

TEACHERMATIC



PEDAGOGY-FIRST AI IN EDUCATION

SAFE

Framework

A pedagogy-first model for the responsible use of AI in feedback and assessment
teachermatic.com

Published for institutional review. We welcome feedback from colleges, universities and sector partners as SAFE continues to develop.

To discuss SAFE in your institution, contact the TeacherMatic team via teachermatic.com.

VERSION 1 — MAY 2026

Contents

01	Executive summary	2
02	Rationale	3
03	Introduction	4
04	Safeguarding data and privacy	5
05	Augmentation of professional judgement	7
06	Fairness and inclusion	10
07	Ethics and transparency	13
08	Implementing SAFE: institutional governance	16
09	Implementation aids	18
10	Ongoing development	21
11	Frequently asked questions	22
12	Glossary	25
13	Disclaimer	27

01 Executive summary

The TeacherMatic SAFE Framework provides a clear, institution-ready approach to the responsible use of AI, artificial intelligence, in feedback and assessment. It has been developed in response to the rapid emergence of AI-supported feedback across education and the growing need for clear ethical and operational guardrails in this area.

SAFE ensures that AI-supported feedback remains secure, educator-led, fair and transparently governed. It is built for institutions that want to benefit from AI without compromising academic integrity, professional judgement or learner trust. Institutions should ensure that their policies on student AI use and educator-facing AI tools are coherent and communicated together, so that students understand the distinction between declared, educator-governed AI feedback and undeclared AI use in their own submitted work.

SAFE STANDS FOR

S	Safeguarding data and privacy
A	Augmentation of professional judgement
F	Fairness and inclusion
E	Ethics and transparency

Unlike general-purpose AI tools that can generate unstructured or autonomous outputs, TeacherMatic's feedback workflows operate within clear pedagogical and governance boundaries. AI-generated feedback is positioned as draft and advisory, never as an automated academic decision.

SAFE supports a hybrid model of marking where professional expertise remains central and technology reduces administrative burden without compromising academic integrity or learner trust.

02 Rationale

AI-supported feedback is emerging rapidly across the education sector, yet there is still limited clarity around how it should be used responsibly within assessment. Feedback and marking sit at the heart of academic integrity, professional judgement and institutional reputation, so it is understandable that schools, colleges and universities are cautious about introducing AI into such a sensitive area.

For many educators, this remains uncharted territory. Teachers are often navigating AI use independently and may be unsure what is acceptable, where the boundaries lie or whether they may inadvertently be breaching institutional expectations or guidelines. Many institutions also do not have the capacity to develop detailed frameworks quickly and are looking for practical guidance and reassurance.

The case for AI-supported feedback is not only about reducing workload and delays in returning work to students. Prompt, detailed feedback has a direct impact on learning. When feedback is returned weeks after submission, learners have often moved on from the thinking that produced the work, and the opportunity to act on feedback is diminished. AI-supported feedback can close this gap by enabling educators to return higher-quality, criteria-aligned feedback in a shorter timeframe, while retaining final professional judgement.

The SAFE Framework was developed to provide clarity in this space. It offers a structured, principled approach to ensure that AI-supported feedback remains secure, educator-led, fair and transparently governed, and that innovation strengthens rather than compromises learning outcomes, academic integrity and learner trust.

As TeacherMatic introduces its AI-supported feedback generators, including the Advanced Feedback Generator, which enables structured feedback on uploaded student work, it is important that clear ethical guardrails accompany this capability. SAFE provides a proactive response to the questions institutions are already asking about data protection, academic judgement, transparency and intellectual property.

03 Introduction

SAFE helps schools, colleges and universities adopt AI tools in ways that protect academic integrity, safeguard data and preserve professional judgement while enabling innovation and reducing workload. The framework reflects both pedagogical best practice and emerging expectations around responsible AI use in assessment.

Across this document, the term AI-supported feedback is used as the primary descriptor. This signals that the educator remains the author of the feedback decision, with AI operating in a supportive role.

Note on regulatory positioning

TeacherMatic's feedback and assessment tools are designed to operate consistently with the principles of the EU Artificial Intelligence Act (2024) UK Department for Education guidance on AI in education, and equivalent emerging international frameworks. The platform's core function is to support educators in generating draft feedback content and to avoid autonomous, high-stakes academic decisions such as final grading or admissions.

