

The Oxford Rubric™: Assessing Artificial Intelligence Usage in Schools

Harwar, M.D., June 2026

Abstract

The rapid emergence of artificial intelligence (AI) tools in educational contexts presents both significant opportunities and material risks for schools. While AI has the potential to enhance teaching, personalise learning and reduce administrative burden, its uncritical or poorly governed adoption may undermine safeguarding, pedagogical integrity, professional judgement, teacher/pupil agency, and budgets.

This paper presents the **Oxford Rubric™** as a principled, school-ready framework for assessing whether the use of AI in school settings is appropriate, ethical, and educationally sound. The framework comprises six criteria:

- **Safety**
- **Efficacy**
- **Accountability**
- **Transparency**
- **Agency**
- **Cost Management**

Together, these criteria provide a coherent lens through which school leaders, teachers, governors, regulators, and policymakers can evaluate AI tools, practices, and policies in a way that is aligned with child-centred education, professional standards, and long-term learning outcomes.

Contents

Introduction.....	1
Purpose and Scope of The Oxford Rubric™	1
The Oxford Rubric™ Framework.....	2
Safety	3
Efficacy (Teaching and Learning).....	5
Accountability	6
Transparency	7
Agency (Teachers and Pupils)	8
Cost Management (Affordability and Sustainability)	9
Discussion and Implications	12
Applying the Oxford Rubric™ in Practice	12
Conclusion	12

Introduction

Artificial intelligence (AI) has moved rapidly from the margins of educational technology into mainstream classroom and school operations. Generative AI systems, adaptive learning platforms, automated assessment tools, and administrative agents are increasingly accessible to teachers and pupils alike. In many cases, adoption has preceded the development of shared standards or robust evaluative frameworks, leaving schools to navigate complex ethical, pedagogical, and safeguarding questions independently.

Existing discussions of AI in education often focus on technical capability or efficiency gains while under-examining the deeper educational implications (such as how AI reshapes the roles of teachers, alters pupil learning behaviours, redistributes responsibility, and influences trust within the school system). School-level education in particular demands a higher threshold of care than adult or corporate settings given children’s developmental vulnerability, the statutory safeguarding duties of schools, and the formative nature of schooling.

The Oxford Rubric™ is proposed as a normative framework rather than a technical standard. Its purpose is not to rank AI systems, but to assess **appropriateness of use** in educational contexts. The framework is intentionally technology-agnostic and future-proofed, focusing on enduring educational values rather than transient tools.

Purpose and Scope of The Oxford Rubric™

The Oxford Rubric™ is designed to:

- Support schools in making defensible, principled decisions about AI adoption and use in all aspects of school-level education.
- Provide a shared language for educators, school leaders, governors, parents, pupils, and regulators.
- Shift discussion from “***Can we use AI?***” to “***Should we use AI, in this way, for this purpose?***”
- Complement (rather than replace) statutory guidance, safeguarding policy, and professional judgement.

The framework applies across all aspects of school-level education, including:

- Classroom teaching and learning
- Assessment and feedback
- Behaviour and pastoral systems
- Special educational needs and inclusion
- School administration
- Leadership decision-making.

The Oxford Rubric™ Framework

Overview

The Oxford Rubric™ consists of the following six interdependent criteria:

- Safety
- Efficacy
- Accountability
- Transparency
- Agency
- Cost Management

AI-enabled practices in a school setting should be considered appropriate **only if it satisfies all six criteria** for the given scenario under evaluation. Failure in any single criterion indicates a need for redesign, restriction, or rejection of use. The following sections provide details of the six criteria.

Safety

Definition

Safety refers to the extent to which pupils are protected from physical, psychological, emotional and informational harm when AI is implemented in any school scenario. It also includes how safeguarding is upheld and data protection obligations fulfilled.

Rationale

Schools have a non-delegable duty of care. AI systems must not introduce new safeguarding risks, expose pupils to harmful content, or compromise personal data. Importantly, safety includes **foreseeable misuse** and unintended consequences (not just intended functionality).

Key Assessment Questions

The following is a non-exhaustive list of broad questions that apply to this criterion:

- Does the AI system comply with child-specific data protection and privacy standards?
- Are content filters, moderation, and guardrails appropriate for the intended age group?
- Could reliance on the system create emotional harm, dependency, or exclusion?
- Are there clear protocols that can be activated in the event of failure, error, or misuse?

Indicators

The following indicators operationalise the Oxford Rubric™ for practical evaluation. They are intended to support procurement decisions, lesson design, policy review, and inspection readiness.

