request evidence or proceed to bounded test.
Pilot	Limit population, duration, data and scope; provide training and support.	Documented evidence and incidents.
Review	Compare benefit, harm, workload, equity and independent capability against baseline.	Scale, revise, contain or stop.
Operate	Publish approved conditions; monitor changes, incidents and usage.	Continued conditional approval.
Renew / retire	Reassess terms, model behavior, cost, need and data disposition.	Renewal, migration or closure plan.

10.3 Minimum procurement questions

- What precise problem does the service solve, and what evidence supports that claim for learners like ours?
- What personal, behavioral, content and metadata are collected; where are they stored; who controls them; are they used for model training?
- What are the age requirements, parental-consent conditions, moderation controls and incident-response commitments?
- Can users understand material AI involvement, correct errors, appeal outcomes and obtain human support?
- How does the service perform across CAS languages, cultures, disabilities, devices and connectivity conditions?

- What teacher workload, training, licensing, integration, environmental and exit costs accompany adoption?
- How will CAS retrieve or delete data, preserve records and continue if the vendor changes terms, price or functionality?
- Who at CAS owns the decision and the ongoing monitoring?

CAS implications

- Establish one visible list of approved, restricted and unapproved AI services with conditions of use.
- Treat approval as conditional and time-limited; monitor material changes in products and terms.
- Keep the AI Literacy Framework vendor-neutral while providing replaceable Gemini job aids for current practice.
- Require data minimization and prohibit sensitive uploads unless an explicitly approved workflow exists.
- Plan for exit, data deletion and continuity before adoption.

11. Implementation, pilots and evidence

11.1 A staged implementation logic

UNESCO policy and curriculum guidance consistently favors readiness assessment, stakeholder participation, capacity-building, local pilots, monitoring and iterative review [UN01; UN04–UN05; UN14; UN21]. CAS's preliminary outline follows this logic: confirm mandate and governance; establish a baseline; publish interim guidance; define the framework; build an annual action plan; consult and develop; run targeted learning and pilots; evaluate and refine; and secure continuity.

Phase	Primary question	Minimum evidence before moving
1. Foundation	Do we know who decides, what is happening now and what immediate risks need guidance?	Governance charter, baseline, interim guidance, incident/approval pathway.
2. Definition	What should CAS teachers and students know and be able to do?	Draft literacy framework, learning-continuum architecture, policy principles and consultation record.
3. Capacity	Can people apply the expectations in real work?	Differentiated learning participation, scenario performance, coaching capacity and usable resources.
4. Pilot	Does a bounded use improve the intended outcome in our context?	Baseline/comparison, artifacts, participant voice, workload, equity, incidents and limitations.
5. Scale or revise	Is the benefit sufficiently repeatable and responsible?	Decision record, conditions, support plan, monitoring indicators and review date.
6. Institutionalize	Can CAS sustain, govern and improve the work?	Annual report, updated policy/resources, budget/ownership, handover and archived evidence.

11.2 What to measure

Domain	Possible indicators	Guard against
Learning	Quality of reasoning, transfer, independent performance, revision, disciplinary understanding.	Equating polished output or speed with learning.
Agency and integrity	Disclosure quality, explanation, appropriate non-use, confidence to question outputs.	Counting prompts or surveillance events as competence.
Teaching and workload	Planning/feedback quality, time shifted, teacher knowledge of students, job satisfaction.	Reporting time saved without examining work added or professional knowledge lost.

Domain	Possible indicators	Guard against
Equity and inclusion	Access, participation, outcomes and experience by division, language and learning need.	Averages that conceal uneven benefit or harm.
Safety and rights	Incidents, sensitive-data exposure, inappropriate content, appeal, well-being and accessibility.	Assuming no report means no harm.
Implementation	Staff readiness by competency aspect, support use, pilot fidelity, resource quality, decision turnaround.	Training attendance as the only measure of capacity.

11.3 Evidence standards and review rhythm

A school pilot does not require a formal research study, but it does require a claim, baseline, bounded intervention, evidence plan and decision rule. Evidence should triangulate artifacts or performance, participant experience and implementation conditions. Novelty effects and self-reported enthusiasm are insufficient on their own. Negative and null findings are valuable because they prevent unproductive scale.