By focusing on educator-led workflows, TeacherMatic supports the kind of transparency, oversight and safety expectations these frameworks set out, namely:

- › Transparency. Users are clearly informed when they are interacting with AI-generated content.
- › Contextual guidance. Educators are given the framework to interpret and refine AI outputs.
- › Usage records. Internal logging supports institutional accountability and traceability.
- › Human oversight. The system architecture prevents autonomous grading, so teachers retain final authority.

Regulatory expectations for AI in assessment are evolving and may classify AI-assisted evaluation as higher risk in some jurisdictions. SAFE responds to this by treating educator review, transparent disclosure and auditable records as non-negotiable, regardless of the tier under which a given system is ultimately classified.

Awarding body guidance on AI use in assessment varies across organisations and continues to evolve. The framework sets out principles and good practice; institutions must always verify the guidance of each relevant awarding body before implementing AI-supported assessment workflows.

PILLAR 04

S

Safeguarding data and privacy**CORE PRINCIPLE**

Student work is sensitive academic and personal data. It must be protected with the highest standards of security and care.

In this framework, the term Safeguarding refers to the protection of student work, identifiers and assessment data within AI-supported workflows. It is distinct from the institution's statutory safeguarding duties for the welfare of children and vulnerable adults, which sit outside the scope of SAFE.

What this means in practice**TEACHERMATIC**

- › Uses instances of AI models that do not train on user-submitted data
- › Operates within secure, controlled environments supported by internationally recognised security frameworks
- › Does not claim ownership of student submissions or generated outputs and does not reuse submitted content beyond the feedback generation task
- › Processes uploaded content solely for the purpose of generating feedback on the submitted work
- › Aligns with UK GDPR, EU GDPR and equivalent international data protection standards
- › Hosts processing on UK-based secure infrastructure with appropriate controls
- › Supports secure bulk upload and inline annotation workflows
- › Provides role-based access controls and appropriate data retention and deletion processes
- › Supports transparency through platform-level logging available to authorised organisational administrators

- › Maintains documented incident response procedures, including timely notification to authorised institutional administrators in the event of a data security incident
- › Maintains a transparent record of subprocessors and underlying AI model providers, with appropriate contractual safeguards in place and notification of material changes
- › Applies retention and deletion policies that continue to apply once a learner has completed their programme or left the institution
- › Encourages institutions to follow good practice when uploading work, such as avoiding unnecessary personal identifiers and adhering to institutional data protection policies

Institutional good practice

INSTITUTIONS EMBEDDING SAFE ARE ENCOURAGED TO

- › Ensure AI-supported feedback is reflected within assessment and data protection policies
- › Confirm retention and deletion procedures are clearly defined and continue to apply when learners complete their programme or leave the institution
- › Ensure incident response procedures cover AI-supported feedback workflows alongside other data systems
- › Provide clarity to staff about appropriate data handling practices
- › Encourage staff to minimise or anonymise unnecessary personal data when uploading work
- › Maintain oversight of how student work is processed within AI-supported workflows, including any subprocessors or AI model providers in the chain
- › Obtain student consent for AI-supported feedback processes; a tick box on enrolment forms or learner agreements offers a practical and consistent mechanism
- › Verify that any AI tools used for feedback processing do not claim ownership of uploaded student work, do not use submitted content to train, update or fine-tune any AI model and do not process the material in any other way that would affect its ownership or privacy
- › Seek contractual guarantees from any AI provider around non-training, IP ownership and deletion terms, independent of vendor choice

PILLAR 05

A

Augmentation

CORE PRINCIPLE

AI should support educators, not replace them.

What this means in practice

TEACHERMATIC

- › Ensures teachers remain in full control of the final output
- › Positions AI-generated feedback as draft, advisory or supportive
- › Requires educator review, editing and approval before feedback is finalised
- › Does not fully automate marking, assign grades or replace the professional judgement of the educator. TeacherMatic generates written, developmental comments on student work – the assignment of marks, grades and pass/fail decisions remains entirely with the educator
- › Provides refine and additional-information features that give educators direct control over tone, emphasis and content
- › Supports verifiable human review through acknowledgement at the point of export, so released feedback always reflects educator oversight
- › Supports colour-coded annotation so that AI-generated, human-added and hybrid comments can be visually distinguished
- › Supports transparency through platform-level logging available to authorised organisational administrators

The 80 / 20 principle in practice

SAFE reflects a hybrid marking model in which AI generates approximately 80 per cent of the first-draft feedback and the educator reviews, refines and evaluates the final 20 per cent. The 20 per cent is where professional judgement is applied: interpreting context, moderating tone, challenging AI-generated phrasing and making any final grading decision. The 80 / 20 split is

illustrative rather than prescriptive. The underlying principle is that AI handles the heavy lifting and the educator retains authorship of the decision.

