Journal of
Risk and Financial Management

MDPI

Article

Triple-Entry Accounting and Other Secure Methods to Preserve User Privacy and Mitigate Financial Risks in AI-Empowered Lifelong Education

Konstantinos Sgantzos [1,*], Panagiotis Tzavaras [2], Mohamed Al Hemairy [3] and Eva R. Porras [4]

[1] Department of Humanities, Social Sciences and Law, School of Applied Mathematical and Physical Sciences, National Technical University of Athens, Zografou Campus, 9, Iroon Polytechniou Str., 15772 Athens, Greece
[2] Department of Management and Marketing, School of Business Administration, European University Cyprus, P.O. Box 22006, 1516 Nicosia, Cyprus; p.tzavaras@external.euc.ac.cy
[3] Research Institute of Science and Engineering [RISE], University of Sharjah, Sharjah P.O. Box 27272, United Arab Emirates; malhemairy@sharjah.ac.ae
[4] Department of Business Economics, Applied Economics II, and Fundamentals of Economic Analysis, Universidad Rey Juan Carlos, P.° de los Artilleros 38, Vicálvaro, 28032 Madrid, Spain; eva.porras@urjc.es
* Correspondence: ksgantzos@mail.ntua.gr; Tel.: +30-693-657-6979

Abstract: Within the past five years, and as Artificial Intelligence (AI) increasingly pervades the academic and educational landscape, a delicate balance has emerged between leveraging AI's transformative potential and safeguarding individual privacy, which needs to be carefully maintained. The preservation of user privacy entails severe financial risks via penalties for the violation of directives such as General Data Protection Regulation (GDPR). This manuscript examines three neoteric approaches to data privacy protection in AI-empowered lifelong education. The first method uses Triple-Entry Accounting (TEA) together with Distributed Ledger Technology (DLT); the second method uses a transaction Merkle tree that can be used as a "proof of existence" so that the users can safeguard their personal information; and the third approach examines the advantages and disadvantages of an offline AI-tutor multimodal model that can operate without internet access. Finally, the ethical implications of deploying such technologies are critically discussed, emphasizing the necessity of achieving privacy while retaining the human factor in education.

Keywords: artificial intelligence; triple-entry accounting; distributed ledger technology; lifelong education; Merkle trees; data analysis; privacy

Academic Editor: Thanasis Stengos

Received: 11 February 2025
Revised: 20 March 2025
Accepted: 21 March 2025
Published: 26 March 2025

Citation: Sgantzos, K., Tzavaras, P., Al Hemairy, M., & Porras, E. R. (2025). Triple-Entry Accounting and Other Secure Methods to Preserve User Privacy and Mitigate Financial Risks in AI-Empowered Lifelong Education. *Journal of Risk and Financial Management*, *18*(4), 176. https://doi.org/10.3390/jrfm18040176

1. Introduction

Lifelong education, as a dynamic and evolving concept, extends beyond formal learning to encompass personal development, professional growth, and social advancement across all stages of life. It aims to enhance the quality of life through structured, semi-structured, and informal learning experiences. Scholars have differentiated lifelong education from lifelong learning, where the former often aligns with institutional frameworks, and the latter reflects a more individualized, socio-personal process (Dave, 1976). The characteristics of the concept of lifelong education are also present in definitions of lifelong learning. For Billett (2010), the latter forms a socio-personal procedure and a personal fact conceptually divergent from lifelong education, which forms a kind of institutional fact.

In the context of lifelong education, AI has emerged as an important factor. AI-based applications have brought in the dynamic and automated nature of their accessibility. They interact with human activity at an unprecedented speed and scale, representing not just

a means but a complementary part of the learning activity (Poquet & De Laat, 2021). By fostering learning groups and communities, advancing accessibility and delivering personalized learning experiences, AI is shaping a promise of a bright educational future ahead of us.

The issue of personalized learning seems to be the hallmark of this dynamic technology. AI analyzes vast amounts of data about a student's performance, preferences, and conduct, shaping dynamic, individualized learning patterns. This kind of tailoring is integral to lifelong education, where the needs, goals and interests of trainees and students evolve over time. Thus, AI-empowered technologies can affect social and even human cognitive processes, going above and beyond the tasks they were designed to perform (De Laat et al., 2020).

During the COVID-19 pandemic it became obvious that education needs to be transformed in a remote but also more personalized manner. Due to social distancing, educators and students were forced to metamorphose the teaching procedure via remote sessions which were able to elegantly compensate, for the most part, the absence of a personal face-to-face experience of the school class (Almpanis & Paul, 2022). Recently, the educational system had another highly transformative element added, and that is none other than the creation and deployment of Large Language Models (LLMs).

Personalized AI models can be constructed to teach specific scientific fields, trained and designed on big data carefully collected and tokenized. These neural networks are usually trained on large datasets of text and code, and they can generate personalized learning experiences for each learning field and student. The procedure includes individualized feedback on the learning process, recommending specialized learning materials, and adapting the pace of instruction to each student's needs depending on the learning feedback (Khan Academy, 2023; Su & Yang, 2023). The unprecedented emergence of different LLMs over this last year has provided a vast opportunity of advancement to the education field, but nevertheless creating also a great deal of ethical dilemmas (Gan et al., 2023).

One of the main drawbacks is privacy concerns of participants in certain personalized procedures, like grading or progress marks. Another is the visibility of sensitive personal data of the participants, such as birthday dates, names, family contact information and others. Personal student details, grading and educational report leaks are common in the education field due to either accidental data sharing or cyber-attacks and can cause shaming and harassment to students. Current centralized systems seem to perform poorly regarding student privacy; there are numerous cases of data breaches, with some of them being high-profile (Nowicki, 2020). Moreover, we must consider that in the AI era, educational data are considered "the new gold" because they can be used for data analysis and machine learning by software companies; as a result, such companies often offer minor updated terms, which compromise data security, to which the user is required to consent to keep using the software (Vigliarolo, 2024). Rarely does the student pay attention to such details, and often they do not know or care about the consequences of sharing their personal information, and the potential data mining they will be subjected to by data brokers.

Under the General Data Protection Regulation (GDPR) directives, educational institutions can face significant financial fines for mishandling student personal data. The regulation establishes a two-tiered system of administrative fines. Severe violations that breach fundamental GDPR principles, data subject rights, or international transfer requirements can incur fines of up to EUR 20 million or 4% of the organization's total worldwide annual revenue, whichever is higher. Less severe infractions, particularly those related to organizational and technical measures, may result in fines of up to EUR 10 million or 2% of total worldwide annual revenue, whichever is higher. These substantial penalties underscore the critical need for robust data protection frameworks in educational settings, where

institutions regularly process sensitive personal information including academic records, health data, and demographic details of their student populations (GDPR, 2025).

We examine three methodologies to mitigate these problems by employing a set of neoteric approaches. The first is Triple-Entry Accounting (Grigg, 2005) to ensure integrity together with Distributed Ledger Technology records (Sgantzos et al., 2023) to ensure transparency and auditability. Since each record forms a transaction which is recorded in multiple copies of a public ledger, TEA a priori ensures that it is impossible for any single participant or third entity to alter records without detection. DLT, due to its decentralized infrastructure, ensures that there are no weak points of failure, making it much harder for hackers to compromise the system. The utilization of cryptographic technology in this system ensures that sensitive data is securely encrypted, protecting students' personal information from unauthorized access. The combination of TEA with DLT provides a system which is highly resistant to cyber-attacks and data breaches. The proposed solution also aligns with GDPR directives on data protection and privacy. It provides data integrity, transparency, and security, so that the institutions involved can better comply with regulatory requirements and avoid large fines associated with data breaches.

The second method is Merklized transactions (J. Davis, 2024) to pseudonymize the sensitive data of the participants and ensure their privacy. The proposed method enables a comprehensive approach to data privacy protection for online LLMs, encircling data minimization, pseudonymization, access control, and auditability. By separating educational data from personal identifiers, the first two methods minimize the potential for identity disclosure. Using pseudonymity, we can further enhance user privacy by replacing personally identifiable information with unique identifications (IDs) (i.e., a cryptographic signature) that cannot be linked back to individuals. Access control mechanisms restrict data access to authorized parties only, while auditability features allow for transparent tracking of data usage and potential breaches.

