

Privacy-by-Design in AI-Powered Online Tutoring: A Framework for GDPR-Compliant Learning Analytics for Children

პერსონალურ მონაცემთა დაცვის პრინციპების ინტეგრირება ხელოვნურ ინტელექტზე დაფუძნებულ ონლაინ რეპეტიტორობაში: ბავშვთა სწავლის ანალიტიკის GDPR-თან შესაბამისობის უზრუნველყოფის ჩარჩო

 <https://doi.org/10.52340/healthecosoc.2026.10.03.14>

Nino Shengelia^{1,2a}

ნინო შენგელია^{1,2}

¹ Ivane Javakhishvili Tbilisi State University, Tbilisi, Georgia

¹ ივანე ჯავახიშვილის სახელობის თბილისის სახელმწიფო უნივერსიტეტი, თბილისი, საქართველო.

* ninashengelia18@gmail.com

Abstract

Introduction: The rapid growth of online tutoring has created demand for artificial intelligence (AI) tools that monitor learner engagement and support tutors' instructional decisions. Many engagement-analytics systems, however, rely on video surveillance, facial and emotion recognition and continuous behavioral tracking, which raises serious concerns for children's privacy and rights. **Methods.** The study combines doctrinal analysis of the General Data Protection Regulation (GDPR), the EU Artificial Intelligence Act, international standards on children's rights in the digital environment and Georgia's 2023 Law on Personal Data Protection with the Privacy-by-Design approach. From this analysis it derives a conceptual framework for privacy-preserving engagement analytics and applies it to a qualitative single case study of a voice-only AI tutoring assistant that operates without video capture, facial recognition or emotion inference. **Results.** The framework comprises seven design principles: modality minimisation, pedagogical purpose binding, human oversight with non-evaluative outputs, ephemerality, visible presence with child-legible transparency, control by default and an accountable supply chain. The case study shows that most principles can be embedded in system architecture, while residual risks remain in raw audio retention, third-party processing and age-appropriate transparency. **Conclusions.** Educational AI can provide meaningful pedagogical support without intrusive surveillance when privacy constraints are treated as design requirements rather than as compliance afterthoughts. The framework offers practical guidance for educational technology developers, tutoring providers, regulators and policymakers, including those in EU candidate countries aligning national law with the EU data protection acquis.

Keywords: Privacy-by-Design; GDPR; children's data protection; online tutoring

აბსტრაქტი

შესავალი: ონლაინ რეპეტიტორობის სწრაფმა განვითარებამ გაზარდა მოთხოვნა ხელოვნური ინტელექტის (AI) ინსტრუმენტებზე, რომლებიც აკვირდებიან მოსწავლეთა სასწავლო პროცესში ჩართულობას და ეხმარებიან რეპეტიტორებს სასწავლო გადაწყვეტილებების მიღებაში. თუმცა, ჩართულობის ანალიტიკის მრავალი სისტემა ეფუძნება ვიდეომეთვალყურეობას, სახისა და ემოციების ამოცნობას და ქცევის უწყვეტ მონიტორინგს, რაც

Caucasus University; PRINT ISSN 2960-9992; ONLINE ISSN 2960-9984



© 2026 by the authors. Licensee Health Policy, Economics and Sociology. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license (<https://creativecommons.org/licenses/by/4.0/>).