In practice, a meaningful educator review can take from a couple of minutes for short-form feedback to ten minutes or more for substantial assignments, depending on assessment type and stakes. Institutions should set local minimum review expectations rather than relying on time alone, and combine these with sampling and moderation to confirm review quality.

When a formal rubric is not available

Formal rubrics are not always available. AI-supported feedback can still be generated where the educator provides an assignment brief, learning outcomes, marking criteria, a mark scheme or other context that anchors the feedback. The less formal the criteria, the more important the educator's professional judgement becomes at the review stage. Educators should be confident that they can ground AI-supported feedback in whatever assessment information they have, and refine the output against their own expertise before release.

When AI and educator judgement diverge

Where the educator's professional judgement differs from the AI-generated draft, the educator's decision is final and must be reflected in the feedback returned to the student on tone, emphasis, criteria interpretation or factual content. Educator judgement always takes precedence. Institutions should put a simple protocol in place so educators know how to override AI-generated content with confidence, and so that material disagreements can be noted where useful for moderation, quality assurance or platform improvement.

Alignment with assessment expectations

Across awarding bodies and regulatory guidance, a single principle is consistent: AI must not replace human academic judgement in assessment. AI may support feedback generation, but final academic judgement and marking decisions remain with qualified educators.

Awarding body rules on AI use differ materially between organisations and continue to develop. Institutions should consult the relevant guidance directly for each awarding body they work with, rather than assuming uniform permission or prohibition across the sector. Where an institution is its own awarding body, the institution's academic regulations set the terms.

Institutional good practice

INSTITUTIONS EMBEDDING SAFE ARE ENCOURAGED TO

- › Reinforce that final academic decisions remain with qualified educators
- › Provide guidance on responsible AI-supported marking practices

- › Integrate AI use into moderation and quality assurance processes
- › Avoid policies that promote automated grading without human oversight
- › Put in place verifiable human review mechanisms, for example, an acknowledgement tick box at export, a minimum review standard, or a workflow-level record that an educator has opened and reviewed each piece of feedback before release
 - † *TeacherMatic fulfils this requirement as standard.*
- › Use colour coding or equivalent annotation conventions to distinguish AI-generated comments from human-added and hybrid comments within annotated student work
- › Invest in staff induction and ongoing support so educators engage meaningfully with AI-generated feedback rather than releasing it without substantive review
- › Recognise that staff attitudes to AI-supported assessment vary, and that resistance often decreases when colleagues see practical benefits first-hand in peer observation or shared sessions

IN PRACTICE Large undergraduate cohort, first-year module

An engineering lecturer used AI-supported feedback on a group design assignment with four tasks, including a hand-drawing element. They uploaded the assignment brief and rubric, scaffolded the prompt with specifics such as tense, referencing style and a request for two positive and two constructive comments per task, and flagged a subtle risk that smaller groups should not be penalised on the drawing task. Every feedback sheet was reviewed against their own professional judgement before release. A short statement on each feedback sheet made clear that the lecturer had marked the work and used AI to support feedback generation. No students raised concerns.

PILLAR 06

F

Fairness and inclusion

CORE PRINCIPLE

AI-supported feedback should promote consistency, inclusivity and criteria alignment.

What this means in practice

TEACHERMATIC

- › Aligns feedback with rubrics, learning outcomes and assessment criteria
- › Applies multi-layered content filtering to reduce harmful, biased or inappropriate outputs
- › Applies prompt-level and model-level measures to mitigate known biases carried by underlying large language models (LLMs – AI systems trained on large volumes of text to understand and generate language)
- › Encourages educator review to identify and address any unintended bias in individual outputs
- › Supports consistency across markers, cohorts and departments by applying the same criteria regardless of cohort size or marking sequence
- › Reduces markers of fatigue-related and relational bias that can emerge in extended marking sessions
- › Allows educators to refine the phrasing, tone, emphasis and accessibility of generated feedback so that it reflects institutional, disciplinary and learner-specific needs

SEND (Special Educational Needs and Disabilities) considerations, accessibility of feedback formats and inclusive design sit with the institution and educator within SAFE. The platform supports inclusive practice by enabling educators to adapt generated feedback for tone, language and accessibility before release; it does not itself provide SEND profiling or automatic compatibility with assistive technologies. Educators should be particularly mindful of learners

who are neurodivergent (including those with ASD or ADHD) and those for whom English is an additional language, ensuring feedback tone and framing is appropriate for their needs.