Indicators of Appropriate Use

- A documented child-centred safeguarding and data protection impact assessment has been completed.
- Age-appropriate content filtering, moderation, and guardrails are enabled by default.
- Human supervision is maintained for all high-risk or sensitive interactions.
- Data collection is minimal, proportionate, and clearly justified educationally.
- Clear procedures exist for reporting, responding to, and learning from incidents.

Indicators of Inappropriate Use

- Pupils interact with AI systems without adequate content moderation or supervision.
- Personal or sensitive pupil data is processed unnecessarily or opaquely.
- Safeguarding responsibilities are implicitly (or explicitly) delegated to technology.
- Failure modes, errors, or misuse scenarios have not been adequately considered.
- The system exposes pupils to harmful, biased, or developmentally inappropriate outputs.

Efficacy (Teaching and Learning)

Definition

Efficacy concerns whether AI use demonstrably enhances teaching quality and learning outcomes, rather than merely increasing efficiency or introducing novelty with little benefit to pupils.

Rationale

Not all efficiency gains are educational gains. AI should serve learning, not displace the pedagogical relationships, struggle, reflection and practice through which learning occurs.

Key Assessment Questions

The following is a non-exhaustive list of broad questions that apply to this criterion:

- Does the AI use align with sound pedagogical principles?
- Is there evidence that it improves learning outcomes or teaching effectiveness?
- Does it risk narrowing learning, over-scaffolding, or reducing cognitive effort?

Indicators

The following indicators operationalise the Oxford Rubric™ for practical evaluation. They are intended to support procurement decisions, lesson design, policy review, and inspection readiness.

Indicators of Appropriate Use

- The educational purpose is defined before adoption.
- AI use is explicitly aligned with sound pedagogical principles.
- Evidence is gathered on the impact on learning and teaching quality.
- Impact is evaluated using educational metrics (not merely technical ones).
- AI supports feedback, insight, or differentiation without replacing learning.
- Usage is regularly reviewed against educational outcomes, not novelty.

Indicators of Inappropriate Use

- AI is adopted primarily for efficiency or convenience.
- Learning is over-scaffolded, reducing productive struggle.
- Cognitive demand is reduced rather than enhanced.
- Assessment integrity is undermined.
- No evaluation of learning impact is undertaken.

Accountability

Definition

Accountability refers to the clarity of responsibility for decisions, outcomes and harms arising from AI use.

Rationale

AI systems do not bear moral or legal responsibility, but people and institutions do. In schools, accountability must never be displaced onto technology. Clear lines of responsibility are essential for ethical practice, redress, and continuous improvement.

Key Assessment Questions

The following is a non-exhaustive list of broad questions that apply to this criterion:

- Who is responsible for decisions influenced or supported by AI?
- Can outcomes be challenged, reviewed, and overridden by humans?
- Are governance structures in place for monitoring impact over time?

Indicators

The following indicators operationalise the Oxford Rubric™ for practical evaluation. They are intended to support procurement decisions, lesson design, policy review, and inspection readiness.

Indicators of Appropriate Use

- School leaders maintain oversight and ownership of AI-related decisions.
- Teachers retain final professional judgement over all educational decisions.
- Decisions influenced by AI can be reviewed, challenged, and overturned.
- Governance structures exist for monitoring impact over time.
- There are documented escalation and review processes.
- Clear escalation routes are defined for concerns or complaints.

Indicators of Inappropriate Use

- Responsibility for decisions is implicitly shifted onto the AI system.
- Teachers feel compelled to follow AI recommendations without challenge.
- There is no mechanism for appeal or review.
- Governance relies solely on vendor assurances.
- Accountability becomes diffused or unclear in practice.

Transparency

Definition

Transparency concerns the visibility and explainability of AI use to teachers, pupils, parents, school leaders, governors and inspectors.

Rationale

Educational relationships depend on trust. When AI systems operate opaquely (such as making decisions, generating outputs or influencing outcomes without clear explanation), trust is eroded and informed consent becomes impossible.

Key Assessment Questions

The following is a non-exhaustive list of broad questions that apply to this criterion:

- Is it clear when AI is being used, and for what purpose?
- Can teachers and pupils understand, at an appropriate level, how outputs are generated?
- Are limitations, uncertainties, and known risks openly communicated?

Indicators

The following indicators operationalise the Oxford Rubric™ for practical evaluation. They are intended to support procurement decisions, lesson design, policy review, and inspection readiness.

Indicators of Appropriate Use

- Schools can articulate what the system does and does not do.
- AI use is explicitly disclosed to teachers, pupils, and parents.
- The purpose and scope of AI involvement is clearly articulated.
- Outputs are explainable at developmentally appropriate levels.
- Known limitations, uncertainties, and risks are openly acknowledged.
- Schools can clearly state when human judgement overrides AI outputs.