Lastly, we also examine the usage of an offline model based on Large Language and Vision Assistant (LLaVA) that can operate as an offline AI-educator, and we evaluate the advantages and disadvantages of this particular technology. This method operates without the need for continuous internet connectivity and aims to mitigate the risks of data exposure to external threats. This is particularly important in educational environments where sensitive student data are handled. By processing data locally, offline models minimize the chances of data interception during transmission (man-in-the-middle attacks), which is a common vulnerability in online systems. The isolated nature of the proposed method ensures that the system is GDPR-compliant, while at the same time offers consistent performance and low latency in its usage.

The impetus for authoring this study comes from many case studies that have highlighted the vulnerabilities that centralized educational systems bear in their conceptual design. One of the most important reasons is that high-profile data breaches in educational institutions have underscored the need for more robust privacy solutions. Another point of interest is that the rapid adoption of AI in education during the pandemic has also revealed both the potential and the challenges of personalized learning. There is a gray area around what happens with user prompts, or how they are used by online AI providers. The need for privacy-preserving solutions has become more pressing as AI technologies become more integrated into educational practices. Finally, considering well-documented ethical dilemmas associated with AI in education, such as data mining and the potential misuse of personal data, we aim to address these concerns by proposing the aforementioned privacy-preserving methodologies. Last but not least, while a vast amount of literature exists on the benefits of AI in education, there is also a notable gap in research that specifically addresses privacy-preserving technologies. This study aims to fill this gap by

building on the insights of previous case studies, and by introducing new approaches in personal privacy. It also tries to establish a secure, ethical, and human-centered framework for AI in lifelong education.

All proposals aim to operate within AI-empowered lifelong education environments and provide a holistic solution to data privacy concerns. By implementing these privacy-preserving methodologies, this research aims to establish a secure, ethical, and human-centered framework for AI in lifelong education. The proposed solutions not only mitigate privacy risks but also emphasize the enduring importance of the human educator, ensuring that technology complements rather than replaces the personal and ethical dimensions of teaching. Finally, we need to state that this study lacks empirical evaluation and real-world data sources. Our approach is primarily theoretical at this stage, with no actual real-world data collected to validate its effectiveness. As such, the study should be considered purely as a Proof of Concept.

2. Materials and Methods

Most educators struggle to understand the difference between security, privacy and confidentiality, which many researchers consider as the triptych of safeguarding personal privacy in education (Figure 1).

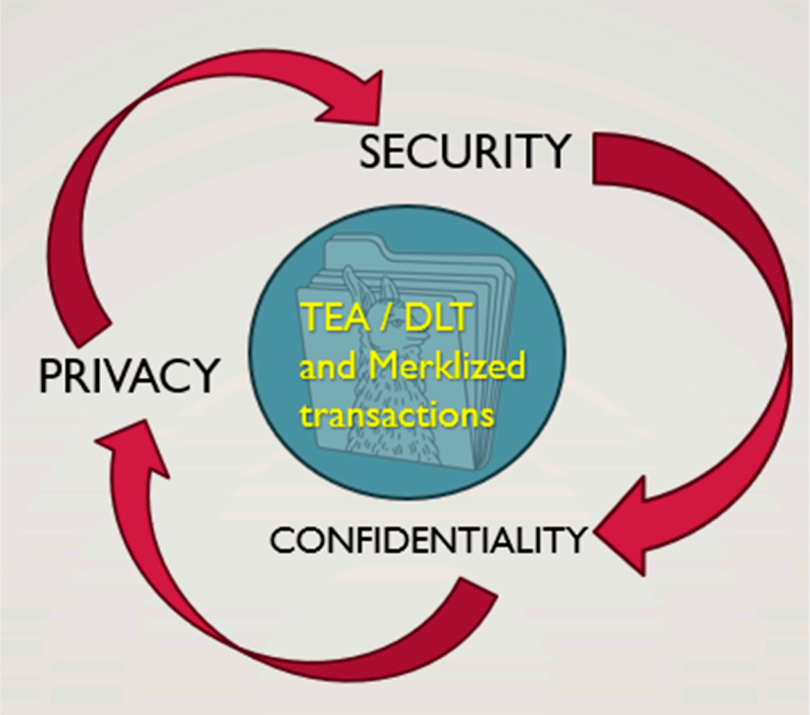

Figure 1. The triptych of safeguarding personal privacy in education.

The definition of privacy corroborates that every student's personal details belong solely to the student and therefore should never be made known to a third party. That information may include, but is not limited to, the evaluation of student grades, progress, and individual characteristics such as age, birthday and place of birth. To safeguard the privacy of students, their personal information should be disclosed only when it is absolutely necessary. It is essential to communicate the importance of confidentiality to teachers and students, in addition to how their personal information can be misused or disclosed. In the classroom, the teacher is usually responsible for ensuring the security of student data by keeping their information private and confidential by implementing firm practices in the classroom (Scheid, 2019). However, in the age of online Large Language Models (LLMs),

the responsibility falls on the companies that operate them. The question here is rather evident: Should we trust the companies who operate such models with our children's personal details? (Melendez & Pasternack, 2019; Holmes et al., 2022; COPPA, 2013).

2.1. TEA and DLT Method

In the history of accounting, double entry bookkeeping has been in use for almost 2500 years (Tarquini, 2016). The true innovation came with the introduction of triple-entry accounting (TEA), introduced by Financial Cryptographer Ian Grigg, which is a financial model that records transactions in three separate accounts by adding a cryptographic signature as a digital receipt (Grigg, 2005). He reasoned that the classical system of double-entry accounting is susceptible to deception and mistakes, as it is established on the honesty and legitimacy of the parties involved. It builds over traditional double-entry accounting by adding a third entry that is recorded in a Distributed Ledger technology medium, (i.e., a blockchain). DLT technology offers significant advantages for data privacy protection. Also, it offers permanent records, transparency, and the ability to audit. Permanent records on a DLT cannot be altered or deleted, transparency enables real-time visibility into data usage, and auditability facilitates tracking of data access and potential breaches. In TEA, each transaction is based on a cryptographic proof, verified by multiple parties, eliminating the reliance on a single trusted intermediary.

More specifically, the steps from Double Entry to Triple Entry are the following:

- TEA incorporates a third entry, known as the 'reputation entry', to provide additional transparency and trust in financial transactions.
- In triple-entry accounting, each transaction is recorded not only in the debit and credit entries but also in a separate entry that captures the 'reputation' or 'proof' of the transaction (Sunde & Wright, 2023).
- This reputation entry is created using cryptographic techniques and serves as an immutable record that can be verified by all parties involved in the transaction.
- TEA forms the foundation of blockchain technology (Sgantzos et al., 2023; Ibañez et al., 2023; Arunda, 2023).

We utilized a form of this technology for controlling LLMs in our previous work (Sgantzos et al., 2023), which can be altered and establish an effective framework for tracking and managing sensitive educational data (in this case, students' grades, or personal details), and enhance privacy, and security. On the other hand, decentralized and transparent systems like DLT provide immutable and verifiable online records, further supporting TEA's objectives.

In a similar fashion, we present an analogous setup:

1. We begin with Pseudonymization of student's personal details.
2. We make a string containing Pseudonymous ID + Student grades for each student.
3. We store the data on the blockchain (either for each student or as a class)

Result: Grades can be reviewed, without revealing the name (Figure 2).

In contemplation of maintaining the confidentiality of the participants, we utilized the principle defined in the Chatham House Rule (Chatham House, 2022). As the rule specifically describes:

> "When a meeting, or part thereof, is held under the Chatham House Rule, participants are free to use the information received, but neither the identity nor the affiliation of the speaker(s), nor that of any other participant, may be revealed".

Maintaining the TEA principle, we secure the three respective parties:

1st party: Teacher/AI Educator
2nd party: Students
3rd party: DLT record.

A thorough review of scientific research on data privacy in AI-empowered education (Cullican, 2023; Holmes et al., 2022; Jennings, 2023; Akgun & Greenhow, 2021; Chan, 2023; Archambault, 2021) shows that multiple researchers have investigated the challenges associated with data collection, data storage, sharing and aggregation. The biggest issue seems to be data usage from companies that employ LLMs for improving these models via algorithmic techniques like Reinforcement Learning from Human Feedback (RLHF), to eliminate hallucinations and algorithmic bias (Boutin, 2022; Manyika et al., 2019). Another technique is the Retrieval Augmented Generation (RAG) that uses the retrieval of information from some confirmed sources (e.g., the Internet or a database) to prevent the model producing incorrect or illogical answers (Proser, 2023). These studies propose various solutions such as user education, anonymization, pseudonymization, access control and encryption to address these issues.