Monthly Working Team meetings can review a small evidence dashboard and one or two cases. Each term, the committee can examine progress against the annual action plan, emerging incidents, policy questions and access patterns. The annual report should state what was tried, what was learned, what changed, what stopped and what remains uncertain. This creates institutional memory rather than a sequence of disconnected experiments.

CAS implications

- Give every pilot an owner, question, baseline, success/harm indicators, data boundary, end date and scale/revise/stop decision.
- Measure learning and independent capability, not only engagement, output quality or time saved.
- Use teacher and student evidence together and disaggregate where relevant.
- Review implementation monthly, report to the committee each term and produce an annual evidence-informed update.
- Document decisions that do not scale so CAS learns from restraint as well as success.

12. Tensions, gaps and local decisions

12.1 Where the research agrees

- Human agency, dignity, accountability and educational purpose must remain central.
- Students and teachers need broad AI literacy, not only operational skill.
- Transparency, verification, privacy, fairness and human oversight require explicit practice.
- Teacher capacity and participation are essential to implementation.
- Curriculum should be developmental, contextualized and tested locally.
- High-stakes decisions should not be fully automated.
- Evidence, monitoring and revision are continuing responsibilities.

12.2 Where CAS must choose

The sources do not prescribe a single school policy, platform, disclosure format, timetable allocation or grade-by-grade sequence. UNESCO offers global frameworks; the IB gives requirements and disciplinary guidance but leaves significant school-level judgment. CAS must decide the level of independent use by division, how disclosure works in different age bands, which data categories are prohibited, who approves tools, how required access is made equitable, and which competencies are common versus subject-owned.

There is also a productive tension between preparation for an AI-rich world and protection of foundational human capacities. The answer is not a permanent percentage of 'AI work' versus 'non-AI work'. It is deliberate task design. CAS should name the capability under development, identify when assistance is educationally useful and preserve independent practice when fluency, memory, attention, interpretation, expression or relationship is the objective.

12.3 Open questions for governance

Decision	Why it cannot be copied from the research	Proposed CAS process
Independent-use thresholds	Provider terms, law, maturity, task and safeguarding all matter.	Cross-divisional proposal reviewed by safeguarding, technology, leadership and families.
Disclosure standard	Appropriate detail varies by age, subject and assessment.	Common minimum plus program/subject exemplars tested with students.
Sensitive-data boundary	Local systems, contracts and responsibilities determine risk.	Data classification aligned with IT/security and safeguarding procedures.
Assessment conditions	The assessed construct differs by task and program.	Teacher-set task conditions within schoolwide authorship principles and IB rules.
Tool approval	Product terms and capabilities change quickly.	Time-limited approval with named owner and change monitoring.
Curriculum time	Integration depends on existing CAS scope and sequence.	Curriculum audit followed by integration maps and protected advanced pathways.

13. Consolidated CAS recommendations

The recommendations below consolidate the chapter-level implications into a sequenced agenda. They are preliminary recommendations based on the supplied research and the CAS mandate; they should inform, not bypass, consultation and formal approval.

13.1 Immediate foundations

Priority	Action	Primary link to evidence chapters
1	Approve a governance charter: mandate, roles, decision rights, consultation, tool approval, incident handling, review cycle and reporting.	Sections 3, 10 and 11
2	Issue concise interim guidance covering approved access, privacy/data, independent-use conditions, authorship/disclosure, assessment, high-impact decisions and help pathways.	Sections 3 and 4
3	Complete a baseline of staff/student use, confidence and needs organized by competency aspects and division; include access and concern data.	Sections 5, 6 and 11
4	Create a visible approved/restricted/unapproved tool register and a time-limited review process, beginning with managed Gemini conditions.	Section 10
5	Select a small portfolio of bounded pilots with clear questions, owners, evidence, data limits and end dates.	Sections 8 and 11