A note on model-level bias

A reasonable concern raised by academics and senior leaders is that LLM, large language models, carry their own biases from the data they were trained on. SAFE does not dismiss this concern. There is well-documented evidence that interaction with a large language model is, in effect, interaction with a system that reflects a particular set of cultural, linguistic and cognitive styles, often skewed towards educated, urban, neurotypical, English-language and Western-centric norms. This can shape the tone, examples and reasoning patterns that surface in AI-generated feedback, in ways that may not always serve every learner equally.

TeacherMatic mitigates this through several mechanisms, including prompt engineering that grounds outputs in the uploaded brief, rubric, learning outcomes or other assessment criteria provided by the educator; content filtering applied on top of model outputs; and the requirement for educator review before any feedback is released. These measures reduce, but do not eliminate, the possibility of biased or culturally narrow phrasing. Institutional processes should treat educator review as the primary safeguard against any residual bias in AI-generated feedback, with particular attention to language, examples and cultural framing where these matter most to the learner cohort.

Equitable access and cohort fairness

AI-supported feedback should support an equitable learner experience across cohorts. Where it is used in some programmes or cohorts and not others, institutions should monitor whether learner experience and outcomes remain comparable, and avoid creating an implicit two-tier feedback culture. Particular care should be taken where AI-supported feedback would be the primary route for richer, more timely feedback, so that learners on programmes without it are not disadvantaged. Accessibility of feedback formats, including for learners using assistive technologies, should be reviewed alongside fairness of content.

Institutional good practice

INSTITUTIONS EMBEDDING SAFE ARE ENCOURAGED TO

- › Combine AI-supported feedback with the rubrics, marking guidance and essential assessment information provided by each relevant awarding body, recognising that specifications differ across organisations
- › Monitor consistency of AI-supported feedback across programmes and departments as part of quality assurance

- › Keep inclusive assessment policies central to how AI is used, including SEND, language, accessibility and cultural context considerations
- › Ensure that feedback formats are compatible with assistive technologies and inclusive of learners using screen readers, alternative input devices and translation tools
- › Consider equitable access across cohorts: where AI-supported feedback is used in some programmes and not others, monitor whether learner experience and outcomes remain comparable
- › Provide staff development on bias awareness, inclusive feedback practice and critical review of AI outputs
- › Engage students proactively about how AI-supported feedback works and explain the fairness benefits, including reduced marker variability and greater consistency of marking
- › Recognise that attitudes vary among student groups, and that resistance is higher in some higher education disciplines and in subjects directly affected by AI in the labour market, and adapt communications accordingly
- › Sample a proportion of AI-supported feedback through moderation to confirm consistency with human professional judgement

IN PRACTICE Teacher education, reflective portfolio

A teacher training lecturer marked thirty-one lesson reflections in a single batch using AI-supported feedback, reviewing each one against professional standards before release. They reported that the feedback recognised reflective models, theory-to-practice links and critical analysis, and produced specific, developmental suggestions. Time saved was not reinvested in additional marking. It was returned to rest and family life. Students described the feedback as more personalised and constructive than feedback they had received in other contexts.

PILLAR 07

E

Ethics and transparency

CORE PRINCIPLE

Students and staff should understand how AI is used in feedback and assessment.

What this means in practice

TEACHERMATIC

- › Encourages clear communication with students about AI-supported feedback
- › Supports institutional transparency statements and policy alignment
- › Avoids hidden or autonomous marking practices
- › Embeds documented risk assessments and governance processes
- › Provides clear mechanisms for monitoring, reporting and escalation
- › Provides platform-level logging available to authorised organisational administrators, so institutions can evidence where AI has been used
- › Supports disclosure to external examiners and awarding bodies that AI-supported feedback tools are used within the institution's assessment workflow
- › Supports institutions in developing honest, student-facing communications that explain how and why AI is used in feedback

Framing AI-supported feedback to students

Evidence from sector pilots suggests that student acceptance is shaped more by how AI-supported feedback is framed than by whether AI is used at all. Framing the use of AI purely as a way to save teacher time tends to be received poorly, particularly in higher education, where students feel they are paying for teacher time. Framing the use of AI around consistency of marking, timeliness of return, the depth and volume of feedback and the reduction of marker fatigue bias tends to be received more positively.