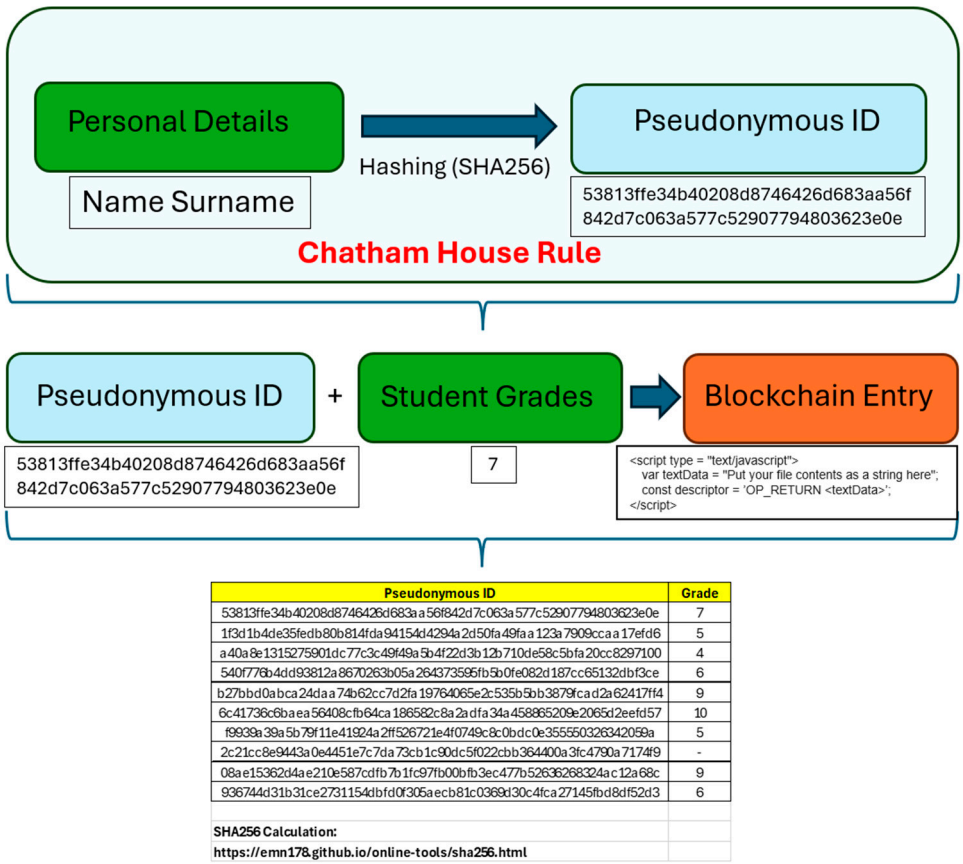

Pseudonymous ID	Grade
53813ffe34b40208d8746426d683aa56f842d7c063a577c52907794803623e0e	7
1f3d1b4de35fedb80b814fda94154d4294a2d50fa49faa123a7909ccaa17efd6	5
a40a8e1315275901dc77c3c49f49a5b4f22d3b12b710de58c5bfa20cc8297100	4
540f776b4dd93812a8670263b05a264373595fb5b0fe082d187cc65132dbf3ce	6
b27bbd0abca24daa74b62cc7d2fa19764065e2c535b5bb3879fcad2a62417ff4	9
6c41736c6baea56408cfb64ca186582c8a2adfa34a458865209e2065d2eefd57	10
f9939a39a5b79f11e41924a2ff526721e4f0749c8c0bdc0e355550326342059a	5
2c21cc8e9443a0e4451e7c7da73cb1c90dc5f022cbb364400a3fc4790a7174f9	-
08ae15362d4ae210e587cdfb7b1fc97fb00bfb3ec477b52636268324ac12a68c	9
936744d31b31ce2731154dbfd0f305aecb81c0369d30c4fca27145fbd8df52d3	6
SHA256 Calculation:	
https://emn178.github.io/online-tools/sha256.html	

Figure 2. Graph of the TEA and DLT record method (Sgantzos et al., 2023).

2.2. Enhanced Privacy for DLT Stored Data Using Merklized Transactions

The second method was recently presented by J. Davis (2024) and offers a very efficient and relatively simple solution to the data privacy problem. Davis presented a novel scheme to assemble blockchain transactions by utilizing Merkle trees composed of transaction fields. The process ensures that the transactional data are verified field-wise using Merkle proofs. The procedure can be applied either at the system level or as a second layer protocol that does not require changes to the underlying DLT medium (i.e., blockchain). The benefit of this technique is that it allows users to efficiently verify stored information

by separately checking targeted individual data items stored in transactions. The system provides a lightweight "Proof of Existence" for anything stored on a blockchain by introducing a secondary block Merkle tree (Figure 3).

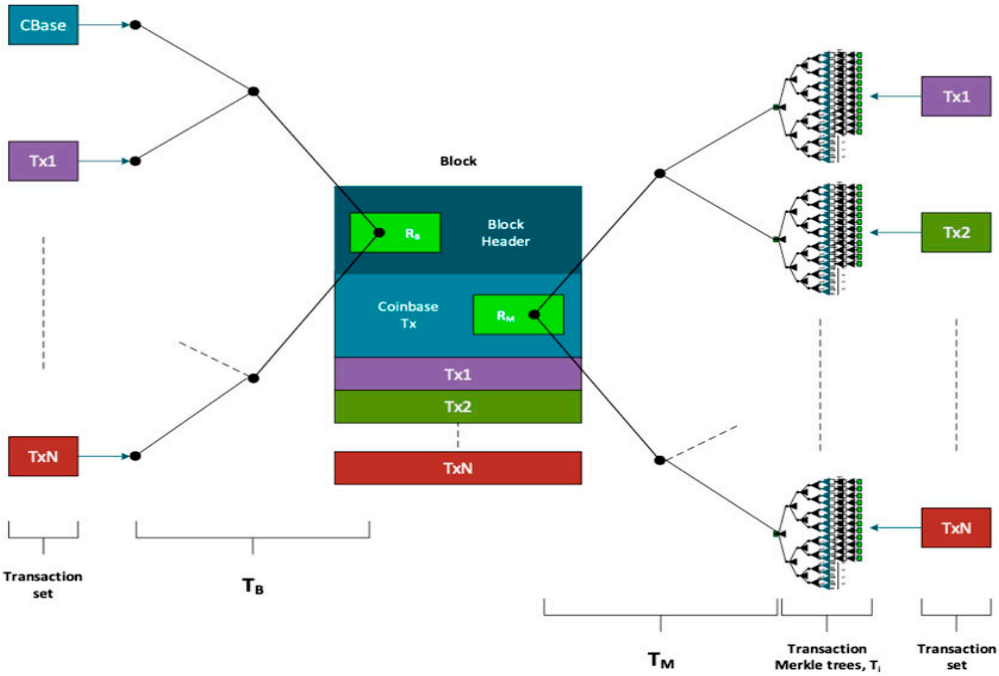

Figure 3. A schematic outline of the block generation protocol for Merklized transactions (J. Davis, 2024).

In a very similar way to the DLT and TEA method presented in Section 2.1, this methodology can be used to solve the privacy problem in tracking and managing sensitive educational data and ensuring privacy and security.

2.3. The Use of an Offline LLM to Ensure Privacy

The third method employs a freely accessible isolated (i.e., offline) model. Offline Large Language Models [LLMs] offer several distinct benefits compared to online models like Bing (Thompson, 2023), Bard (Manyika & Hsiao, 2023), and online Generative Pretrained Transformers (GPT), such as ChatGPT and its counterparts. These benefits are primarily attributed to their ability to process uninterrupted and continuous data without requiring constant connectivity or real-time updates. Different implementations give different opportunities, but at the cost of added complexity. For instance, GPT4All (NomicAI, 2023), can parse a directory of documents which can then be processed as a "short memory" from the LLM. Other implementations need a full installation of a programming environment like Python 3.10 together with the respective libraries to function effectively.