13.2 2026–2027 development priorities

Priority	Action	Completion evidence

Priority	Action	Completion evidence
6	Build the AI Literacy Framework by adapting the UNESCO student and teacher architectures to the IB Learner Profile, ATL and CAS context.	Published competency statements, examples and evidence descriptors.
7	Develop the Nursery–Grade 12 continuum through curriculum mapping, developmental bands, bridge experiences and pilot demonstrations.	Cross-divisional progression with subject/advisory ownership and curriculum maps.
8	Create differentiated professional-learning pathways: Building Foundations, Advancing Practice, and leadership/creation opportunities.	Task-based demonstrations, coaching network and documented transfer to practice.
9	Redesign assessment guidance around protected constructs, task conditions, process evidence, explanation and age-appropriate disclosure.	Common principles plus program/subject exemplars and moderation cases.
10	Embed inclusion, language, accessibility and alternatives in every tool, curriculum and pilot review.	Completed review criteria and disaggregated participation/outcome evidence.
11	Use monthly Working Team inquiry, termly committee review and an annual public-to-the-committee evidence report.	Decision log, updated resources, annual findings and next-year action plan.
12	Design continuity: owners, budget, repositories, review dates, induction and handover.	Sustainable responsibilities and maintained records beyond individual champions.

13.3 Decisions to defer or constrain

- Do not set a schoolwide requirement for frequent AI use or evaluate teachers by adoption volume.
- Do not make one product's interface or prompt syntax the organizing structure of the literacy framework.
- Do not rely on AI detectors as the primary academic-integrity response.
- Do not automate final grading, progression, discipline, safeguarding, well-being or employment decisions.
- Do not require personal consumer accounts or sensitive-data uploads for school learning.
- Do not scale pilots on enthusiasm, polished outputs or time saved without evidence of learning, equity, workload and harm.
- Do not finalize a detailed grade-by-grade continuum before auditing existing curriculum and testing developmental assumptions.

14. Blueprint for CAS deliverables

14.1 AI policy and interim guidance

Required element	What it should establish
Purpose and principles	Human-centered, inquiry-driven, rights-respecting, transparent and evidence-led use.
Scope and definitions	Who, which systems and which school activities are covered; distinguish assistance, generation and automated decision-making.
Roles and accountability	Student, staff, leader, IT/data, safeguarding, committee and vendor responsibilities.
Access and age	Independent, supervised and prohibited use by division/service; consent and alternatives.
Data and privacy	Classification, prohibited uploads, approved accounts, minimization, retention, incident response.
Authorship and assessment	Disclosure, citation, task conditions, authentication and educationally proportionate responses.

Required element	What it should establish
High-impact uses	Human review, explanation, appeal and categories that may not be fully automated.
Approval and monitoring	Tool lifecycle, pilot requirements, change review, complaints and annual policy revision.

Interim guidance should be shorter than the final policy and answer immediate operational questions. It should clearly state where to seek help and how uncertainty is handled. The final policy can then incorporate consultation and pilot evidence rather than pretending the school already knows every future case.

14.2 AI literacy framework

The framework should contain two connected profiles: student competencies across human-centered mindset, ethics, techniques/applications and system design; and teacher competencies across human-centered mindset, ethics, foundations/applications, pedagogy and professional learning. Each competency should include a concise statement, level descriptors, age/role-relevant examples, observable evidence and links to the Learner Profile, approaches to learning and applicable policies.

CAS should avoid turning the framework into a checklist of tools. The progression should describe transferable judgment and performance. Gemini job aids and classroom examples can sit in an accompanying resource layer that is easier to update.

14.3 Nursery–Grade 12 learning continuum

The continuum should translate the student framework into developmental expectations without forcing identical experiences in every grade. Begin with the five proposed bands in Section 7, then map existing curriculum, identify gaps and overlaps, assign ownership, define a small number of common demonstrations and create bridge experiences. Publish both a one-page progression and a detailed implementation version.

A balanced model gives all students integrated foundational learning and provides optional technical depth through electives, clubs, competitions, independent projects and mentorship. Access to advanced pathways should be actively broadened.

14.4 Professional learning, pilots and reporting

Professional learning should be annual and cyclical rather than a launch event. Foundation learning establishes common safety, integrity and evaluation practices; advancing pathways focus on subject design and assessment; leaders support coaching, governance and evidence. Pilots should use a common one-page protocol. Reporting should connect activity to learning, agency, equity, workload and rights, and should explicitly record revise/scale/stop decisions.

