Journal of
Risk and Financial Management

MDPI

Article

Triple-Entry Accounting as a Means of Auditing Large Language Models

Konstantinos Sgantzos [1,*], Mohamed Al Hemairy [1], Panagiotis Tzavaras [2] and Spyridon Stelios [3]

[1] Research Institute of Science and Engineering [RISE], University of Sharjah, Sharjah P.O. Box 27272, United Arab Emirates; malhemairy@sharjah.ac.ae
[2] Department of Management and Marketing, School of Business Administration, European University Cyprus, P.O. Box 22006, Nicosia 1516, Cyprus; ptzavaras@gmail.com
[3] Department of Human Sciences, Social Sciences and Law, School of Applied Mathematical and Physical Sciences, National Technical University of Athens, 9, Iroon Polytechniou Str., Zografou Campus, 15772 Athens, Greece; stelioss@mail.ntua.com
* Correspondence: sgacos@gmail.com

Abstract: The usage of Large Language Models (LMMs) and their exponential progress has created a Cambrian Explosion in the development of new tools for almost every field of science and technology, but also presented significant concerns regarding the AI ethics and creation of sophisticated malware and phishing attacks. Moreover, several worries have arisen in the field of dataset collection and intellectual property in that many datasets may exist without the license of the respective owners. Triple-Entry Accounting (TEA) has been proposed by Ian Grigg to increase transparency, accountability, and security in financial transactions. This method expands upon the traditional double-entry accounting system, which records transactions as debits and credits in two separate ledgers, by incorporating a third ledger as an independent verifier via a digitally signed receipt. The utilization of a digital signature provides evidentiary power to the receipt, thus reducing the accounting problem to one of the presence or absence of the receipt. The integrity issues associated with double-entry accounting can be addressed by allowing the parties involved in the transaction to share the records with an external auditor. This manuscript proposes a novel methodology to apply triple-entry accounting records on a publicly accessed distributed ledger technology medium to control the queries of LLMs in order to discourage malicious acts and ensure intellectual property rights.

Keywords: Triple-Entry accounting; AI Auditing; LLM transparency; AI ethics; blockchain; Web3

check for updates

Citation: Sgantzos, Konstantinos, Mohamed Al Hemairy, Panagiotis Tzavaras, and Spyridon Stelios. 2023. Triple-Entry Accounting as a Means of Auditing Large Language Models. *Journal of Risk and Financial Management* 16: 383. https://doi.org/10.3390/jrfm16090383

Academic Editors: Eva R. Porras and Thanasis Stengos

Received: 27 June 2023
Revised: 23 August 2023
Accepted: 24 August 2023
Published: 27 August 2023

1. Introduction

LLMs, like OpenAI's GPT-3 (Brown et al. 2020) or Microsoft's Sydney (Thompson 2023), are powerful algorithmic models that can generate human-like text based on user-input data using a combination of natural language processing (NLP) algorithms and stochastic analysis. These models employ tokenization, lemmatization, and part-of-speech tagging for identifying and classifying words, as well as predictive modeling for determining the relevance of each word to the context (Mearian 2023). In essence, a large language model takes into consideration the statistical probability of each word based on the dataset it was trained on. This helps in finding the most suitable words relevant to the input sentence, producing impressive continuance that can closely imitate genuine human writing. However, there have been concerns raised regarding the potential of LLMs to generate non-accurate results, fake statements, prejudiced or injurious material, or even malevolent code which can result in the creation of viruses or malicious software (Zacharakos 2023; McGuffie and Newhouse 2020).

In the NLP industry, almost all the data sets used for training AI models are carefully supervised and tokenized. The supervision process refers to the labeling or the feedback of the model, while tokenization is the compartmentalization of the text into smaller units called tokens. Tokenization is a heuristic process and can be carried out at different depths and levels, such as words, half-words, characters, and symbols (Menzli 2023). Unfortunately, the procedure does not guarantee that each token represents an accurate description or that the dataset is highly diverse, since supervision is sometimes inaccurate and the tokenization often comes with errors and ambiguities, such as incorrectly splitting words, missing critical information, or generating too few or too many tokens. Therefore, since humans are involved in the process, there is always a high probability of mistakes. Those mistakes are reflected in the results of the model as AI bias that originates from the supervision (Dilmegani 2022) and AI hallucination that derives from the tokenization (Marr 2023).

The approaches to overcome the problems of AI hallucinations and AI bias include diverse and representative datasets, rigorous quality control, and validation methods to ensure accuracy and consistency. Current research tactics aim to avoid any potential harm or errors on a societal level by monitoring and evaluating the performance of the model against different groups and outcomes. Finally, using active learning techniques such as RLHF to correct or avoid hallucinations showed promising results (Boutin 2022; Manyika et al. 2019). However, all these methods have social and privacy elements that need to be addressed without compromising the accuracy of the procedures.