As a case study, we chose to have a stand-alone and offline "AI-educator" to simulate the environment of a remote education session. For this purpose, we employed a decades-old, rather cheap system based on 2013 technology (such systems can be found used at the range of about EUR 300–600), just to demonstrate that the hypothesis stands even in low-income households. The technical specifications of the system in our experimental setup were as follows:

Computer Model: Apple, Los Altos, California, USA, MacPro (Late 2013)
CPU: Intel, Santa Clara, California, USA, Xeon E5 2.7 GHz, 12-Core,
GPU: AMD, Santa Clara, California, USA, FirePro D300, 2 Gb
RAM: OWC, Woodstock, Illinois, USA, 64 Gb, 1866 MHz DDR3.
HDD: Apple, Los Altos, California, USA, 512 Gb NVMe

The average response speed for our answers came at a speed of 8 tokens per second, which compares to the average speed of an online model like ChatGPT-3.5 or Microsoft's, Redmond, WA, USA, Co-Pilot. We selected a program called "llamafile" (Hood, 2023), which is easy to install and easy to use, with the absolute minimal requirements for an efficient usage of LLaVA model as an AI-Educator. llamafile allows the user to turn large language model weights into executables. It is formed by a single executable that the user can download on a home PC, and it supports all the common platforms (Microsoft's, Redmond, WA, USA, Windows™, Apple, Los Altos, CA, USA, OSX™, San Francisco, CA, USA, Linux). LLaVA is a multimodal LLM that can do more than just chat functionalities. Multimodality in AI implementations seems to be the main recent trend in LLMs, as we also proposed in a previous work (Sgantzos et al., 2022). In this particular model, the user can also upload images and ask questions about them. With LLaVA, the process happens locally; no data ever leave the computer. The model can answer questions ranging from simple elementary school problems to advanced higher education classes (Figure 4).

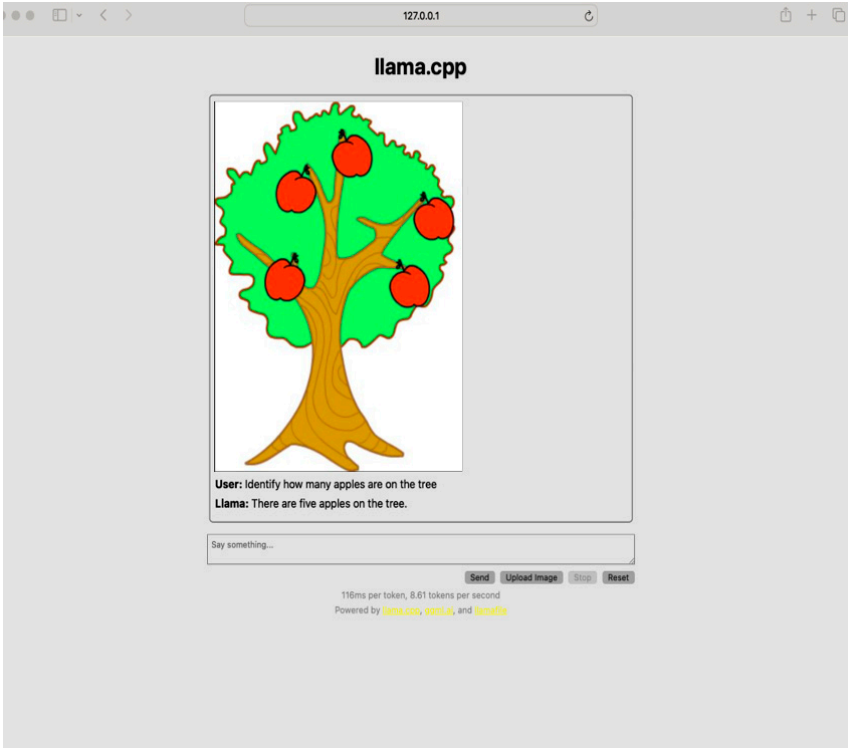

Figure 4. LLaMA counts successfully how many apples are in the picture.

However, in more advanced mathematical cases, the model fails to correctly answer when asked what the square root is of 225, both visually and textually (Figure 5).

However, with newer versions of the model and observing the results of a quantitative analysis based on MMMU (Massive Multi-Discipline Multimodal Understanding) (Yue et al., 2023) of this model's larger versions (i.e., LLaVA-NeXT-34B) with 34 billion tokens, it can perform at the levels of GPT-4V and Gemini (Liu et al., 2023). We see great improvement in mathematical problems such as the one we demonstrated above, but we also argue that there is still a long way to go for a single student to have an AI tutor without human supervision.

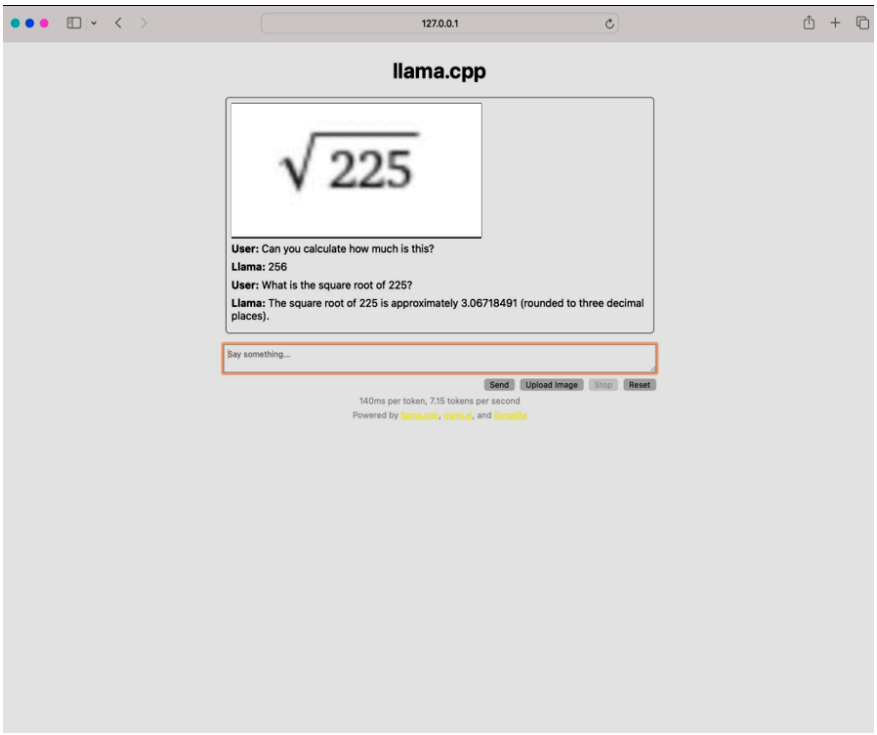

Figure 5. LLaMA fails to answer correctly in both cases.

We next provide a systematic comparison to justify the usage of TEA-DLT, Merklized transactions and offline LLMs over other existing privacy-preserving methods:

Double Entry Bookkeeping compared to TEA-DLT:
Limitations: Traditional accounting methods such as the double-entry bookkeeping system are basically based on the trust of the parties involved, making those systems susceptible to fraud and errors. They do not provide a standalone verification mechanism, meaning that discrepancies can often go unnoticed.
Justification for the use of TEA and DLT: TEA introduces a third entry, adding a layer of transparency and traceability that double entry bookkeeping method lacks. DLT records entries on a decentralized ledger, ensuring data integrity and immutability which significantly mitigates the risk of unauthorized alterations.

User education compared to TEA-DLT/Merklized transactions/Offline LLMs:
Limitations: There is an inalienable truth in the history of computing: that almost always, with a few exceptions of machine failure, the single point of failure in system security is the user. Therefore, it is impossible for somebody to ensure that the user is fully aware of all vulnerabilities.
Justification for the use of our proposed methods: TEA introduces a trustless way to mitigate user errors and non-compliance. Using a cryptographic proof for each entry, it isolates the trust from the associated parties and therefore ensures system security. The same stands for the Merklized transactions and offline LLMs.

Privacy-preserving Techniques (anonymization, pseudonymization):
Limitations: While classic techniques like anonymization and pseudonymization help reduce identifiable information via obscuring personal details, these techniques remain vulnerable to re-identification through advanced data analysis methods. This creates an ongoing privacy risk.
Justification for our proposed methods: TEA coupled with DLT offers stronger privacy protection. This combined approach allows transaction verification using cryptographic

signatures, for example Secure Hashing Algorithm (SHA-256) (Khovratovich et al., 2011) or Elliptic Curve Digital Signatures Algorithm (ECDSA) (Johnson et al., 2001), while keeping sensitive identifiers secure. By processing data in isolated, encrypted environments and recording transactions on a distributed ledger, the system significantly reduces the risk of data being traced back to individuals. Finally, the offline LLMs provide a priori the isolation of personal data.