Deliverable	First useful version	Mature version
Interim guidance / AI policy	Immediate guardrails and help pathways.	Consulted, approved policy with lifecycle governance and review schedule.
AI Literacy Framework	Adapted domains, levels and shared vocabulary.	Evidence descriptors, exemplars, self-assessment and links to curriculum/PD.
N–G12 continuum	Developmental bands and pilot demonstrations.	Mapped progression, subject ownership, bridge experiences and evaluated resources.
Professional learning	Common foundation plus first advancing cohorts.	Differentiated pathways, peer coaching and organizational learning cycle.
Evidence system	Baseline, pilot template and decision log.	Disaggregated indicators, term reviews, annual report and continuity records.

Appendix A. CAS mandate–evidence crosswalk

This crosswalk shows where the main design claims originate. Source keys resolve to the complete register in Appendix B. The references are intentionally selective: they identify the strongest supporting sources rather than every document that touches the issue.

CAS workstream	Primary evidence	Key design claims	Relevant recommendations
Governance and policy	UN01, UN02, UN14, UN22, UN24; IB34	Whole-system leadership; rights across the lifecycle; participation; human oversight; adaptive review.	1, 2, 4, 10–12
Academic integrity and assessment	IB01, IB02, IB33, IB35, IB37–IB44; UN14	AI output is not automatically student work; disclose; authenticate process; protect the construct; redesign assessment.	2, 9
Teacher competency	UN05; IB03–IB32, IB37–IB44	Five aspects; Acquire/Deepen/Create; differentiated, contextual, task-based and career-long learning.	3, 8, 11
Student competency	UN04, UN26; UN22, UN24	Four aspects; Understand/Apply/Create; agency, ethics, foundations, design and citizenship.	6, 7
N–G12 curriculum	UN04, UN06, UN21; IB04–IB32	Spiral progression; integration maps; electives for specialized creation; teacher involvement; pilot testing.	7
Inclusion and rights	UN02, UN07, UN08, UN14, UN22	Access, privacy, language, culture, accessibility, alternatives, protection and participation.	2, 4, 10
Tool and vendor governance	UN01, UN02, UN14, UN22; IB34	Purpose, data minimization, auditability, age, contractual controls, transparency, exit and monitoring.	1, 2, 4
Implementation and evidence	UN01, UN04–UN06, UN21; IB03, IB34	Readiness diagnosis, bounded pilots, teacher/student feedback, disaggregated evidence, iterative scale.	3, 5, 8, 11, 12

Appendix B. Complete source register

All 70 research files in the two supplied folders were opened and reviewed. The register records how each was used and the main applicability caution. Six Word files are activity companions and contain no PDF page count; they are nevertheless included. CAS mandate and brand documents are inputs to the assignment, not part of the 70-source research corpus.

B1. International Baccalaureate sources

Key	Source	Tier	Contribution	Applicability caution	pp.
IB01	Artificial intelligence (AI) in learning, teaching, and assessment - International Baccalaureate®	1 — Core	Current IB position or operational academic-integrity guidance.	Apply directly where it governs IB work; CAS still needs local procedures.	3
IB02	Artificial intelligence in IB assessment and education a crisis or an opportunity - International Baccalaureate®	1 — Core	Current IB position or operational academic-integrity guidance.	Apply directly where it governs IB work; CAS still needs local procedures.	4
IB03	How AI is reshaping our classrooms - International Baccalaureate®	2 — Applied	Practitioner account of purpose-first use, formative feedback and phased adoption.	IB-hosted example, not formal policy or independently evaluated evidence.	8
IB04	Using-ai-to-create-an-instructional-video	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	7