Evidence also suggests that direct conversation from the educator using the tool is the most effective way to address student apprehension. Written statements are necessary but not sufficient. Institutions should support programme leaders and module leaders to have short, honest conversations with cohorts about what AI-supported feedback is, why it is being used and how the educator remains the author of the decision.

Transparency is particularly important where students are expected to declare their own AI use. Institutions should be prepared to explain why AI-supported feedback by educators is treated differently from unattributed AI use by students. The honest answer is that educator use is declared, bounded, reviewed and recorded, and SAFE is the framework that makes that visible.

Transparency with staff

Transparency applies to staff as well as students. Educators should understand when and how AI is used within their assessment and feedback workflows, including in moderation, sampling and quality assurance processes that touch their work. Hidden or undisclosed AI use, even at institutional level, undermines the trust the framework relies on. Institutions should make AI use visible to staff through governance documentation and routine communication.

Learner appeals and complaints

AI-supported feedback should not introduce new categories of academic appeal. Existing institutional appeal and complaint procedures apply equally to feedback generated with AI support, with no special carve-out. Institutions should make this explicit to learners and ensure that any review of disputed feedback considers the educator's decision-making rather than the AI-generated draft alone.

Institutional good practice

INSTITUTIONS EMBEDDING SAFE ARE ENCOURAGED TO

- › Provide student-facing statements explaining how AI-supported feedback is used
- › Ensure AI use aligns with academic integrity regulations
- › Communicate clearly with staff about expectations and boundaries, including where AI is used in moderation and quality assurance workflows
- › Confirm that institutional academic appeal and complaint procedures apply equally to feedback generated with AI support, with no special carve-out
- › Support programme leaders to have direct conversations with cohorts about AI-supported feedback rather than relying on written statements alone
- › Review AI practices regularly as guidance evolves

- › Maintain a documented record of where and how AI has been used within assessment workflows, including which assignments were processed using AI-supported feedback, to support quality assurance, external review and awarding body compliance requirements
- › Exercise significant caution with AI-generated plagiarism or academic integrity detection tools, which currently produce high rates of false positives; professional judgement and contextual evidence should remain central to any academic integrity decision

08 Implementing SAFE: institutional governance

SAFE compliance is an institutional commitment, not an individual one. Evidence from FE and HE sector pilots shows that AI-supported assessment cannot be implemented safely at the level of individual teachers working in isolation. It requires organisational readiness, clear governance and active senior leadership support.

Institutions planning to embed AI-supported feedback are strongly encouraged to establish a cross-functional working group bringing together senior leaders, middle managers, teaching staff, business support staff and student representatives. This group should address four foundational areas before rollout begins: permissions, policies, procedures and people.

Permissions

Students should be informed that AI tools may be used to support feedback generation and should provide consent. For institutions with upcoming intake cycles, the most practical approach is a tick box on enrolment forms or learner agreements. For institutions wishing to begin sooner, a clearly worded Microsoft Form or Google Form can serve as an interim consent mechanism. Transparency with students about the purpose and benefits of AI-supported feedback is not just a regulatory consideration; it is also the most effective way to build acceptance.

Policies

Existing institutional policies may not yet fully address the use of AI in assessment and feedback. Institutions should review whether their current AI, assessment, teaching and learning, academic integrity or other relevant policies adequately cover AI-supported feedback practices, or whether additional guidance is needed.

Any institutional approach should reflect the hybrid marking model, the human-in-the-loop requirement and the institution's position on which assignment types and awarding bodies are in scope. Policies should be reviewed periodically as awarding body guidance evolves and sector practice develops.

Procedures

Staff need practical guidance on how to use AI-supported feedback tools responsibly within their institutional context. Procedures should cover which assignment types are in scope (for example formative versus summative, and by level and awarding body), how to handle student

work containing personal data, the minimum review standard required before feedback is released and how AI use should be recorded for quality assurance purposes.

Where institutions have broad access to AI-supported feedback tools across their teaching staff, ensuring all staff have access to guidance and support is as important as ensuring all staff have access to the tool itself.