Indicators of Inappropriate Use

- AI use is hidden, implied, or presented as fully autonomous.
- Teachers or pupils cannot explain how or why outputs are generated.
- Marketing claims from software vendors are accepted without critical scrutiny.
- Limitations or risks are downplayed or undisclosed.
- Parents are unable to understand or question AI use meaningfully.

Agency (Teachers and Pupils)

Definition

Agency refers to the extent to which AI use preserves and strengthens the autonomy, judgement, and developmental agency of both teachers and pupils.

Rationale

Education is fundamentally a human endeavour. AI should augment, not diminish, the professional agency of teachers or the developing independence of learners. Over-automation risks deskilling educators and fostering passive dependency in pupils.

Key Assessment Questions

The following is a non-exhaustive list of broad questions that apply to this criterion:

- Do teachers retain control over how AI is used in their practice?
- Are pupils encouraged to think, decide and create independently?
- Does the system support metacognition rather than replace it?

Indicators

The following indicators operationalise the Oxford Rubric™ for practical evaluation. They are intended to support procurement decisions, lesson design, policy review, and inspection readiness.

Indicators of Appropriate Use

- Teachers can adapt, refuse, or modify AI-supported processes.
- Pupils understand AI as a tool, not an authority.
- Pupils are taught to question and critically engage with outputs.
- Use of AI is framed explicitly as support, not substitution.
- Human judgement is augmented, not replaced.
- AI use supports independence and metacognition.

Indicators of Inappropriate Use

- Teachers are deskilled or over-reliant on automation.
- Pupils defer thinking or decision-making to AI.
- AI outputs are treated as authoritative or unquestionable.
- Creativity, voice, or autonomy is constrained.
- Dependency on the system increases over time.

Cost Management (Affordability and Sustainability)

Definition

Cost concerns whether the full financial commitment required to adopt, operate, and sustain an AI system is understood, predictable, proportionate to the educational benefit, and affordable within the school's budget over time. It extends beyond the headline price to the total cost of ownership (TCO) and to the financial risks created by changing pricing models, vendor incentives, or the possible withdrawal of a tool on which the school has come to depend.

Rationale

Schools operate under fixed, publicly accountable budgets that are typically set annually and in advance. This makes them especially exposed to AI pricing that is volatile or difficult to forecast. The sector is moving away from flat-rate, "all-you-can-eat" subscriptions towards usage-based ("token") metering, under which costs rise with use and can spike unpredictably. This penalises the sustained, embedded use that effective adoption produces.

Cost is therefore not a narrow procurement detail but an educational and ethical question. Money committed to AI is money not available for teachers, learning resources, or evidence-based interventions. Every commitment carries an opportunity cost.

Affordability also has an impact on equity. Cost pressures can widen gaps between better- and under-resourced schools and can create inequity within a school where only some pupils are able to access paid tiers.

School leaders are particularly vulnerable to bait-and-switch dynamics. Free pilots, introductory offers, and generous educational discounts can establish dependency before prices rise, before features are moved behind premium tiers, or before metering is introduced. Headline pricing and vendor marketing routinely understate the true, long-run cost.

A tool that is affordable today may be repriced, deprecated, or withdrawn tomorrow for commercial, regulatory, or geopolitical reasons. Capable models have been launched amid considerable marketing fanfare and then rapidly restricted or made unavailable. For example, access to Claude Fable 5 was suspended following a government export-control directive. Where teaching and learning have been built around such tools, its loss is disruptive and costly. Continuity, contractual protection, and a viable exit plan are therefore integral to assessing cost.

Key Assessment Questions

The following is a non-exhaustive list of broad questions that apply to this criterion:

- Is the full total cost of ownership understood, including indirect, recurring, and hidden costs (not only the headline subscription or licence fee)?
- Is the pricing model predictable and affordable across the school's budget cycle under realistic, rather than promotional, usage?
- What is the school's exposure to future price increases, the removal of features, or a shift to usage-based metering?
- Would the repricing or withdrawal of the tool disrupt teaching and learning, and is there a viable contingency or exit plan?
- Does the expenditure represent demonstrable value for money relative to the educational benefit and to lower-cost or non-AI alternatives?

Indicators

The following indicators operationalise the Oxford Rubric™ for practical evaluation. They are intended to support procurement decisions, lesson design, policy review, and inspection readiness.