სერიოზულ შემფოთებას იწვევს ბავშვთა პერსონალური მონაცემების დაცვისა და მათი უფლებების უზრუნველყოფის თვალსაზრისით. მეთოდები: კვლევა აერთიანებს მონაცემთა დაცვის ზოგადი რეგულაციის (GDPR), ევროკავშირის ხელოვნური ინტელექტის აქტის (EU AI Act), ციფრულ გარემოში ბავშვთა უფლებების დაცვის საერთაშორისო სტანდარტებისა და საქართველოს 2023 წლის „პერსონალურ მონაცემთა დაცვის შესახებ“ კანონის სამართლებრივ-დოქტრინულ ანალიზს პერსონალურ მონაცემთა დაცვის სისტემის დაპროექტების ეტაპიდანვე უზრუნველყოფის (Privacy-by-Design) მიდგომასთან. აღნიშნული ანალიზის საფუძველზე შემუშავებულია კონცეპტუალური ჩარჩო, რომელიც მიზნად ისახავს სასწავლო პროცესში ჩართულობის ანალიტიკის განხორციელებას პერსონალური მონაცემების დაცვით. ჩარჩო გამოყენებულია ერთი შემთხვევის თვისებრივ კვლევაში, რომელიც განიხილავს მხოლოდ ხმოვან კომუნიკაციაზე დაფუძნებულ ხელოვნური ინტელექტის სასწავლო ასისტენტს. სისტემა ფუნქციონირებს ვიდეოჩაწერის, სახის ამოცნობისა და ემოციების ავტომატური შეფასების გარეშე. შედეგები: შემუშავებული ჩარჩო მოიცავს შვიდ ძირითად პრინციპს: მონაცემთა შეგროვების მოდალობების მინიმიზაცია; მონაცემთა დამუშავების მკაფიოდ განსაზღვრული პედაგოგიური მიზნებით შეზღუდვა; ადამიანის ზედამხედველობა არაშეფასებითი შედეგების გამოყენებით; მონაცემთა დროებითი შენახვა; სისტემის ფუნქციონირების თვალსაჩინოება და ბავშვებისთვის გასაგები გამჭვირვალობა; მონაცემთა დაცვაზე ორიენტირებული ნაგულისხმევი პარამეტრები და მომხმარებლის კონტროლის შესაძლებლობა; აგრეთვე, ანგარიშვალდებული მიწოდების ჯაჭვი. შემთხვევის კვლევამ აჩვენა, რომ აღნიშნული პრინციპების უმეტესობა შესაძლებელია უშუალოდ სისტემის არქიტექტურაში ინტეგრირდეს. ამავდროულად, გარკვეული რისკები კვლავ რჩება დაუმუშავებელი აუდიოჩანაწერების შენახვის, მესამე მხარეების მიერ მონაცემთა დამუშავებისა და ბავშვების ასაკისთვის შესაფერისი, გასაგები ინფორმაციის მიწოდების თვალსაზრისით. დასკვნები: განათლებაში ხელოვნური ინტელექტის გამოყენებას შეუძლია მნიშვნელოვანი პედაგოგიური მხარდაჭერის უზრუნველყოფა ინტრუზიული მეთვალყურეობის გარეშე, თუ პერსონალური მონაცემების დაცვის მოთხოვნები სისტემის შემუშავების საწყის ეტაპზე იქნება გათვალისწინებული და არ შემოიფარგლება მხოლოდ შემდგომი სამართლებრივი შესაბამისობის უზრუნველყოფით. შემუშავებული ჩარჩო პრაქტიკულ სახელმძღვანელო პრინციპებს სთავაზობს საგანმანათლებლო ტექნოლოგიების დეველოპერებს, ონლაინ რეპეტიტორობის მომსახურების მიმწოდებლებს, მარეგულირებელ ორგანოებსა და პოლიტიკის შემუშავებლებს, მათ შორის ევროკავშირის წევრობის კანდიდატ ქვეყნებში, რომლებიც ეროვნული კანონმდებლობის ევროკავშირის პერსონალურ მონაცემთა დაცვის სამართლებრივ სტანდარტებთან ჰარმონიზაციას ახორციელებენ.

საკვანძო სიტყვები: პერსონალური მონაცემების დაცვა; მონაცემთა დაცვის ზოგადი რეგულაცია; ბავშვთა პერსონალური მონაცემების დაცვა; ონლაინ რეპეტიტორობა.

1. Introduction

Online tutoring has moved from the margins of education to its everyday infrastructure. School closures during the COVID-19 pandemic normalised synchronous video lessons, and many families continued to rely on live online tutors after schools reopened. The shift brought clear benefits in access and flexibility, but it also removed much of what teachers rely on in a physical classroom. On a video call, a tutor cannot easily read posture, notice a wandering gaze or sense the quiet withdrawal of a learner behind a muted microphone. This loss of cues has created a market for artificial intelligence (AI) tools that promise to observe engagement on the tutor's behalf and to prompt timely changes in teaching.

Engagement is a multidimensional construct that combines behavioural, emotional and cognitive components (Fredricks et al., 2004). Automated approaches to its measurement seek to infer these states at a fine temporal resolution from machine-readable signals such as facial expressions, eye movements and interaction logs (D'Mello et al., 2017). These techniques are attractive to developers because they appear to turn an elusive pedagogical quality into a stream of numbers. In practice, however, they tend to require continuous camera access, biometric processing and the accumulation of detailed behavioural records about

learners, many of whom are children. Education is thereby drawn further into the broader process of datafication that has reshaped policy and practice in the sector (Williamson, 2017).

The risks of this model are well documented. Critical scholarship warns that facial recognition in schools normalises surveillance and may push schooling in authoritarian and divisive directions (Andrejevic & Selwyn, 2020). The scientific basis for reading emotions from facial movements is itself contested, because the same configuration of facial muscles can express different states in different people and contexts (Barrett et al., 2019). Empirical investigations reinforce these concerns: a global review of educational technology products endorsed by governments during the pandemic found that 145 of the 163 products examined surveilled children, or had the capacity to do so, in ways that put their rights at risk (Human Rights Watch, 2022). Regulators have also intervened. In 2019 the Swedish data protection authority fined a municipality for a school pilot that used facial recognition to register attendance, holding that sensitive biometric data had been processed unlawfully and that consent was not a valid legal basis given the imbalance of power between pupils and the school (European Data Protection Board, 2019).