People

Technology adoption succeeds or fails on culture. Alongside permissions, policies and procedures, institutions should plan for staff development, peer learning and honest dialogue with students. Practical measures include induction sessions for new users, short peer-observation slots where staff can see the tool in use, communities of practice for shared refinement of prompting and review habits, and regular feedback loops into the institutional working group so practice continues to improve. Building staff AI literacy is not a one-off event – it is an ongoing professional development commitment that evolves alongside the technology.

Staged adoption: formative first

Sector pilot evidence suggests a clear starting point. Formative feedback, including draft review, practice essays and reflective writing, is the lowest-risk entry point for AI-supported feedback and where workload benefits are most readily achieved. Institutions are encouraged to begin with formative use cases, build staff and student confidence, and extend into summative contexts only when governance, moderation and staff capability are demonstrably in place.

09 Implementation aids

The resources below are provided as indicative starting points. Institutions should adapt them to their own context, policies and awarding body requirements.

The educator-led workflow

AI-supported feedback under SAFE follows a consistent five-step workflow. The educator remains the author of the feedback decision at every stage.

01	02	03	04	05
UPLOAD	GENERATE	REVIEW	RELEASE	RECORD
Brief, rubric or assessment criteria uploaded by educator	AI generates draft feedback aligned to criteria	Educator reviews, refines and approves	Feedback released with AI-use declaration	Use logged for quality assurance and audit

SAFE at a glance

PRINCIPLE	WHAT TO DEMONSTRATE
S Safeguarding	Secure processing, no model training on uploaded work, clear retention and deletion, IP retained by institution or learner
A Augmentation	Educator-led workflow, verifiable human review before release, AI positioned as draft and advisory
F Fairness	Rubric-aligned outputs, inclusive practice, mitigation of fatigue, relational and model-level bias
E Ethics	Student-facing statements, staff guidance, documented record of AI use, honest framing

Sample policy clause

A short clause that institutions can adapt and drop into an existing AI, assessment or academic integrity policy.

SAMPLE Policy clause

This institution is authorised to use approved AI-supported feedback tools to assist educators in the generation of draft feedback on student work. All AI-supported feedback is reviewed, edited and approved by a qualified educator before release. Final academic judgement remains with the educator, and AI-supported feedback is not used to determine final grades or other high-stakes academic decisions. Use of AI-supported feedback is governed by the institution's data protection, academic integrity and assessment policies, and is recorded for quality assurance purposes.

Institutional self-assessment checklist

Use the prompts below to gauge readiness. Each item can be rated as in place, in progress or not yet started.

PERMISSIONS

- › Student consent mechanism is agreed and in place

Note: Where the student cohort may include individuals under the age of 18, institutions should apply additional safeguarding steps and seek appropriate parental or guardian consent in line with their data protection obligations.

- › Student-facing statement on AI-supported feedback is drafted

POLICIES

- › AI use is reflected in assessment, academic integrity and data protection policies
- › Position on formative and summative use is clear
- › Awarding bodies in scope have been checked individually

PROCEDURES

- › Minimum educator review standard is defined and communicated
- › Record of AI-supported feedback use is maintained for quality assurance
- › Retention and deletion procedures for uploaded work are documented

PEOPLE

- › Induction and CPD are available to all staff with access to the tool
- › Community of practice, CoP, or peer-learning route is in place
- › Programme leaders are equipped to have direct conversations with cohorts

Sample student-facing statement

Adapt the wording below to your institutional voice, policy and awarding body context.

SAMPLE

As part of your assessment, your teacher may use TeacherMatic, an AI-supported feedback platform, to help generate draft feedback on your work. Your teacher reviews, edits and approves all feedback before it is returned to you. Your work is processed securely, is not used to train any AI model and remains your intellectual property, or that of the institution as set out in your enrolment agreement. If you have any questions about how AI is used in your feedback, please speak to your programme leader.

Sample consent wording**SAMPLE**

I understand that my tutors may use AI-supported tools, such as TeacherMatic, to help generate draft feedback on my submitted work. I understand that my work is processed securely, is not used to train any AI model and that my tutor reviews and approves all feedback before it is released. [tick box] Yes, I consent to this use.

10 Ongoing development

The SAFE Framework will continue to evolve as sector practice develops, informed by institutional feedback, regulatory guidance and ongoing research into responsible AI use in education.