Key	Source	Tier	Contribution	Applicability caution	pp.
IB05	Using-ai-to-generate-a-unit-plan	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	8
IB06	Using-ai-chatbots-to-generate-culturally-relevant-lesson-ideas	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	5
IB07	Using-ai-to-generate-questions	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	6
IB08	CP - Using-diffit-to-scaffold-content-for-students-learning-in-their-native-language	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	2
IB09	CP - Using-trello-and-its-automation-butler-to-develop-self-management-skills-in-cp-students	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	3
IB10	CP - Using-chat-gpt-to-refine-and-format-resumes-and-cover-letters-for-students-linked-in-profiles	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	2
IB11	CP - Using-chat-gpt-to-generate-ideas-for-students-service-learning-projects-in-the-cp	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	3
IB12	DP - The-five-paragraph-essay-framework	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	4
IB13	DP - Point-of-view-change	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	3
IB14	DP - Interview-with-an-historical-figure-or-famous-author-scientist-mathematician	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	3
IB15	DP - You-are-the-artist-image-creation-from-literature-imagery-setting-scene-s-and-character	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	4
IB16	MYP - Unraveling-theigma-mastering-complex-questions-and-perplexity-for-creative-problem-solving	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	4
IB17	MYP - The-moral-machine	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	4
IB18	MYP - innovation-unleashed-showcasing-new-inventions-with-gamma-app-	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	4

Key	Source	Tier	Contribution	Applicability caution	pp.
IB19	MYP - Mastering-the-art-of-refining-ai-responses	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	6
IB20	MYP - Empowering-minds-exploring-ai-development-with-teachable-machine	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	3
IB21	MYP - Using-ai-chatbots-to-reflect-on-the-impacts-of-science	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	4
IB22	PYP -Unpacking-lines-of-inquiry	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	6
IB23	PYP - Line of Inquiry Sub Statements	3 — Supporting	Companion classroom material for a PYP activity.	Useful only with its parent activity; not standalone evidence.	DOCX
IB24	PYP -LOI Sub Statements on Hexagons	3 — Supporting	Companion classroom material for a PYP activity.	Useful only with its parent activity; not standalone evidence.	2
IB25	PYP - Expository-texts-reading-and-writing-using-ai	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	5
IB26	PYP - Paragraph-scramble	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	3
IB27	PYP - Scavenger-hunt-ignition-activity	2 — Applied	IB classroom activity illustrating a possible teacher or student use of AI.	Exemplar rather than endorsement; verify age, privacy, accessibility and learning value.	4
IB28	PYP - Scavenger Hunt Activity CLUES	3 — Supporting	Companion classroom material for a PYP activity.	Useful only with its parent activity; not standalone evidence.	DOCX
IB29	PYP - Scavenger Hunt Answers	3 — Supporting	Companion classroom material for a PYP activity.	Useful only with its parent activity; not standalone evidence.	DOCX
IB30	PYP - Scavenger Hunt Answers No details	3 — Supporting	Companion classroom material for a PYP activity.	Useful only with its parent activity; not standalone evidence.	DOCX
IB31	PYP - Scavenger Hunt Rhyming Clues	3 — Supporting	Companion classroom material for a PYP activity.	Useful only with its parent activity; not standalone evidence.	DOCX
IB32	INTRODUCTION	2 — Applied	Orientation to the IB AI in the Classroom activity collection.	Activity guidance, not policy; tools and data practices require current review.	DOCX
IB33	IB AI Usage Guide	1 — Core	Current IB position or operational academic-integrity guidance.	Apply directly where it governs IB work; CAS still needs local procedures.	2
IB34	Shaping our approach to AI as a community - International Baccalaureate®	1 — Core	Draft IB AI design principles and emerging red-line direction.	Draft consultation material; do not present as final IB policy.	6