While it is not proper to abate the technological gifts of AI and LLMs, we also need to address that in the wrong hands, they also provide the ability of intentional misuse and malice. Europol Innovation Lab published a report that identified three crime areas of major concern for the agency, based on the expertise of their specialists: Fraud and social engineering, Disinformation, and Cybercrime (Europol 2023). The report's research team aimed to acquire a profound comprehension of the technology from the Europol officers, as well as to develop their own LLMs and enhance the officers' existing knowledge, in order to equip them with the necessary competencies to address the new challenges of the forthcoming "age of AI". However, since the current AI development is following an exponential curve, we argue that more thorough methods need to be employed.

TEA is an advanced framework intended to augment traditional double-entry accounting by incorporating a third entry for each transaction. The fundamental principle of triple-entry accounting dictates that every transaction must involve three parties and three entries. As in double-entry accounting, the first two entries record the debit and credit effects on the accounting equation. The third entry serves as a cryptographic receipt that validates the transaction and is stored on a distributed ledger, such as a blockchain (Wright 2008). This system aspires to foster greater trust, transparency, and accuracy in financial reporting, while also minimizing opportunities for fraudulent activities and errors (Tyra 2023).

TEA was introduced by Grigg as a method to enhance transparency, accountability, and security of financial transactions (Grigg 2005). He argued that the conventional system of double-entry accounting is vulnerable to fraud and errors, as it depends on the integrity and veracity of the parties involved. He proposed that a third entry, which is distributed and verified by a network of independent nodes, could enhance the security and reliability of accounting records by creating a single source of truth for all transactions (Surana and Bhanawat 2021). This system, which he named triple-entry accounting, could eliminate the need for trust among the parties involved (Ibañez et al. 2021).

The Double-Entry accounting system was popularized by Luca Pacioli in 1494 (Lehrke 2023) nevertheless, it was based on previous methods dated back to ancient times (Tarquini 2016). TEA extends this conventional system, which records transactions as debits and credits in two distinct ledgers, by adding a third ledger that acts as an independent validator through a digitally signed receipt. The digital signature of TEA offers evidential value to the receipt and thereby simplifies the accounting problem to the existence or nonexistence of the receipt. The approach offers unassailable value in the accounting field due to its ability to provide far more accurate and reliable records compared to Double-Entry financial transactions. By merging TEA and Distributed Ledger Technology (DLT) records (i.e., a Blockchain) an audit trail is formed that any auditor can use to trace back transactions and identify any inconsistencies. The combination of both technologies helps to reduce the risk of errors and fraud and provides independent, third-party verification for every record (Andersen 2016). TEA combined with a publicly accessed DLT recording medium offers the capability to group multiple entries, such as financial dealings, contracts, or simple transactions, in a single-entry record and ensures a more efficient and accurate auditing process.

This paper introduces a novel methodology that leverages Triple-Entry Accounting (TEA) to create a verifiable audit trail of the interactions between users and Large Language Models (LLMs). The methodology consists of recording the user's input (prompt), the LLM's output (answer), and a third entry on a Distributed Ledger Technology (DLT) medium that ensures the integrity of the records. The third entry is generated by an algorithm that also pseudonymizes the records to protect the privacy of the parties involved. The parties are the user, the LLM, and a public immutable records ledger (e.g., a blockchain). The records are stored as entries on the blockchain, which can be accessed by anyone, thus enhancing transparency and accountability.

The paper presents an algorithmic TEA method that can be implemented as a smart contract on the same DLT medium as the records (on-chain) or on a separate system that can access the blockchain contents (off-chain). The method aims to prevent or expose malicious use of LLMs for purposes such as malware, virii, or fake news generation, which could have serious implications in domains such as finance, healthcare, and national security. The method also has potential applications for intellectual property protection of the sources used by each LLM. The paper argues that TEA offers a promising approach to address some of the ethical and social challenges associated with LLMs and to promote their responsible use. The method can be expanded in the future to control the intellectual property of the people whose work has been included in each LLM. The idea of TEA offers a promising approach to addressing some of the challenges associated with LLMs and could help to ensure that these systems are used in a responsible and ethical manner.

2. Materials and Methods

In this section, we describe the nature of LLMs and Transformer-based technologies together with the possible problems and Ethical concerns of their usage (Sections 2.1 and 2.2). We then proceed towards our method analysis via the Anonymization/Pseudonymization procedure (Section 2.3) and the methodology of gathering the information needed from the LLM input and output for creating a Triple-Entry Accounting record on a DLT medium (Section 2.4).