European law now places explicit obligations on those who design such systems. The General Data Protection Regulation (GDPR) requires controllers to implement data protection by design and by default (European Parliament & Council of the European Union, 2016), and it recognises that children merit specific protection. The Artificial Intelligence Act (AI Act) adds a separate prohibition: from 2 February 2025, Article 5 has prohibited the placing on the market, putting into service or use of AI systems that infer the emotions or intentions of natural persons on the basis of biometric data in workplace and education institutions, except where use is intended for medical or safety reasons (European Parliament & Council of the European Union, 2024). The statutory scope is limited to education institutions; whether, and when, a private online tutoring provider falls within that term requires a context-specific assessment. Nevertheless, the underlying rights and reliability concerns also support excluding emotion inference by design in private tutoring settings. At the international level, the United Nations Committee on the Rights of the Child (2021) has confirmed that children’s rights, including the right to privacy, apply fully in the digital environment. Yet a persistent gap remains between these principles and engineering practice. Ethical frameworks in learning analytics have been criticised for offering little concrete guidance to the teams who actually build analytics infrastructure (Kitto & Knight, 2019). What is missing is a design-level account of how real-time engagement analytics for children can be built so that privacy is a property of the architecture rather than a promise in a policy document.

The question also has a specific regional dimension. Georgia committed to approximating its legislation with the EU *acquis* under the Association Agreement (European Union, 2014) and was granted EU candidate status in December 2023 (European Council, 2023). In June 2023, the Parliament of Georgia adopted a new Law on Personal Data Protection. The law is substantially aligned with GDPR-style data-protection principles and contains dedicated safeguards for minors’ data. Its core provisions entered into force on 1 March 2024, while the provisions on data protection impact assessments and data protection officers became applicable on 1 June 2024 (Parliament of Georgia, 2023). Educational technology companies operating from Georgia and other candidate countries therefore face converging national and European expectations, often while seeking to serve learners in the EU itself.

Against this background, this article addresses three research questions. RQ1: Which legal and ethical requirements govern AI-based engagement analytics in online tutoring that involves children? RQ2: How can these requirements be translated into a practical Privacy-by-Design framework? RQ3: To what extent can an operational system implement that framework, and which residual risks remain? Section 2 describes the materials and methods. Section 3 presents the requirements, the resulting framework and the case study. Section 4 discusses the findings, their legal and policy implications and the limitations of the study, and Section 5 concludes.

2. Materials and Methods

2.1. Research design

The study follows a qualitative, two-stage design. The first stage is doctrinal and conceptual: it identifies binding legal requirements and authoritative guidance and translates them into design principles. This stage is informed by design science research, which treats the construction of purposeful artefacts, including conceptual models and frameworks, as a legitimate mode of inquiry when the artefact is grounded

in existing knowledge and evaluated against its intended use (Hevner et al., 2004). The second stage applies the framework to a single illustrative case study, a method suited to examining a contemporary phenomenon in its real-world context where the boundary between phenomenon and context is not sharply defined (Yin, 2018).

2.2. Materials

The legal corpus comprised the GDPR (European Parliament & Council of the European Union, 2016); the European Data Protection Board (2020) guidelines on data protection by design and by default; the Article 29 Data Protection Working Party (2017) guidelines on data protection impact assessments; the AI Act (European Parliament & Council of the European Union, 2024); General Comment No. 25 of the Committee on the Rights of the Child (2021); the United Kingdom's Age Appropriate Design Code (Information Commissioner's Office, 2020), included as the most developed regulatory translation of children's data rights into design standards; and Georgia's Law on Personal Data Protection (Parliament of Georgia, 2023).

The scholarly corpus covered three bodies of literature: Privacy-by-Design and its legal codification (Bygrave, 2017; Cavoukian, 2009; Hoepman, 2014; Spiekermann, 2012); the ethics of learning analytics (Drachsler & Greller, 2016; Kitto & Knight, 2019; Pardo & Siemens, 2014; Selwyn, 2019; Slade & Prinsloo, 2013); and children's privacy and consent online (Livingstone et al., 2019; van der Hof, 2016). Literature on AI regulation (Veale & Zuiderveen Borgesius, 2021) and on surveillance-based business models (Zuboff, 2019) informed the interpretation of results.

2.3. Framework development

Framework development proceeded in three steps. First, obligations and recommendations relevant to real-time analytics of learner behaviour were extracted from the legal corpus and recorded together with their source provisions. Second, the extracted requirements were coded thematically and mapped onto the eight privacy design strategies proposed by Hoepman (2014): minimise, hide, separate, aggregate, inform, control, enforce and demonstrate. This mapping ensured that each requirement could be expressed as an engineering choice rather than as an abstract value. Third, overlapping codes were consolidated into a small set of design principles, each linked to its legal anchors and to concrete design controls. The ethics literature was used to test whether the principles addressed the concerns most frequently raised about learning analytics, including transparency, student agency, purpose creep and the commercial exploitation of learner data.

2.4. Case study

The case was selected purposively as an information-rich example of a system whose architecture deliberately excludes video capture and facial analysis. The case is AI Whisper, a real-time transcription and engagement coaching assistant for live online sessions, developed by a Tbilisi-based company. The analysis drew on system architecture documentation, product configuration settings, descriptions of data flows and data retention, and internal compliance planning materials. Each principle of the framework was compared with the documented design using pattern matching and classified as implemented, partially implemented or not yet implemented. The analysis is descriptive and design-focused; it does not evaluate the pedagogical effectiveness of the system.