Key	Source	Tier	Contribution	Applicability caution	pp.
IB35	Statement from the IB about ChatGPT and artificial intelligence in assessment and education - International Baccalaureate®	1 — Core	Current IB position or operational academic-integrity guidance.	Apply directly where it governs IB work; CAS still needs local procedures.	2
IB36	The changing face of teaching & learning The impact of technologyhanced learning on pedagogy, assessment and learning environments in IB Diploma Programme schools - International Baccalaureate®	3 — Context	Earlier context on technology-enhanced pedagogy, assessment and learning environments.	Predates current generative AI conditions; indirect relevance.	2
IB37	Using AI in the learning and teaching of DP classical languages	2 — Applied	Subject-specific examples and cautions for DP classical languages.	DP subject guidance; adapt rather than generalize unchanged across divisions.	6
IB38	Using AI in the learning and teaching of DP dance	2 — Applied	Subject-specific examples and cautions for DP dance.	DP subject guidance; adapt rather than generalize unchanged across divisions.	8
IB39	Using AI in the learning and teaching of DP film	2 — Applied	Subject-specific examples and cautions for DP film.	DP subject guidance; adapt rather than generalize unchanged across divisions.	5
IB40	Using AI in the learning and teaching of DP language A language and literature	2 — Applied	Subject-specific examples and cautions for DP language A language and literature.	DP subject guidance; adapt rather than generalize unchanged across divisions.	7
IB41	Using AI in the learning and teaching of DP language A literature	2 — Applied	Subject-specific examples and cautions for DP language A literature.	DP subject guidance; adapt rather than generalize unchanged across divisions.	6
IB42	Using AI in the learning and teaching of DP mathematics analysis and approaches	2 — Applied	Subject-specific examples and cautions for DP mathematics analysis and approaches.	DP subject guidance; adapt rather than generalize unchanged across divisions.	6
IB43	Using AI in the learning and teaching of DP philosophy	2 — Applied	Subject-specific examples and cautions for DP philosophy.	DP subject guidance; adapt rather than generalize unchanged across divisions.	7
IB44	Using AI in the learning and teaching of DP theory of knowledge	2 — Applied	Subject-specific examples and cautions for DP theory of knowledge.	DP subject guidance; adapt rather than generalize unchanged across divisions.	8

B2. UNESCO sources

Key	Source	Tier	Contribution	Applicability caution	pp.
UN01	AI and education guidance for policy-makers	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	50
UN02	AI and education protecting the rights of learners	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	100
UN03	AI and the future of education disruptions, dilemmas and directions	3 — Context	Futures-oriented framing of education, technology and the public good.	Strategic or conceptual; not an operational school framework.	165
UN04	AI competency framework for students	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	80

Key	Source	Tier	Contribution	Applicability caution	pp.
UN05	AI competency framework for teachers	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	52
UN06	AI in the United Arab Emirates' computing, creative design and innovation K-12 curriculum a case study	2 — Applied	Regional K–12 curriculum case study showing progression, integration and enabling conditions.	National case; transfer principles, not the curriculum wholesale.	71
UN07	Artificial intelligence and inclusion, compendium of promising initiatives Mobile Learning Week 2019	2 — Applied	Case compendium on inclusion, accessibility and equitable participation.	Examples vary in context and evidence strength; some predate generative AI.	62
UN08	Artificial intelligence and inclusion, compendium of promising initiatives Mobile Learning Week 2020	2 — Applied	Case compendium on inclusion, accessibility and equitable participation.	Examples vary in context and evidence strength; some predate generative AI.	48
UN09	Artificial intelligence and the Futures of Learning UNESCO	3 — Context	Futures-oriented framing of education, technology and the public good.	Strategic or conceptual; not an operational school framework.	9
UN10	Artificial intelligence for Sustainable Development synthesis report, Mobile Learning Week 2019	3 — Context	Multi-stakeholder synthesis of emerging priorities, opportunities and concerns.	Conference synthesis; broad and partly aspirational rather than implementation standard.	51
UN11	Beijing Consensus on Artificial Intelligence and Education	2 — Applied	International consensus framing human-centered, equitable AI in education.	High-level commitments need translation into school procedures.	70
UN12	Futures of Education UNESCO	3 — Context	Futures-oriented framing of education, technology and the public good.	Strategic or conceptual; not an operational school framework.	6
UN13	Generative AI and the future of education	3 — Context	Futures-oriented framing of education, technology and the public good.	Strategic or conceptual; not an operational school framework.	8
UN14	Guidance for generative AI in education and research UNESCO	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	48
UN15	Guidelines for ICT in education policies and masterplans	3 — Context	General policy-design guidance relevant to infrastructure, openness, governance and sustainability.	Broader than AI; use selectively.	184
UN16	Guidelines on the development of open educational resources policies	3 — Context	General policy-design guidance relevant to infrastructure, openness, governance and sustainability.	Broader than AI; use selectively.	107
UN17	International conference on Artificial intelligence and Education, Planning education in the AI Era Lead the leap final report	3 — Context	Multi-stakeholder synthesis of emerging priorities, opportunities and concerns.	Conference synthesis; broad and partly aspirational rather than implementation standard.	46
UN18	International Forum on AI and Education Ensuring AI as a Common Good to Transform Education, 7-8 December 2021; synthesis report	3 — Context	Multi-stakeholder synthesis of emerging priorities, opportunities and concerns.	Conference synthesis; broad and partly aspirational rather than implementation standard.	52
UN19	International forum on AI and education steering AI to empower teachers and transform teaching, 5-6 December 2022; analytical report	3 — Context	Multi-stakeholder synthesis of emerging priorities, opportunities and concerns.	Conference synthesis; broad and partly aspirational rather than implementation standard.	47
UN20	International Forum on AI and the Futures of Education, developing competencies for the AI Era, 7-8 December 2020 synthesis report	3 — Context	Multi-stakeholder synthesis of emerging priorities, opportunities and concerns.	Conference synthesis; broad and partly aspirational rather than implementation standard.	39