Centralized systems:
Limitations: Current privacy protection systems often concentrate data storage and management in centralized locations. This centralization creates vulnerable targets for attackers. In case of a compromised system, it can result in massive data breaches and privacy violations, since all protected information is stored in one place. The concentration of sensitive data in a single location essentially creates a "honey pot" that attracts malicious actors and increases the potential impact of any successful breach.
Justification for our proposed methods: By decentralizing data storage through distributed ledgers, TEA enhances security against potential single-point failures. The strong point of this approach enables multiple parties to validate and audit transactions without relying on a central authority. Merklized transactions also rely on the same TEA principles, while offline LLMs preserve privacy through isolation.

The proposed architectures represent three neoteric approaches towards ensuring privacy preservation in AI-empowered lifelong education by combining the benefits of TEA and DLT technologies, for the online models, thus enabling personalized learning experiences while protecting sensitive educational data effectively. By combining the strengths of TEA and DLT, we can effectively protect sensitive educational data while enabling AI to deliver personalized and transformative learning experiences. In the case of offline LLMs, there are no privacy issues.

3. Results

The first two methods, presented in (Section 2.1) Triple-Entry Accounting combined with Distributed Ledger Technology, and (Section 2.2) Merklized Transactions, can be used separately with any online model or combined for optimal results. The third method presented in (Section 2.3), Offline LLMs, aims to present a privacy alternative to their online counterparts.

One of the most impressive features of the Large Language Models is their ability to generate coherent and continuous speech, even when the logic behind the sentence is paradoxical. For instance, in the question "how a child's height affects the height of the father?" it will try to answer thoroughly as though the previous sentence was sound. The reason this happens is because the next word is chosen from the enormous "vocabulary" (tokens) they have been trained with, to match the previous one more closely.

But is it a problem or a feature of the technology? Andrej Karpathy, who worked as a senior engineer at OpenAI, believes that AI Hallucination is the modus operandi of these models and characterizes them as "dreaming machines" (Karpathy, 2023). Other pioneers in AI such as Yann LeCun argue that if we are to improve AI, we must recognize that the technology in question has begun to reach a "plateau", and it would be good to start looking for other ways and new ideas. He also argues that AI will conquer the world, but not enslave people (LeCun, 2023).

On the side of offline models, there are a lot of advantages, which include the following:

1. Enhanced efficiency with low latency due to the absence of network delays and communication bottlenecks that can hinder online models' performance in processing large amounts of data simultaneously.

2. Improved privacy, as offline models do not require continuous internet connectivity or access to user-specific information, thus reducing potential security risks associated with online data transmission and storage.
3. Increased scalability since offline models can be preprocessed and stored locally, allowing for faster loading times and reduced resource consumption during inference compared to their online counterparts that require continuous updates via the internet.
4. No maintenance costs: offline models eliminate the need for cloud infrastructure, data storage, and bandwidth expenses associated with maintaining an online model's functionality.

Drawbacks of Offline LLMs: despite these advantages, offline models face several limitations:

1. Limited real-time responsiveness and adaptability since offline models do not receive continuous updates or access to new data sources, which can reduce their ability to respond quickly to evolving situations or provide up-to-date information in real-time scenarios.
2. Increased storage requirements due to the need for large amounts of memory and processing power to store preprocessed offline models, potentially leading to higher hardware costs and resource constraints during inference.
3. Potential wrong or outdated answers if the model is not properly maintained or updated over time, as old information may become inaccurate or irrelevant without continuous refinement through online updates.
4. Greater dependency on preprocessing steps, since offline models require extensive preparation before they can be used for inference, which might increase the overall computational complexity and time required to process data compared to their online counterparts which can quickly adapt to new inputs in real time.

As a last note, we must state here that the offline LLMs, like LLaMa, present inconsistencies even in simple mathematical problems (Figure 5). It is evident that technology has a long way to go before such a model can be used as a "home-educator". However, we must not berate their capabilities in content creation and their usage as virtual assistants. In most cases their responses are remarkably accurate. It is clear though, that a human is always needed as an evaluator.

Privacy and Ethical Concerns: Privacy is a critical, yet very underrated matter in the field of AI, and more specifically in LLMs. The problem begins with the data contained within a dataset that those models are being trained on, down to the feedback and queries the companies collect and use to make these models better. Human feedback is vastly utilized to improve the answer integrity and accuracy of models such as ChatGPT (Ethayarajh et al., 2023). The users who are asking the questions are actively contributing to the model's future effectiveness by rating the answers they receive. It is a notable fact that about 60% of the weighted pre-training dataset that was used for GPT-3 comes from a filtered version of Common Crawl (Commoncrawl, 2023) which contains 410 billion byte-pair-encoded tokens of webpages from around the world (Huggingface, 2023). Recently, a legal battle began against OpenAI led by the author and entertainer Sarah Silverman, as well as horror and fantasy author Christopher Golden and the novelist Richard Kadrey. They litigated against both OpenAI and Meta in a US District Court over dual rights of copyright infringement; and this case seems to be only the beginning (W. Davis, 2023). In AI-assisted lifelong education, this issue grows more serious for online AI models. What happens to the personal data of the students? How is it possible to safeguard the spread of sensitive information like birth dates, grades, or lesson progress? The ethical question here is: "Cui bono"? Who benefits most from the innovation this technology has to offer?

Is it the user, or the company who released the model? While it is difficult to have a definitive answer, we should make sure that certain failsafes are in place. But what if, instead of asking if we can trust a company that provides the AI model and how they may use the students' personal data and information, we would introduce methods to remove the trust question via a trustless protocol?

Proposed Methods for Privacy Preservation:

The three methods we proposed in this manuscript can isolate personal identifiers of students from the class while using AI models. All proposed methods ensure compliance with the current GDPR regulations and mitigate the risk of extremely high fines for the educational organization. Even in non-severe cases, such as those related to organizational and technical issues, may result in fines of up to EUR 10 million or 2% of total worldwide annual revenue, whichever is higher. However, none of these methods answers another ethical dilemma, which is the question of whether any of the online or offline models now available are able to replace the human teacher factor in education. This, along with other questions of an ethical nature, are discussed in the Discussion section. We next present a comparative analysis of all three methods (Table 1).

Table 1. Comparative Analysis of all three proposed methods.

Criteria	Triple-Entry Accounting (TEA)	Merkle Trees	Offline Models
Data Integrity	Ensures integrity through cryptographic verification and consensus mechanism among parties.	Uses cryptographic hashes to verify individual transaction data integrity through tree structure.	Depends on the integrity of the model's preloaded data and requires manual updates for changes.
User Privacy	Provides privacy through pseudonymization; sensitive data is not stored in plain text.	Offers privacy by allowing verification of data without revealing underlying information due to hash function.	Enhances privacy by storing data locally without network exposure; user data are not transmitted.
Scalability	Potentially scalable but could be challenged by increased transaction volumes requiring processing power.	Highly scalable due to efficient storage and verification of large data sets using tree structures.	Scalability may be limited by local hardware capabilities and the need for frequent updates.
Real-time Accessibility	Provides real-time access to data and transactions, important for educational settings.	Not inherently designed for real-time usage, verification may cause slight delays.	Limited real-time flexibility; dependent on prior updates and local processing power.
Implementation Complexity	Requires more sophisticated infrastructure and expertise for setup and maintenance.	Easier to implement in existing blockchain structures; relatively straightforward in integration.	Implementation complexity varies; requires extensive preprocessing, but can be simpler for isolated environments.
Compliance with Regulations	Designed to align with data protection regulations (e.g., GDPR) by explicitly managing consent and access rights.	Generally complies depending on implementation; requires careful management of user data within the blockchain.	Can comply with data regulations by isolating sensitive information and ensuring local storage without transmission.
Cost-Effectiveness	May incur higher costs with initial setup and ongoing operational expenses.	Generally more cost-effective for verifying large datasets due to minimal computational overheads.	Potentially cost-effective as it reduces dependence on cloud services and network infrastructure.
User Experience	Users may require knowledge of the system; the learning curve can impact usability initially.	Generally offers a seamless experience for users as they benefit from efficient data verification.	Enhanced user satisfaction due to ease of access to resources without potential online restrictions.

Finally, we present a (theoretical) quantitative analysis for each method considering the respective metrics presented in the comparative analysis (Table 2).

Table 2. Quantitative analysis of all three proposed methods.

Criteria	Triple-Entry Accounting (TEA)	Merkle Trees	Offline Models
Data Integrity (Score out of 10)	9	8	7
User Privacy (Score out of 10)	8	9	9
Scalability (Transactions per second)	Blockchain-dependent	Blockchain-dependent	Token-dependent
Implementation Cost (EUR)	Fees-dependent. Can be as low as EUR 0.00001 per record	Fees-dependent. Can be as low as EUR 0.00001 per record	EUR 300–600 for the current setup
Implementation Time (Months)	1	1	1
User Experience (Satisfaction Score out of 10)	7	8	9
Compliance Score (out of 10)	9	9	10
Efficiency (Resource Usage %)	Blockchain-dependent: 1–70%	Blockchain-dependent: 1–70%	40% compared to the online models.