2.5. Positionality and conflict of interest

The author is the founder and director of the company that develops the case system. This position provided detailed access to design documentation but also creates a risk of bias. To mitigate this risk, the analysis relies on documented architectural features rather than on performance claims, reports identified gaps alongside implemented controls and makes no claims about learning outcomes. The case is presented as an illustration of how the framework can be applied, not as evidence that the system is compliant or superior to alternatives. Independent evaluation remains necessary and is identified as a direction for future research.

3. Results

3.1. Legal and ethical requirements (RQ1)

Core data protection principles. Article 5 of the GDPR requires that personal data be processed lawfully, fairly and transparently; collected for specified, explicit and legitimate purposes; limited to what is necessary; kept accurate; retained no longer than necessary; and protected by appropriate security, with the

controller able to demonstrate compliance. Article 25 turns these principles into a design duty. Controllers must implement appropriate technical and organisational measures, both when determining the means of processing and during processing itself, and must ensure that by default only the personal data necessary for each specific purpose are processed. The European Data Protection Board (2020) stresses that these measures must be effective rather than merely formal and must reflect the state of the art, the cost of implementation and the risks to data subjects. Bygrave (2017) observes that Article 25 is drafted in open-textured language, which makes its practical meaning depend heavily on how designers interpret it.

Protections specific to children. Recital 38 of the GDPR states that children merit specific protection because they may be less aware of the risks and of their rights, particularly in marketing and profiling. Article 8 sets conditions for a child’s consent in relation to information society services, with a default age threshold of 16 that Member States may lower to 13. Article 12 requires information to be provided in clear and plain language, in particular where it is addressed to a child, and Article 17 strengthens the right to erasure for data collected in the context of such services. Recital 71 indicates that decisions based solely on automated processing, including profiling, that produce legal or similarly significant effects should not concern a child. Scholarship has cautioned that consent-centred approaches are ill-suited to children, whose capacity and practical freedom to refuse are limited (van der Hof, 2016), and evidence reviews show that children’s understanding of commercial data practices develops gradually and unevenly (Livingstone et al., 2019). Guidance on impact assessments identifies vulnerable data subjects, systematic monitoring and the innovative use of technology as factors that may indicate likely high risk (Article 29 Data Protection Working Party, 2017). A real-time engagement-analytics service for children may combine several of these factors, particularly where it monitors interactions at scale, generates inferences about behaviour, retains recordings or relies on third-party AI services. A DPIA under Article 35 will therefore be highly likely to be required and should be completed before deployment, with the final assessment depending on the system’s actual purposes, data flows, scale and safeguards.

Biometric and emotion-related processing. Facial images and voice data are not automatically biometric data under the GDPR. They become biometric data where they are subject to specific technical processing relating to physical, physiological or behavioural characteristics and that processing allows or confirms the unique identification of a person. When biometric data are processed for the purpose of uniquely identifying a natural person, Article 9 treats them as a special category of personal data. Ordinary audio recording and speech-to-text processing may therefore involve personal data without necessarily constituting biometric-data processing; the classification depends on the functionality and purpose of the processing. The AI Act adds a separate regulatory layer. It defines an emotion-recognition system as an AI system intended to identify or infer emotions or intentions of natural persons on the basis of their biometric data. Article 5(1)(f) prohibits placing on the market, putting into service or using such systems in workplace and education institutions, except where the system is intended to be used for medical or safety reasons; the prohibition has applied since 2 February 2025. Annex III also classifies specified education and vocational-training uses as high-risk where they are used, for example, to determine access or admission, evaluate learning outcomes, determine an individual’s appropriate level of education or materially influence that level. AI systems intended to detect prohibited behaviour during tests in educational institutions are separately identified as high-risk. Whether a tutoring-related tool is within these categories depends on its intended purpose and practical use, not merely its technical capability (European Parliament & Council of the European Union, 2024).

Children’s rights standards. General Comment No. 25 affirms that children’s rights to privacy, protection and participation apply in the digital environment and that the best interests of the child should guide the design and regulation of digital services (Committee on the Rights of the Child, 2021). The Age Appropriate Design Code translates similar commitments into operational standards, among them high-privacy default settings, data minimisation, profiling switched off by default and transparency adapted to the child’s age (Information Commissioner’s Office, 2020).

The Georgian framework. Georgia’s 2023 Law on Personal Data Protection incorporates data-protection principles and institutional mechanisms comparable to those found in the GDPR, while remaining a distinct national regime. It includes obligations concerning impact assessment, data-protection officers, security incidents and safeguards for minors. In particular, the law requires controllers to protect the best interests of minors; where processing relies on consent, a minor under 16 generally requires consent from a parent or other legal representative, and special-category processing involving minors generally requires written parental or legal-representative consent unless legislation provides otherwise. The core provisions took effect on 1 March 2024, with the DPIA and DPO provisions applying from 1 June 2024 (Parliament of Georgia, 2023). For Georgian developers, the requirements identified above are therefore not only relevant for access to EU markets but increasingly reflect domestic obligations as well.