Key	Source	Tier	Contribution	Applicability caution	pp.
UN21	K-12 AI curricula a mapping of governmentdorsed AI curricula	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	63
UN22	Recommendation on the Ethics of Artificial Intelligence	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	28
UN23	Reimagining our futures together a new social contract for education	3 — Context	Futures-oriented framing of education, technology and the public good.	Strategic or conceptual; not an operational school framework.	186
UN24	The algorithm in the room seeing and confronting the implications of AI for the future of education	1 — Core	Direct framework or policy guidance for competencies, rights, ethics, governance or implementation.	Global guidance requiring localization to CAS, Moroccan law and school context.	163
UN25	UNESCO Strategy on technological innovation in education (2021-2025)	3 — Context	UNESCO strategic priorities for inclusive, evidence-informed technology use.	Organizational strategy period ends in 2025; contextual rather than current school guidance.	6
UN26	What you need to know about UNESCO's new AI competency frameworks for students and teachers	3 — Context	Background context for human-centered AI and education policy.	Use as supporting context rather than a direct CAS rule.	12

Source status note. The supplied UNESCO Recommendation on the Ethics of Artificial Intelligence file is a 2021 draft-stage document in this local corpus; later UNESCO publications refer to the adopted recommendation. Before CAS cites it formally in policy, the final adopted text should be used. The 2026 IB AI Design Principles source explicitly presents draft principles for community consultation.

Appendix C. Working glossary

Term	Working meaning in this document
AI system	A machine-based system that produces predictions, content, recommendations or decisions from inputs for explicit or implicit objectives.
Generative AI	AI that generates new content such as text, images, audio, video or code in response to inputs.
Human agency	The capacity and practical freedom to understand, choose, question, direct or refuse AI use and remain responsible for action.
Human oversight / determination	Meaningful human authority to examine, change or reject an AI-supported result, especially in high-impact contexts.
AI literacy	Competencies for understanding, critically evaluating, responsibly using and, where appropriate, designing AI systems within ethical and disciplinary contexts.
Assessed construct	The knowledge, skill or capability an assessment is intended to measure.
Authentication	Evidence-based confidence that submitted work and decisions are attributable to the student under the stated conditions.
Disclosure	A proportionate explanation of whether and how AI contributed to a process or product.
Data minimization	Using or sharing only the least amount of data necessary for an approved purpose.
High-impact decision	A decision materially affecting grades, progression, access, discipline, safeguarding, well-being, employment or comparable rights and opportunities.
Pilot	A bounded, time-limited implementation designed to answer a question before broader adoption.
Independent use	A learner's direct interaction with a service without real-time adult mediation; it does not mean absence of school rules or accountability.

Term	Working meaning in this document
School-managed service	A service accessed through institutionally governed accounts and contractual/technical controls.
AI lifecycle	Selection or design, procurement, configuration, use, monitoring, incident response, review, renewal and retirement.
Protected human work	The thinking, practice, relationship, expression or judgment that a learning or professional task is intended to develop or preserve.

Closing note

This synthesis should be treated as a living foundation. The research establishes durable principles and frameworks, but tools, contracts, IB communications, law and local practice will change. CAS's strongest long-term safeguard is therefore not a static list of permitted techniques. It is a community capable of asking good questions, preserving human purposes, documenting decisions, learning from evidence and revising its rules with care.