Empirical Validation through comparison with existing case studies:

Our present work is not the first to try to address the problem of user privacy. Reviewing the current literature regarding privacy-preserving technologies in AI-empowered education, we try to empirically validate our proposed methods.

Several case studies give insights into current alternatives. In their work, Akgun and Greenhow (2021) discuss data privacy challenges in online learning environments. They enumerate the Ethical concerns and potential risks of AI applications in education, and they compare case studies from different educational institutions that have implemented AI. They addressed their successes and failures in dealing with privacy concerns and provided practical insights on an ongoing issue. On privacy-preserving AI implementations, Chan (2023) presents an analysis of various privacy-preserving methodologies in AI applications across different sectors, including education. Moreover, Holmes et al. (2022) conducted longitudinal studies on the impact of security measures in educational settings. On user education and privacy, the findings from Boutin (2022) discuss the impact of user education on privacy outcomes and can be used to create case studies that compare educational outcomes in institutions that prioritize user awareness versus those that do not. Finally, Archambault (2021) outlines privacy regulations and frameworks affecting diverse student populations in privacy frameworks across demographics.

We compare the aforementioned case studies with the manuscript's proposed methods in the following table (Table 3). The comparative analysis evaluates how these three methods improve privacy protection, enhance user trust, address implementation challenges, and increase effectiveness in educational settings.

Table 3. Empirical validation via comparative analysis with other case studies.

Method/Case Study	Privacy Preservation	User Trust	Implementation Challenges	Overall Effectiveness
Proposed Method: Triple-Entry Accounting (TEA)	High	Moderate	Complexity in managing and verifying transactions	High
Proposed Method: Merkle Trees	High	Moderate	Integration into existing systems may be challenging	High
Proposed Method: Offline AI Models	Very High	High	Difficulty in periodic updates and current knowledge	Moderate
Akgun and Greenhow (2021)	Moderate	Low	Varies by institution and technology adopted	Moderate
Chan (2023)	Moderate to High	Moderate	Differences in application effectiveness across sectors	High
Holmes et al. (2022)	Moderate to High	High	Data gathering and longitudinal study complexity	High
J. Davis (2024)	High	Moderate	Complexity of blockchain integration	High
Boutin (2022)	Moderate	Low	Educational implementation varies; reliance on user commitment	Moderate
Archambault (2021)	High	Moderate	Varies significantly by demographic and institutional policy	Moderate to High

The above comparative analysis against current research reveals several key insights:

Privacy Preservation: The evaluated approaches demonstrate strong privacy safeguards, with offline AI models providing the highest level of protection by completely eliminating online vulnerability exposure.

User Trust: Solutions that operate offline generally inspire greater confidence among users, as they eliminate data transmission concerns. Alternative approaches may encounter varying degrees of trust challenges based on algorithmic transparency perceptions.

Implementation Challenges: Our case analysis highlights that deployment difficulties vary significantly depending on technological complexity and context. While our proposed methodologies face distinct obstacles, they often benefit from more structured implementation frameworks.

Overall Effectiveness: When compared to existing approaches documented in the literature, our proposed methods demonstrate superior performance, particularly in scenarios demanding robust privacy protections.

4. Discussion

The three methodologies presented in this work, namely TEA and DLT, Merklized transactions, and offline LLMs, form three separate new frames of reference for assisting the role of the tutor. Undoubtedly, the relationship a teacher may have with the student is a distinctive one. Capturing the nature of that relationship requires an approach which recognizes the characteristic ways in which the dispositions and motivations that govern it are guided (Oakley & Cocking, 2001, p. 51). One such approach is, for example, "what it means to be a good teacher". This conception, which is personal and particularistic and not universalist, operates as a controlling parameter or regulative ideal in determining whether to take on the task of teaching. In this case, the teacher does not look at the situation in a detached manner. He/she does not step back to adopt a stance reflecting an impartial, uni-

versal perspective, figuring out what that viewpoint suggests to anyone situated similarly to her/himself (Stelios, 2020, p. 112). The teacher's motivation is of a more particularistic nature. He/she builds the bond with the student, approaches them and delivers based on the particular special needs and characteristics of the person.

Personalization and/or personalized learning refers to tailoring teaching to accommodate specific individuals through the collection of personal data. Before the advent of information and communication technologies, the personalized manifestation of education depended on the virtues of the teacher. Today, although the individual needs of each student can be captured using sophisticated models, it is not certain that they can be addressed without the presence of the human factor. The role of a human tutor remains particularly important.

Additionally, decent professional roles must be part of a decent profession. A good profession is one which involves a commitment to the human good; a good "which plays a crucial role in enabling us to live a humanly flourishing life" (Sen, 1985, p. 74). But is there a link between serving professional roles and serving human good? We believe so. For instance, if we can consider the enhancement—through formal, non-formal and informal learning—of the quality of life of both individuals and their collectives in different stages and domains of life as the central goal or telos of lifelong education, then teaching certainly counts as a good profession (Stelios, 2020). A good teacher is one whose motivation, planning, and action towards their students are guided by a certain awareness of what the activity of teaching involves. This conception of what it is to be a good teacher is ideally shaped by virtues.

By interacting with human activity, AI-empowered technology can address challenges learners may experience within the framework of their work practices. In particular, the Capabilities approach (Sen, 1985) can serve as a telos—actually, a shift of telos—for lifelong education (Rubenson, 2019; Tuijnman & Boström, 2002). Instead of AI-based personalized technologies in work and everyday life that deal primarily with learning implications identifying skill development as an economic investment, framing learning through the Capabilities approach is focused on agency, equity, freedom and on a broader human development view. The emphasis is on personal satisfaction, aspirations and autonomy and not so much on the development of skills. Human well-being is as important as widespread economic guiding concepts and human capital theory.

A (hypothetical) syllogism emerges from the above. Achieving privacy through the methods mentioned above can lead to a personalized form of education that will consider the diverse data and, ultimately, characteristics of each adult student. This personalization in education is extremely important. It represents the good and virtuous professional role of the teacher which, within lifelong education, takes on a more humanistic character through the Capabilities approach.

It is noteworthy how the above syllogism (Figure 6) starts from an informational and practical approach to privacy to conclude that there is a humanistic conception of AI-empowered education. Of course, to the extent that an online or offline AI-based model can (a) not only ensure privacy but also (b) exploit personalization by offering comprehensive lifelong learning, then it can replace the human teacher. The second option is clearly quite difficult to achieve and for this reason we believe that humans should be present in teaching processes.

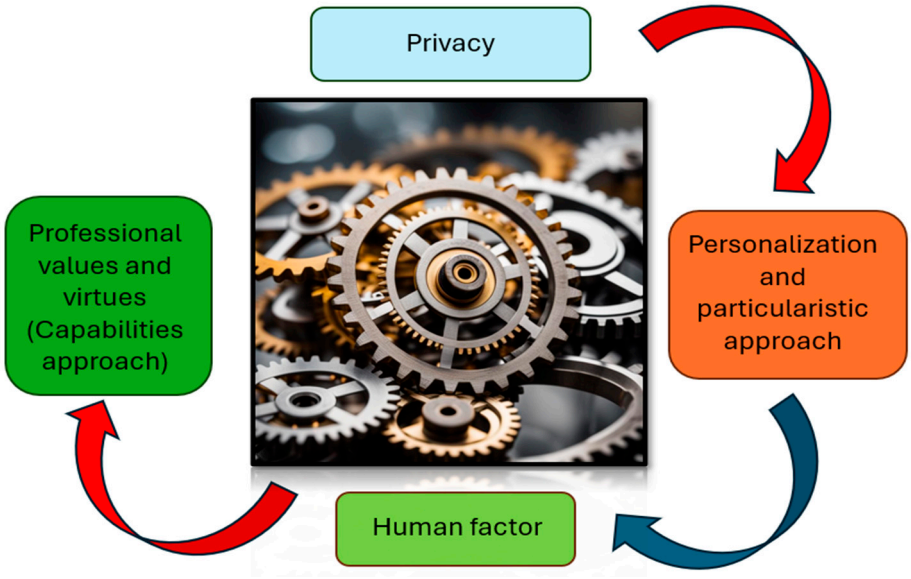

Figure 6. An illustration of the syllogism.