Ethical requirements. The learning analytics literature adds concerns that the law captures only partially. Scholars emphasise transparency, student agency, the recognition that learners are more than their data and the need to avoid using analytics in ways that stigmatise (Slade & Prinsloo, 2013; Pardo & Siemens, 2014). Checklists for trusted learning analytics similarly foreground purpose, consent, anonymisation and the external use of data (Drachsler & Greller, 2016). Critical perspectives warn that analytics can entrench existing inequalities, reduce learning to behaviour management and subordinate education to the interests of the data economy (Selwyn, 2019).

3.2. The Privacy-Preserving Engagement Analytics framework (RQ2)

The consolidation of these requirements produced a framework of seven design principles, referred to here as the Privacy-Preserving Engagement Analytics (PPEA) framework. Its central premise is that the pedagogical purpose of engagement analytics, namely helping a human tutor respond to a learner in time, determines the minimum data needed, and that every design choice beyond that minimum requires specific justification. Table 1 summarises the principles, their legal anchors and representative design controls.

Table 1. *The Privacy-Preserving Engagement Analytics (PPEA) framework*

Principle	Core requirement	Main legal anchors	Representative design controls
P1. Modality minimisation	Use the least intrusive signal sufficient for the pedagogical purpose	GDPR Art. 5(1)(c), 9, 25(2); AI Act Art. 5(1)(f)	No video capture; no facial detection or recognition; no emotion inference from face or voice; no biometric identification
P2. Pedagogical purpose binding	Analytics serve in-session instructional support only	GDPR Art. 5(1)(b); Recital 38; GC 25	Written purpose statement; no advertising or commercial profiling; no reuse of learner data for model training; contractual limits on processors
P3. Human oversight and non-evaluative outputs	Outputs inform an adult educator and never decide about the child	GDPR Art. 5(1)(a), 22 and Recital 71; AI Act Annex III, where the intended use falls within the listed education categories	Prompts addressed to the tutor; indicators describe the interaction, not the learner's traits; no scores used for grading, admission or placement
P4. Ephemerality	Process transiently and retain as little as possible for as short a time as possible	GDPR Art. 5(1)(e), 25(2)	Real-time processing in memory; derived summaries preferred to raw recordings; short fixed retention; automatic deletion
P5. Visible presence and child-legible transparency	Learners know when analytics are active and what they do	GDPR Art. 12–13; ICO Code	Named, visible software participant; notice at session start; age-appropriate explanations; information for parents
P6. Control by default	The system should provide meaningful, age-appropriate agency and should not make participation conditional on unnecessary or disproportionate analytics processing	GDPR Art. 7, 17, 21; GC 25	Clear non-analytics alternative where feasible; host controls; learner and parent/guardian information; accessible rights-exercise and erasure routes; privacy-protective defaults; procedures for objections and withdrawal of consent where consent is relied on
P7. Accountable supply chain	Controller can demonstrate compliance across the whole processing chain	GDPR Art. 5(2), 28, 32, 35, 44–49	Data protection impact assessment; processor agreements; no-training commitments; transfer assessments; security testing

Note. GC 25 = General Comment No. 25 (Committee on the Rights of the Child, 2021); ICO Code = Age Appropriate Design Code (Information Commissioner's Office, 2020). Article 22 is most directly relevant where an AI-supported decision is based solely on automated processing and produces legal or similarly significant effects. In-session prompts presented to, interpreted by and acted on independently by a tutor will ordinarily fall outside that core scenario; the principle is designed to prevent the system's use from escalating into consequential automated decision-making. Source: author's analysis.

The principles operate as a hierarchy. Modality minimisation (P1) and purpose binding (P2) define what may be collected and why. Human oversight (P3) and ephemerality (P4) constrain what the system may

do with the data and for how long. Transparency (P5) and control (P6) secure the position of the learner, while the accountable supply chain (P7) ensures that the guarantees are not undermined by the third-party components on which modern AI systems depend. In terms of Hoepman’s (2014) strategies, P1 and P4 implement minimise and aggregate, P2 and P3 implement separate and enforce, P5 and P6 implement inform and control, and P7 implements demonstrate.

3.3. The case system (RQ3)

AI Whisper is offered to tutoring providers and other organisations that conduct live online sessions. In the tutoring setting, a software agent joins the session as a named participant. It must be admitted by the session host and is then visible to all participants. The agent processes only the audio stream. Speech is transcribed, and a language model evaluates indicators of engagement drawn from the dialogue, such as the balance of tutor and learner talk and the pattern of questions and responses. When the indicators suggest that the learner may be disengaging, the system delivers a short prompt to the tutor in real time, supported by a colour-coded status indicator. After the session, the tutor receives a written report. The system does not capture video frames, does not perform facial detection or recognition and is not designed to infer emotions. Session audio is available to the tutor in the web portal for 30 days and is then automatically deleted. The provider has decided not to offer the service to learners under 13 and positions it primarily for adult learners and, subject to appropriate safeguards, supervised secondary-school learners. This product boundary does not itself determine the applicable legal regime. Where the service is offered to minors, the provider and tutoring organisation must identify the applicable lawful basis, assess the child’s age and jurisdiction, provide age-appropriate information, and implement any required parental or guardian involvement. For consent-based processing in Georgia, the statutory threshold is generally 16 rather than 13. The company has also committed to building full compliance with the GDPR and with comparable US frameworks rather than claiming compliance prematurely.