Indicators of Appropriate Use

- The full total cost of ownership has been calculated before adoption, including licensing, infrastructure, devices, connectivity, staff training, integration, ongoing IT support, data-protection compliance, and eventual decommissioning.
- Pricing is predictable, and the commitment is affordable across the full budget cycle under realistic usage (not only during an introductory or discounted period).
- Contractual protections are in place against sudden price rises, feature removal, or changes to metering (for example, price-lock terms, clear notice periods, and usage caps).
- The school retains a viable exit strategy, including data portability and a documented contingency plan should the tool be withdrawn, deprecated, or made unavailable.
- Cost is weighed explicitly against educational benefit, and against lower-cost or non-AI alternatives, before money is committed.
- Total costs are independently verified rather than taken from headline pricing or vendor marketing.
- Expenditure is reviewed regularly against actual usage, outcomes, and continued value for money.

Indicators of Inappropriate Use

- Adoption is driven by a free pilot, introductory offer, or time-limited educational discount, without modelling the cost once full or metered pricing applies.
- Only the headline price is considered. Indirect, recurring, and hidden costs are overlooked.
- The school becomes dependent on a tool whose pricing is volatile, usage-metered, and difficult to forecast or cap.
- Significant vendor lock-in exists through proprietary data formats, integrations, or workflows, raising the cost and difficulty of switching.
- No contingency exists for the tool being repriced, deprecated, or withdrawn (whether for commercial, regulatory, or geopolitical reasons) as illustrated recently by capable models being heavily marketed and then restricted or made unavailable.
- Spending on AI displaces investment in staffing, resources, or interventions of greater or better-evidenced educational value.
- Cost forecasts rely on vendor assurances or marketing claims accepted without independent scrutiny.
- Cost pressures create or widen inequities between schools, or within a school where only some pupils can access paid tiers.

Discussion and Implications

Adopting the Oxford Rubric™ encourages schools to move beyond reactive or tool-led adoption toward principled, educationally grounded governance of AI. It aligns technological innovation with safeguarding, professional ethics, and long-term educational aims. The framework also provides a foundation for dialogue between schools, regulators, and technology providers. It makes explicit the conditions under which AI can be considered educationally legitimate.

Applying the Oxford Rubric™ in Practice

The Oxford Rubric™ can be operationalised through an evaluation framework that includes:

- School leader operational checklists
- AI impact assessment templates
- Procurement evaluation checklists
- Staff training and professional development
- Pupil and parent communication materials
- Governor/trustee communication materials and assurance packs
- Ofsted-aligned policy packs
- A maturity model that is aligned with the Oxford Rubric™
- AI decision trees and use case approval frameworks
- Ongoing review cycles aligned with school improvement planning.

Crucially, the framework supports **context-sensitive judgement** rather than binary approval. A use case may be appropriate for one age group, subject, or purpose, but inappropriate for another.

The Oxford Rubric™ is intended as a living framework, capable of refinement as technologies, evidence, and educational understanding evolve. We are actively working with trusted partners to fully operationalise and future-proof the Oxford Rubric™, including creating tools and resources for each item in the above list. Appendix A includes the rubric that can be used as a starting point.

Conclusion

AI will continue to shape the educational landscape. The central question for school education is not whether AI will be used, but how (and under what conditions) it should be allowed to shape teaching and learning. The Oxford Rubric™ offers a clear, values-driven framework to support that judgement. By insisting on Safety, Efficacy, Accountability, Transparency, Agency, and Cost Management schools can ensure that AI remains a servant of education rather than its master.

Appendix A. Simple Oxford Rubric

The following rubric enables rapid, consistent evaluation of any AI tool or practice. A use case should be considered acceptable **only when all criteria are Green**.

Criterion	Green – Appropriate	Amber – Caution	Red – Inappropriate
Safety	Safeguarding and data impacts assessed. Strong protections in place.	Partial safeguards. Risks mitigated but not eliminated.	Safeguarding risks unmanaged or unclear.
Efficacy	Clear evidence of improved teaching/learning.	Claimed benefit but limited evidence.	No educational benefit or negative impact.
Accountability	Human judgement clearly retained and documented.	Accountability shared but blurred.	Decisions delegated to AI.
Transparency	AI use is explicit, explainable, and understood.	Some disclosure but limited explainability.	AI use hidden or poorly understood.
Agency	Teachers and pupils remain autonomous.	Some risk of dependency.	Agency diminished or replaced.
Cost Management	Full total cost of ownership understood, predictable, and affordable across the budget cycle, with strong contractual and continuity protections.	Some cost uncertainty or reliance on promotional pricing; partial protection against escalation or withdrawal.	True cost unclear or unaffordable; exposed to metering spikes, bait-and-switch, or service withdrawal with no exit plan.