Below, we present the real-world challenges that occur for each method, based on scalability, difficulty of implementation, costs and usability. Table 4 outlines the challenges associated with each method, making it easier to compare and understand the key issues that may arise during implementation, which we try to address below.

Table 4. Real-world challenges for each proposed method.

Method	Security	Privacy	Computational Cost
Triple-Entry Accounting (TEA) and DLT	High; employs cryptographic verification and consensus mechanisms, ensuring data integrity and security among parties.	Provides privacy through pseudonymization; sensitive data are not directly stored.	May incur higher initial setup costs and ongoing operational expenses.
Merkle Trees	High; utilizes cryptographic hashes to verify individual transaction data integrity, maintaining security.	Allows verification of data without revealing sensitive information, enhancing user privacy.	Generally more cost-effective for large datasets due to minimal computational overhead.
Offline Models	Moderate; relies on the inherent security of the local environment and lack of external threats.	Enhances privacy by storing data locally and not requiring online access.	Potentially cost-effective initially but long-term costs for hardware maintenance and data storage must be considered.

Practical Considerations for Privacy Implementation

To successfully integrate privacy-preserving methods in AI-enhanced lifelong education, several key practical factors must be addressed:

Institutional Factors

Educational institutions need strong leadership support and mind-changing practices to foster innovation while prioritizing data privacy. They should consider upgrading the technological infrastructure, faculty training needs, and take into account existing legal policies. By making partnerships with industry, they can gain expertise and resources for implementing privacy solutions.

Regulatory Factors

Institutions must comply with privacy laws like GDPR. A thorough and detailed compliance strategy should be developed, with legal expertise to navigate regulatory complexities and reduce risks.

Economic Factors

Privacy-preserving implementations are expensive and require financial investment, and the cost depends on the method. Institutions should conduct cost-benefit analyses, consider long-term savings from preventing data breaches, explore funding opportunities, and apply changes gradually throughout a multi-year policy to manage budgets.

Policy Contexts

The existing educational policies also affect technology implementation. Institutions should review their current policies for compatibility with new methodologies and engage in advocacy for regulations that support innovation while protecting privacy.

Industry Collaboration

Partnerships with specialized technology companies provide tools, training, and support for implementing privacy-preserving methods. Such collaborations provide knowledge sharing and development of solutions aligned with educational needs.

By addressing institutional, regulatory, economic, and policy considerations while fostering industry partnerships, we believe that a much needed framework will be created, for educational institutions to effectively leverage AI while protecting student privacy.

5. Conclusions

The integration of AI in education has opened up new possibilities for lifelong learning. However, it is crucial to address concerns regarding privacy and data protection when utilizing AI-powered technologies to maintain a secure environment. This research paper presents three distinct novel approaches to ensure privacy in AI by employing techniques such as Triple-Entry Accounting with Distributed Ledger Technology, Merkle trees, and an offline multimodal AI tutor model that operates without internet access. These solutions demonstrate the potential for balancing privacy concerns while harnessing the benefits of AI in lifelong education. However, this study lacks empirical evaluation and real-world data sources. Our approach is primarily theoretical at this stage, with no actual real-world data collected to validate its effectiveness. As such, the study should be considered purely as a Proof of Concept.

This research aims to mitigate privacy concerns in AI for lifelong education and the enormous financial risks that come in case of privacy violations for the educational organization. It has examined methods using advanced technologies like Triple-Entry Accounting, blockchain and Merkle trees to achieve this. Lastly, an isolated AI model is proposed that is able to create an AI assistant that can function offline, protecting user data while still offering the advantages of AI in education. There are, however, real-world challenges that cannot be ignored. For instance, there are certain difficulties in the implementation of TEA and DLT regarding the complexity of the method. The same happens with the Merklized transactions. On the other hand, the difficulty of periodic updates to attain current knowledge is the burden of Offline LLMs and can be challenging for remote schools where these models could be more useful. In any case, further research will have to include real-life application of each method to evaluate the strong and weak points presented in this manuscript.

Within lifelong education, personalized learning is an important factor directly related to privacy. What seems to be the case is that it also represents a guide to achieving the ideal AI-empowered lifelong education as it relates to the virtues of the human tutor. This humanitarian underlying nature of education should be considered when referring

to technological approaches and models. The syllogism presented can be a foundation for further research which we plan to address separately and thoroughly in a forthcoming manuscript.

Below, we enumerate and summarize the manuscript's contributions, limitations, and future work directions:

Contributions:

We present three new privacy-preserving approaches for AI in education:

- Triple-Entry Accounting with Distributed Ledger Technology
- Merkle trees implementation
- Offline multimodal AI tutor model without internet connectivity

All three methods are meant to balance privacy protection with AI benefits in lifelong learning and they provide solutions for mitigating financial risks associated with privacy violations for educational organizations.

Limitations:

- There are implementation challenges with Triple-Entry Accounting and DLT due to method complexity
- Technical difficulties with Merklized transactions implementation
- Offline LLMs face challenges with periodic updates to maintain current knowledge
- Limited accessibility for remote schools where these models could be most beneficial

Future Work Directions:

- Real-world application testing of each method to empirically evaluate strengths and weaknesses
- Further research into personalized learning as a key factor in privacy-preserving AI education
- Exploration of the humanitarian aspects of education in relation to technological approaches
- A planned forthcoming manuscript addressing personalized learning and the ideal AI-empowered lifelong education system

Author Contributions: Conceptualization, K.S.; methodology, K.S.; validation, K.S., P.T., M.A.H. and E.R.P.; writing—original draft preparation, K.S.; writing—review and editing, K.S., P.T., M.A.H. and E.R.P.; visualization, K.S.; supervision, K.S. and M.A.H.; project administration, K.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The authors would like to thank Spyridon Stelios and Georgios Papageorgiou for their helpful discussions and proofreading on early drafts of this work.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
DLT	Distributed Ledger Technology
ECDSA	Elliptic Curve Digital Signature Algorithm
GDPR	General Data Protection Regulation
GPT	Generative Pre-trained Transformers
GPT4All	Generative Pre-trained Transformers for All
ID(s)	Identification(s)
LLaMA	Large Language Model Meta™ AI
LLaVA	Large Language and Vision Assistant
LLM(s)	Large Language Model(s)
RAG	Retrieval Augmented Generation
RLHF	Reinforcement Learning from Human Feedback
SHA-256	256-bit Secure Hashing Algorithm
TEA	Triple-Entry Accounting

References

Akgun, S., & Greenhow, C. (2021). *Artificial intelligence in education: Addressing ethical challenges in K-12 settings*. Available online: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8455229/ (accessed on 17 January 2024).

Almpanis, T., & Paul, J.-R. (2022). Lecturing from home: Exploring academics' experiences of remote teaching during a pandemic. *International Journal of Educational Research Open*, *3*, 100133. [CrossRef]

Archambault, S. (2021). Student privacy in the digital age. *BYU Education & Law Journal*, *2021*(1), 6. Available online: https://scholarsarchive.byu.edu/byu_elj/vol2021/iss1/6/ (accessed on 18 January 2024).

Arunda, B. (2023). *Blockchain in accounting: Triple-entry accounting*. Available online: https://www.linkedin.com/pulse/blockchain-accounting-triple-entry-benjamin-arunda/ (accessed on 17 January 2024).

Billett, S. (2010). The perils of confusing lifelong learning with lifelong education. *International Journal of Lifelong Education*, *29*(4), 401–413. [CrossRef]

Boutin, C. (2022). *There's more to AI bias than biased data. NIST report highlights*. Available online: https://www.nist.gov/news-events/news/2022/03/theres-more-ai-bias-biased-data-nist-report-highlights (accessed on 16 January 2023).

Chan, C. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, *20*, 38. [CrossRef]

Chatham House. (2022). *The Chatham House rule*. Available online: https://www.chathamhouse.org/about-us/chatham-house-rule (accessed on 17 January 2024).

Children's Online Privacy Protection Act (COPPA). (2013). *Children's online privacy protection rule: A six-step compliance plan for your business. Federal trade commission*. Available online: https://www.ftc.gov/business-guidance/resources/childrens-online-privacy-protection-rule-six-step-compliance-plan-your-business (accessed on 23 January 2024).

Commoncrawl. (2023). *The common crawl corpus*. Dataset. Available online: https://commoncrawl.org/get-started (accessed on 20 January 2024).