Table 2. *Application of the PPEA framework to the case system*

Principle	Observed implementation in the case	Assessment
P1. Modality minimisation	Audio only; no video frames, facial detection, recognition or emotion inference. Voice is processed as personal data; biometric-data classification depends on the specific technical processing and purpose. Inference from vocal features must remain excluded.	Implemented
P2. Pedagogical purpose binding	Purpose defined as real-time and post-session support for the educator. Formal published purpose statement and processor-level restrictions on secondary use and model training still to be documented.	Partially implemented
P3. Human oversight and non-evaluative outputs	Prompts go to the tutor, who decides whether and how to act. The system makes no decisions about learners and produces no grades or placements.	Implemented
P4. Ephemerality	Real-time analysis is transient, but raw session audio is retained and downloadable for 30 days before deletion.	Partially implemented
P5. Visible presence and child-legible transparency	The agent is named and visible to all participants. Age-appropriate notices for learners and information for parents are not yet standardised.	Partially implemented
P6. Control by default	The host controls admission. Learner-level opt-out and a dedicated erasure route for minors are not yet formalised.	Partially implemented
P7. Accountable supply chain	Transcription and language-model functions rely on third-party processors. A documented impact assessment, processor terms and transfer assessments are in preparation.	Not yet implemented

Source: author’s analysis of system documentation.

The mapping shows that the most consequential choices, those concerning which signals are collected and who acts on them, were made at the architectural level and are therefore robust (P1 and P3). No configuration change or administrative decision could turn the system into a facial analysis tool without rebuilding it. The remaining principles are partially implemented, and the gaps share a common feature: they

concern documentation, defaults and supply-chain governance rather than core architecture. Three residual risks stand out. First, the 30-day retention of raw audio is longer than the real-time purpose requires, and a derived-data default with shorter retention for sessions involving minors would align better with P4. Second, reliance on external transcription and language-model services creates exposure to secondary use and international transfers unless it is contained by contract and by technical safeguards (P7). Third, a visible software participant is a necessary but not sufficient form of transparency for children, who need explanations adapted to their age (P5).

4. Discussion

4.1. *Surveillance is a design choice*

The central finding of the study is that intrusive surveillance is not a technical precondition for useful engagement analytics. Camera-based systems are often justified by the richness of the data they capture, yet that richness has a doubtful pedagogical return. Inferences about emotion drawn from facial movements rest on contested science (Barrett et al., 2019), and a label such as “confused” or “bored” still leaves the tutor to decide what to do. Dialogue-based indicators are narrower but closer to action: a learner who has not spoken for several minutes can simply be invited to respond. The case therefore suggests a shift in design logic from inferring inner states to supporting instructional moves. This shift reduces the volume and sensitivity of the data processed while keeping the tutor, not the algorithm, responsible for pedagogical judgement.

This finding also has an economic dimension. Surveillance-based models are sustained by the expectation that behavioural data will acquire further value beyond the original service (Zuboff, 2019). Purpose binding (P2) and ephemerality (P4) remove that expectation by design. They impose a cost on the provider, which forgoes a potential data asset, but they also remove a source of legal and reputational risk that has repeatedly affected educational technology (Human Rights Watch, 2022).

4.2. *Legal ambiguities*

The analysis revealed several points on which the law leaves room for interpretation. The first concerns the scope of the AI Act prohibition on emotion recognition. Article 5(1)(f) refers to education institutions, and it is open to debate whether a private online tutoring service falls within that notion. The framework does not depend on the answer: it excludes emotion inference as a matter of principle, because the risks that justify the prohibition are equally present in private tutoring. The second point concerns voice. Ordinary voice recording or transcription is not automatically biometric-data processing under the GDPR; classification depends on whether specific technical processing is used to allow or confirm unique identification. Because the AI Act definition of emotion recognition is tied to biometric data, systems that infer emotional states from vocal features such as pitch or tone may fall within the prohibition where the requisite biometric-data conditions are met, even if they never process an image. Developers of voice-only tools should therefore treat prosodic emotion inference as excluded, as P1 requires.

The third point concerns roles and lawful bases. In a business-to-business tutoring arrangement, the tutoring provider will often determine the purposes of processing and act as controller, while the analytics vendor may process personal data on documented instructions as a processor. The allocation is nevertheless fact-specific: it depends on who determines the essential purposes and means of the processing in practice. A vendor that determines independent purposes—for example, product analytics, model development, benchmarking or retention beyond the customer’s instructions—may be an independent controller or joint controller for those operations. A clear data-processing agreement, transparent allocation of responsibilities and a prohibition or tightly defined limitation on secondary use are therefore essential. Consent may be difficult to rely on in compulsory or dependency-based educational contexts because of the imbalance of power between learners and institutions, as the Swedish school case illustrates (European Data Protection Board, 2019), and because of children’s limited capacity to refuse (van der Hof, 2016); however, no lawful basis should be presumed. The appropriate basis depends on the controller’s role, the relationship with the learner, whether the service is optional, the nature of the processing and applicable national law. Where legitimate interests are considered, the controller must document a genuine necessity and balancing assessment that gives substantial weight to the child’s rights, reasonable expectations and best interests. Where consent is relied upon, it must be demonstrably informed, specific, freely given and withdrawable, with any child-specific requirements satisfied. Finally, while engagement prompts addressed to a tutor are unlikely to fall within the high-risk categories of Annex III, the same data could slide into high-risk territory

if they were used to evaluate learning outcomes or to monitor behaviour during assessments. P3 guards against that drift.