TeacherMatic's ongoing participation in sector pilots, including the JISC AI assessment pilot, is central to this development process. Evidence from these pilots continues to shape how the platform is built and how guidance for institutions is framed. As comparative feedback, BTEC assessment, portfolio-based grading and VLE integration capabilities are developed, the governance and ethical considerations that accompany them will be reflected in future versions of this framework.

Institutions using SAFE-aligned tools are encouraged to share their own experiences and governance approaches with TeacherMatic, contributing to a growing body of shared practice across the sector. The framework is also intended to sit alongside, and remain compatible with, broader sector work from bodies including JISC, QAA, Advance HE and the Russell Group, as well as awarding body and regulator guidance as it continues to develop.

11 Frequently asked questions

Concise responses to the questions institutions most often raise about AI-supported feedback and the SAFE Framework.

WHAT IS AI-SUPPORTED FEEDBACK?

AI-supported feedback is feedback drafted with the help of an AI tool and then reviewed, refined and approved by a qualified educator before it is released to a learner. The educator remains the author of the feedback decision.

WHAT IS THE DIFFERENCE BETWEEN ASSESSMENT, MARKING, FEEDBACK AND GRADING?

These terms are often used interchangeably but mean different things. SAFE positions AI to support the drafting of feedback, with marking judgement and final grading remaining with qualified educators.

Assessment is the broad, ongoing process of evaluating student progress, identifying learning gaps and gauging understanding. It can be diagnostic (where the learner is now), formative (improving learning as it happens) or summative (capturing what has been learned at a defined point).

Formative assessment can also be directed at the teacher's practice, prompting the teacher to adapt their teaching when learning gaps emerge.

Marking is the act of checking a piece of work, such as an essay, test or assignment, against criteria to identify strengths and gaps, and also in the case of summative assessment, is the act of grading. Feedback is the actionable advice returned to a student to help them improve, drawing on the marking and the wider assessment context.

Grading is the assignment of a final mark or grade, such as a 1 to 9 score, an A to E band or a pass/fail outcome, to capture performance at the end of a unit, module or course. In everyday teaching, these activities often happen together, but they are conceptually distinct, which matters when deciding where AI assistance is appropriate. Grading is also the area where awarding body rules and AI-use restrictions are most likely to come into tension, which is why SAFE places it firmly outside the scope of automated assistance.

WHO DECIDES THE FINAL MARK?

The educator. SAFE-aligned tools do not determine grades. Final academic judgement remains with qualified educators, supported by institutional moderation and quality assurance.

WHERE IS STUDENT DATA STORED, AND IS IT USED TO TRAIN AI MODELS?

Within SAFE-aligned tools, student work is processed within a secure environment, is not used to train third-party AI models, and remains the intellectual property of the learner or the institution as set out in their agreement.

WHAT ABOUT AWARDING BODY RULES?

Awarding body rules on AI use in assessment vary across organisations and continue to evolve. Institutions should consult the relevant guidance directly for each awarding body they work with rather than assume uniform permission or prohibition.

WHAT IF A STUDENT DISPUTES AI-SUPPORTED FEEDBACK?

Existing institutional academic appeal and complaint procedures apply, with no special carve-out for AI-supported feedback. Any review of disputed feedback considers the educator's decision-making rather than the AI-generated draft alone.

WHAT IF THE AI GETS IT WRONG?

Educator review exists for exactly this reason. Where AI-generated content is incorrect, biased or misaligned with criteria, the educator overrides it. Institutions should make this expectation explicit in staff guidance.

WHAT IF I DON'T HAVE A RUBRIC?

AI-supported feedback works best when grounded in explicit criteria, but a formal rubric is not always available. Educators can use an assignment brief, learning outcomes, marking criteria, a mark scheme or other context to anchor the feedback. Where criteria are less formal, the educator's professional judgement carries more weight at the review stage, and the educator should expect to refine the AI-generated draft more substantively before release.

DO STUDENTS NEED TO CONSENT?

Students should be informed that AI tools may be used to support feedback generation and should provide consent. A tick box on enrolment forms or learner agreements offers a practical and consistent mechanism. Where a student declines consent, institutions should have a clear process for handling their work outside the AI-supported workflow.

WHAT SHOULD WE TELL STUDENTS?

Tell students clearly and in plain language: that AI is used to help draft feedback, that the educator reviews and approves all feedback, that their work is processed securely and not used to train AI models, and that AI is not used to determine final grades.