Cullican, J. (2023). *AI in education: Privacy concerns and data security: Navigating the complex landscape*. Available online: https://www.linkedin.com/pulse/ai-education-privacy-concerns-data-security-complex-jamie-culican/ (accessed on 15 January 2024).

Dave, R. H. (Ed.). (1976). *Foundations of lifelong education*. UNESCO Institute for Education, and Pergamon Press.

Davis, J. (2024). Enhanced scalability and privacy for blockchain data using Merklized transactions. *Frontiers in Blockchain*, *6*, 2624–7852. [CrossRef]

Davis, W. (2023). *Sarah silverman is suing open AI and meta for copyright infringement*. Available online: https://www.theverge.com/2023/7/9/23788741/sarah-silverman-openai-meta-chatgpt-llama-copyright-infringement-chatbots-artificial-intelligence-ai (accessed on 20 January 2024).

De Laat, M., Joksimovic, S., & Ifenthaler, D. (2020). Artificial intelligence, real-time feedback and workplace learning analytics to support in situ complex problem-solving: A commentary. *The International Journal of Information and Learning Technology*, *37*(5), 267–277. [CrossRef]

Ethayarajh, K., Choi, Y., & Swayamdipta, S. (2023). *Stanford human preferences dataset*. Available online: https://huggingface.co/datasets/stanfordnlp/SHP (accessed on 20 January 2024).

Gan, W., Qi, Z., Wu, J., & Lin, J. C. (2023). *Large language models in education: Vision and opportunities*. Available online: https://arxiv.org/pdf/2311.13160.pdf (accessed on 12 January 2024).

GDPR. (2025). *GDPR fines/penalties*. European Commission. Available online: https://gdpr-info.eu/issues/fines-penalties (accessed on 7 February 2025).

Grigg, I. (2005). *Triple entry accounting*. Systemics Inc. Available online: https://iang.org/papers/triple_entry.html (accessed on 12 January 2024).

Holmes, W., Persson, J., Chounta, I. A., Wasson, B., & Dimitrova, V. (2022). *Artificial intelligence and education: A critical view through the lens of human rights, democracy and the rule of law*. Council of Europe. ISBN 978-92-871-9236-3. Available online: https://rm.coe.int/prems-092922-gbr-2517-ai-and-education-txt-16x24-web/1680a956e3 (accessed on 15 January 2024).

Hood, S. (2023). *Introducing llamafile*. Available online: https://hacks.mozilla.org/2023/11/introducing-llamafile/ (accessed on 16 January 2024).

Huggingface. (2023). *Byte-pair encoding tokenization*. Available online: https://huggingface.co/course/chapter6/5?fw=pt (accessed on 20 January 2024).

Ibañez, J. I., Bayer, C. N., Tasca, P., & Xu, J. (2023). REA, triple-entry accounting and blockchain: Converging paths to shared ledger systems. *Journal of Risk and Financial Management*, *16*, 382. [CrossRef]

Jennings, J. (2023). *AI in education: Privacy and security*. Available online: https://www.esparklearning.com/blog/ai-in-education-privacy-and-security/ (accessed on 15 January 2024).

Johnson, D., Menezes, A., & Vanstone, S. (2001). The Elliptic Curve Digital Signature Algorithm (ECDSA). *International Journal of Information Security*, *1*, 36–63. [CrossRef]

Karpathy, A. (2023). *On the hallucination problem*. Available online: https://twitter.com/karpathy/status/1733299213503787018 (accessed on 16 January 2024).

Khan Academy. (2023). *Khanmingo AI tutor. Online AI tutor*. Available online: https://www.khanacademy.org/khan-labs (accessed on 12 January 2024).

Khovratovich, D., Rechberger, C., & Savelieva, A. (2011). *Bicliques for preimages: Attacks on Skein-512 and the SHA-2 family*. Available online: https://eprint.iacr.org/2011/286.pdf (accessed on 27 February 2025).

LeCun, Y. (2023). *How not to be stupid about AI, with Yann LeCun*. Available online: https://www.wired.com/story/artificial-intelligence-meta-yann-lecun-interview/ (accessed on 16 January 2024).

Liu, H., Li, C., Li, Y., Li, B., Zhang, Y., Shen, S., & Lee, Y. J. (2023). *LLaVA-NeXT: Improved reasoning, OCR, and world knowledge*. Available online: https://llava-vl.github.io/blog/2024-01-30-llava-next/ (accessed on 8 June 2024).

Manyika, J., & Hsiao, S. (2023). *An overview of bard: An early experiment with generative AI*. Available online: https://ai.google/static/documents/google-about-bard.pdf (accessed on 16 January 2024).

Manyika, J., Silberg, J., & Presten, B. (2019). *What do we do about the biases in AI?* Available online: https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai (accessed on 16 January 2023).

Melendez, S., & Pasternack, A. (2019). *Here are the data brokers quietly buying and selling your personal information*. Available online: https://www.fastcompany.com/90310803/here-are-the-data-brokers-quietly-buying-and-selling-your-personal-information (accessed on 23 January 2024).

NomicAI. (2023). *GPT4All: An ecosystem of open-source on-edge large language models*. Available online: https://github.com/nomic-ai/gpt4all (accessed on 16 January 2024).

Nowicki, J. M. (2020). *Data security: Recent K-12 data breaches show that students are vulnerable to harm*; GAO-20-644. Available online: https://www.gao.gov/products/gao-20-644 (accessed on 7 June 2024).

Oakley, J., & Cocking, D. (2001). *Virtue ethics and professional roles*. Cambridge University Press.

Poquet, O., & De Laat, M. (2021). Developing capabilities: Lifelong learning in the age of AI. *British Journal of Educational Technology*, *52*(4), 1695–1708. [CrossRef]

Proser, Z. (2023). *Retrieval augmented generation (RAG)*. Available online: https://www.pinecone.io/learn/retrieval-augmented-generation/ (accessed on 16 January 2023).

Rubenson, K. (2019). Assessing the status of lifelong learning: Issues with composite indexes and surveys on participation. *International Review of Education*, *65*(2), 295–317. [CrossRef]

Scheid, M. (2019). *The educator's role: Privacy, confidentiality, and security in the classroom. Student privacy compass*. Available online: https://studentprivacycompass.org/scheid1/ (accessed on 12 January 2024).

Sen, A. (1985). *Commodities and capabilities*. North-Holland.

Sgantzos, K., Grigg, I., & Hemairy, M. A. (2022). Multiple neighborhood cellular automata as a mechanism for creating an AGI on a blockchain. *Journal of Risk and Financial Management*, *15*, 360. [CrossRef]

Sgantzos, K., Hemairy, M. A., Tzavaras, P., & Stelios, S. (2023). Triple-entry accounting as a means of auditing large language models. *Journal of Risk and Financial Management*, *16*, 383. [CrossRef]

Stelios, S. (2020). Professional engineers: Interconnecting personal virtues with human good. *Business and Professional Ethics Journal*, *39*(2), 253–267. [CrossRef]

Su, J. (苏嘉红), & Yang, W. (杨伟鹏). (2023). Unlocking the power of ChatGPT: A framework for applying generative AI in education. *ECNU Review of Education*, *6*(3), 355–366. [CrossRef]

Sunde, T. V., & Wright, C. S. (2023). Implementing triple entry accounting as an Audit Tool—An extension to modern accounting systems. *Journal of Risk and Financial Management*, *16*(11), 478. [CrossRef]

Tarquini, L. (2016). *Il Falco e il Topo Manualetto di Gestione Aziendale*. Lulu.com. ISBN-10: 1326893939, ISBN-13:978-1326893934.

Thompson, A. D. (2023). *Microsoft bing chat" (Sydney/GPT-4)*. Available online: https://lifearchitect.ai/bing-chat/ (accessed on 16 January 2024).

Tuijnman, A., & Boström, A. K. (2002). Changing notions of lifelong education and lifelong learning. *International Review of Education*, *48*, 93–110. [CrossRef]

Vigliarolo, B. (2024). *Adobe users just now getting upset over content scanning allowance in Terms of Use*. The Register. Available online: https://www.theregister.com/2024/06/06/adobe_users_upset_over_content/ (accessed on 7 June 2024).

Yue, X., Ni, Y., Zhang, K., Zheng, T., Liu, R., Zhang, G., Stevens, S., Jiang, D., Ren, W., Sun, Y., Wei, C., Yu, B., Yuan, R., Sun, R., Yin, M., Zheng, B., Yang, Z., Liu, Y., Huang, W., … Chen, W. (2023). MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI. *arXiv*. [CrossRef]