4.3. Trade-offs and critical perspectives

Privacy-preserving design carries real trade-offs. Audio-only indicators will miss forms of disengagement that are silent, and speech recognition can perform unevenly across languages, accents and age groups, which may produce less accurate prompts for some learners. These limitations should be tested and disclosed rather than hidden. The framework also does not escape the critique that learning analytics can reduce education to behaviour management (Selwyn, 2019). Measures such as talk-time are proxies, not learning itself. Directing prompts to an adult professional rather than to the child, keeping outputs non-evaluative and leaving pedagogical decisions to the tutor mitigate this concern but do not remove it. A further risk is that session reports could be repurposed to evaluate tutors, raising workplace monitoring issues that lie outside the scope of this article but deserve attention.

4.4. Implications for Georgia and other candidate countries

For Georgia, the framework offers a practical bridge between domestic law and the EU *acquis*. The 2023 Law on Personal Data Protection includes impact-assessment obligations and dedicated protections for minors (Parliament of Georgia, 2023), and the PPEA principles can structure how local educational technology companies meet those obligations in a way that is also recognisable to EU regulators and customers. The Personal Data Protection Service could support this process through sector-specific guidance on educational technology and through regulatory sandboxes in which companies test privacy-preserving designs under supervision. Such measures would help small developers in candidate countries compete in the EU market on the basis of trustworthiness rather than on data extraction, and would contribute to the practical, not merely formal, approximation of legislation.

4.5. Limitations and future research

The study has several limitations. It relies on a single case, which limits generalisation. The author’s role in the company that develops the case system creates a potential for bias, which was mitigated but not eliminated. The analysis is based on documentation rather than on an independent technical audit, and it does not assess pedagogical effectiveness or the perspectives of children and parents. This assessment evaluates the documented architecture against a conceptual Privacy-Preserving Engagement Analytics framework. It is not a legal opinion, certification, conformity assessment, data protection impact assessment or independent security audit. Whether the system complies with the GDPR, the AI Act, Georgian law or any other applicable regime depends on the specific deployment context, the identity and role of each actor, the learners’ jurisdictions and ages, contractual arrangements, data flows, technical configuration and actual use. In particular, design choices that reduce risk do not remove the need for a documented lawful basis, appropriate transparency measures, processor governance, security controls and, where required, a DPIA. Future research should include comparative multi-case studies, independent audits and published impact assessments, empirical evaluation of whether privacy-preserving signals improve teaching and learning, and participatory research with children, consistent with their right to be heard in matters that affect them (Committee on the Rights of the Child, 2021).

5. Conclusion

This article set out to establish whether AI-based engagement analytics in online tutoring can serve children’s education without subjecting them to intrusive surveillance. The analysis of European, international and Georgian law shows that such systems are subject to demanding requirements concerning purpose limitation, data minimisation, transparency, automated decision-making and, where the AI Act’s conditions are met, the prohibition on emotion recognition in education institutions (RQ1). These requirements can be translated into seven design principles that together form the Privacy-Preserving Engagement Analytics framework (RQ2). The case study demonstrates that the most consequential principles, modality minimisation and human oversight, can be embedded directly in system architecture, while retention, transparency for children and supply-chain accountability require continuing work (RQ3).

Three recommendations follow. Developers should decide which signals to collect on the basis of the concrete instructional action they intend to support, and should exclude video, facial and emotion analysis by design. Tutoring providers should assess whether a DPIA is required, complete one before deployment where the Article 35 threshold is met, and require no-training commitments and short retention periods from their technology suppliers. Regulators and policymakers, including those in EU candidate countries, should issue sector-specific guidance that treats privacy-preserving architecture as the expected standard for educational AI involving children. When privacy is treated as a design requirement rather than an afterthought, educational AI can support teachers while respecting the rights of the learners it is meant to serve.