HOW DOES THIS FIT WITH ACADEMIC INTEGRITY POLICIES?

AI-supported feedback by educators is declared, bounded, reviewed and recorded. This is different from undeclared AI use by learners. Institutions should explain this distinction in their academic integrity guidance and ensure SAFE practice is reflected.

CAN AI-SUPPORTED FEEDBACK BE USED FOR SUMMATIVE ASSESSMENT?

Sector evidence suggests starting with formative use cases. Summative use is possible where governance, moderation and staff capability are demonstrably in place, and where awarding body and institutional rules allow.

HOW DO WE EVIDENCE GOVERNANCE AND OVERSIGHT?

Through documented policies, recorded consent, platform-level logging, sampled moderation, and a documented record of where AI has been used within assessment workflows. SAFE's self-assessment checklist sets out the full picture.

WHO IN THE INSTITUTION OWNS THIS?

Most successful adopters establish a cross-functional working group bringing together senior leadership, teaching staff, quality, data protection and student representatives. SAFE is an institutional commitment, not an individual one.

WILL AI-SUPPORTED FEEDBACK BE CONSISTENT ACROSS STUDENTS?

TeacherMatic uses consistent underlying models governed by the same educator-defined rubrics and prompts for each assignment. Because all feedback is reviewed and approved by the educator before release, the educator acts as the quality assurance step. Institutions are advised to document which platform and tool version they use as part of their AI governance records, so that any variation over time can be monitored and accounted for.

12 Glossary

Plain-language definitions of the terms used in this framework.

Advanced Feedback Generator TeacherMatic's capability for generating structured feedback on uploaded student work.

AI-supported feedback Feedback drafted with AI assistance and reviewed, refined and approved by an educator before release.

Assessment The broad, ongoing process of evaluating student progress and gauging understanding. May be diagnostic, formative or summative.

BTEC A vocational qualification framework awarded by Pearson, used widely in further education.

DPO Data Protection Officer, the role responsible for overseeing data protection compliance in an organisation.

EU AI Act The European Union's legal framework for the regulation of artificial intelligence, which classifies AI systems by risk tier.

FE / HE Further Education and Higher Education. In the UK, FE typically covers post-16 college provision and HE covers university-level programmes.

Feedback Actionable advice returned to a student to help them improve, drawing on the marking and the wider assessment context.

GDPR General Data Protection Regulation. In the UK this refers to UK GDPR; in the EU, EU GDPR. Both set out rules for the processing of personal data.

Grading The assignment of a final mark or grade to capture performance at a defined point in a programme, such as the end of a unit, module or course.

JISC A UK organisation supporting digital infrastructure and innovation across post-16 and higher education.

Limited Risk A category within the EU AI Act for AI systems that pose limited risk and are subject to transparency obligations rather than the stricter requirements of higher tiers.

Marking Checking a piece of work against criteria to identify strengths and gaps.

PGCE Postgraduate Certificate in Education, an initial teacher training qualification.

Rubric A structured scheme used to mark student work against agreed criteria and standards.

SEND Special Educational Needs and Disabilities, a UK term covering learners with additional support needs.

Subprocessor A third party engaged by a service provider to process data on its behalf, for example an underlying AI model provider or hosting partner.

VLE Virtual Learning Environment, the digital platform an institution uses to deliver and manage learning, for example Moodle, Blackboard or Canvas.

13 Disclaimer

The SAFE Framework is intended as practical guidance to support institutions in adopting AI-supported feedback responsibly. It does not constitute legal advice. Institutions with specific regulatory or compliance questions are encouraged to seek independent professional advice appropriate to their context.

Each institution remains responsible for its own policies, assessment practices and compliance obligations. TeacherMatic makes no representation that following this framework ensures compliance with any specific legal requirement, awarding body rule or regulatory standard. Awarding body guidance on AI use in assessment differs across organisations and continues to develop; institutions should consult the relevant guidance directly and not rely solely on this document for compliance decisions.

This framework reflects the state of sector knowledge and regulatory guidance at the time of publication. It will be reviewed and updated as AI governance in education evolves. Institutions should treat this as a living document and monitor updates from TeacherMatic, their awarding bodies and relevant sector bodies as guidance in this area develops.

THE SAFE FRAMEWORK

Pedagogy-first AI for education.

teachermatic.com