References

- Andrejevic, M., & Selwyn, N. (2020). Facial recognition technology in schools: Critical questions and concerns. *Learning, Media and Technology*, 45(2), 115–128. <https://doi.org/10.1080/17439884.2020.1686014>
- Article 29 Data Protection Working Party. (2017). *Guidelines on data protection impact assessment (DPIA) and determining whether processing is “likely to result in a high risk” for the purposes of Regulation 2016/679 (WP 248 rev.01)*. European Commission.
- Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest*, 20(1), 1–68. <https://doi.org/10.1177/1529100619832930>
- Bygrave, L. A. (2017). Data protection by design and by default: Deciphering the EU’s legislative requirements. *Oslo Law Review*, 4(2), 105–120. <https://doi.org/10.18261/issn.2387-3299-2017-02-03>
- Cavoukian, A. (2009). *Privacy by design: The 7 foundational principles*. Information and Privacy Commissioner of Ontario. <https://www.ipc.on.ca/wp-content/uploads/resources/7foundationalprinciples.pdf>
- Committee on the Rights of the Child. (2021). *General comment No. 25 (2021) on children’s rights in relation to the digital environment (CRC/C/GC/25)*. United Nations. <https://www.ohchr.org/en/documents/general-comments-and-recommendations/general-comment-no-25-2021-childrens-rights-relation>
- D’Mello, S., Dieterle, E., & Duckworth, A. (2017). Advanced, analytic, automated (AAA) measurement of engagement during learning. *Educational Psychologist*, 52(2), 104–123. <https://doi.org/10.1080/00461520.2017.1281747>
- Drachsler, H., & Greller, W. (2016). Privacy and analytics: It’s a DELICATE issue. A checklist for trusted learning analytics. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge* (pp. 89–98). Association for Computing Machinery. <https://doi.org/10.1145/2883851.2883893>
- European Council. (2023, December 15). *European Council meeting (14 and 15 December 2023): Conclusions*. Council of the European Union. <https://www.consilium.europa.eu/>
- European Data Protection Board. (2019, August 22). *Facial recognition in school renders Sweden’s first GDPR fine* [News release]. https://www.edpb.europa.eu/news/national-news/2019/facial-recognition-school-renders-swedens-first-gdpr-fine_en
- European Data Protection Board. (2020). *Guidelines 4/2019 on Article 25: Data protection by design and by default* (Version 2.0). https://www.edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_201904_dataprotection_by_design_and_by_default_v2.0_en.pdf
- European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). *Official Journal of the European Union*, L 119, 1–88. <http://data.europa.eu/eli/reg/2016/679/oj>
- European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L 2024/1689, 12 July 2024. <http://data.europa.eu/eli/reg/2024/1689/oj>
- European Union. (2014). Association Agreement between the European Union and the European Atomic Energy Community and their Member States, of the one part, and Georgia, of the other part. *Official Journal of the European Union*, L 261, 4–743.

Fredricks, J. A., Blumenfeld, P. C., & Paris, A. H. (2004). School engagement: Potential of the concept, state of the evidence. *Review of Educational Research*, 74(1), 59–109. <https://doi.org/10.3102/00346543074001059>

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>

Hoepman, J.-H. (2014). Privacy design strategies. In N. Cuppens-Boulahia, F. Cuppens, S. Jajodia, A. Abou El Kalam, & T. Sans (Eds.), *ICT systems security and privacy protection* (pp. 446–459). Springer. https://doi.org/10.1007/978-3-642-55415-7_38

Human Rights Watch. (2022). “How dare they peep into my private life?”: Children’s rights violations by governments that endorsed online learning during the Covid-19 pandemic. <https://www.hrw.org/report/2022/05/25/how-dare-they-peep-my-private-life/childrens-rights-violations-governments>

Information Commissioner’s Office. (2020). *Age appropriate design: A code of practice for online services*. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/childrens-information/childrens-code-guidance-and-resources/age-appropriate-design-a-code-of-practice-for-online-services/>

Kitto, K., & Knight, S. (2019). Practical ethics for building learning analytics. *British Journal of Educational Technology*, 50(6), 2855–2870. <https://doi.org/10.1111/bjet.12868>

Livingstone, S., Stoilova, M., & Nandagiri, R. (2019). *Children’s data and privacy online: Growing up in a digital age. An evidence review*. London School of Economics and Political Science. <https://eprints.lse.ac.uk/101283/>

Pardo, A., & Siemens, G. (2014). Ethical and privacy principles for learning analytics. *British Journal of Educational Technology*, 45(3), 438–450. <https://doi.org/10.1111/bjet.12152>

Parliament of Georgia. (2023). Law of Georgia on Personal Data Protection (No. 3144-XIმს-Xმმ), adopted 14 June 2023, consolidated version as amended. Legislative Herald of Georgia. <https://matsne.gov.ge/en/document/view/5827307>

Selwyn, N. (2019). What’s the problem with learning analytics? *Journal of Learning Analytics*, 6(3), 11–19. <https://doi.org/10.18608/jla.2019.63.3>

Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist*, 57(10), 1510–1529. <https://doi.org/10.1177/0002764213479366>

Spiekermann, S. (2012). The challenges of privacy by design. *Communications of the ACM*, 55(7), 38–40. <https://doi.org/10.1145/2209249.2209263>

van der Hof, S. (2016). I agree, or do I? A rights-based analysis of the law on children’s consent in the digital world. *Wisconsin International Law Journal*, 34(2), 409–445.

Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the draft EU Artificial Intelligence Act. *Computer Law Review International*, 22(4), 97–112. <https://doi.org/10.9785/crl-2021-220402>

Williamson, B. (2017). *Big data in education: The digital future of learning, policy and practice*. SAGE.

Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.