2.1. LLMs and GPT-3 Variants

LLMs such as Generative Pre-trained Transformer 3 (GPT-3) are autoregressive language models that use deep learning to produce natural language text and computer language code. They consist of a transformer neural network which is designed to process long sequences of data such as text. One of the most impressive implementations of GPT-3, due to its excess popularity, is Chat-GPT from OpenAI (Schulman et al. 2022). About 60% of the weighted pre-training dataset for GPT-3 comes from a filtered version of Common Crawl (Wikipedia 2023, Commoncrawl) consisting of 410 billion byte-pair-encoded tokens of webpages from around the world (Huggingface 2023). It also uses a sophisticated algorithmic technique to ensure grammatical correctness and relevance to the topic questioned. Recently, Reinforcement Learning from Human Feedback (RLHF) has been adopted as a technique that employs human feedback in order to enhance a language model through techniques derived from reinforcement learning. Language models are now able to replicate intricate human values to a model trained on a vast corpus of textual data due to RLHF. Human feedback is utilized to train models such as ChatGPT (Ethayarajh et al. 2023). In essence, the participants who are using the model are also contributing to the model's future efficiency.

There is an ongoing criticism regarding the ethics of the "free AI training contribution" policy of the LLMs participants without their explicit consent in the past few months (Van der Wal et al. 2023). Another unanswered question up to now is if the dataset OpenAI used for the training of its models, contains any intellectual property and if their respective owners have been asked for their permission to use their content. The company refused to release the details of the training of their latest model (GPT-4) claiming security reasons (Mollman 2023). The specific statement and the lack of transparency from Sam Altman raised even more questions and a regulatory ban from Italy (Burgess 2023).

Furthermore, under concerns about GDPR and personal data ownership, the RLHF can be considered as user-owned data so, the latest policy that was adopted by OpenAI has been changed. Data submitted through the API is no longer used in model training or other service advancements, given that users have explicitly opted in. A 30-day data retention policy is adopted as the default for API users, with the option to customize it depending on individual requirements. The pre-launch review has been eliminated, having been made feasible through optimizing automated monitoring systems. The Terms of Service and Usage Policies have also been streamlined, with meticulous enhancements on the terms relating to data ownership, in which users possess the input and output of the models (OpenAI 2023).

2.2. AI Ethics and Concerns of LLM Usage

The rapid changes that AI has brought to people's lives have raised profound ethical concerns. AI-based systems have the potential to give rise, among other things, to cyber-attacks, misinformation, embedding of biases, threats to human rights, and inequalities. These ominous prospects are also evident in the way society interacts with today's smart digital platforms. By assigning tasks to these systems to serve the purposes of daily needs, humans grant them the privilege of shaping the way they act, think, and work. In a way, individuals, as subjects, are developing digital behaviors (see Tzavaras and Stelios 2022). When teachers rely, for example, on an LLM to organize the syllabus of a course, they give up the (mental) obligation to refer to the authentic sources of information (see books, articles, etc.) but also to remember them.

This is the era of online content, where algorithms largely determine the decision-making process regarding the creation and dissemination of information. Various AI-based systems, even indirectly or passively, seem to program our lives. They set the regulatory framework, the ground rule: "If you want my services (e.g., a complete text describing the abovementioned course), you have to learn my operating rules and follow them." Humans adapt—perhaps even submit—to the standards of these platforms, leaving behind their traditional and analog nature. They follow the technology that they have shaped (Stelios 2023). Through this interaction, a moral dimension is manifested.

There should be careful consideration, though. There is not (yet) an AI-based system's intention to perform morally right or wrong actions. It is well known that "moral agency" is a qualification of an intentional system to which one can ascribe beliefs and desires (Dennett 1998, pp. 3, 9 as cited in Gunkel 2012, p. 19). It includes those agents whose actions are directed by or subject to some moral criteria or stipulation. Therefore, it is not clear whether an LLM, for instance, falls into this category. However, this is of little importance. Even as a means to an end (see fast calculations, security, valid results, and countless others), that is, having an instrumental ontological dimension, computational intelligence systems change their creator. They seem to have the last word. For example, the function of an LLM, which was designed by humans and developed through deep learning, will most likely be the standard text generator. The danger here is that in the same way that today's LLMs serve humans, in the future, more intelligent LMMs will manipulate them. By having the ability to generate human-like text at an unprecedented scale and speed, LLMs are able to 'hack' language, which is the basic means of human communication. They have the potential to influence informational processes in subtle and profound ways (Mensier 2023).

A greater focus on these ethical issues strengthens society's collective defense against this potentially threatening technological transformation. In LLMs, which, as mentioned, act as quasi-moral agents, the presence of an independent verifier can limit their ability to produce malicious content. Additionally, the establishment of a smart contract "screens out" users who would use anonymity for malware and phishing attacks.

2.3. Anonymization/Pseudonymization of the Participant(s)

As of today, GPT-3 and its known variants do not store queries. The model processes the queries given and generates an output based on the input. In order to record the usage of the LLM we need to account for the relevant regulations of GDPR regarding the privacy of the persons involved as per the principle described in the Chatham House Rule (Chatham House 2022). As the rule 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."

For example, for the usage of the Chat-GPT an OpenAI API key is needed to direct both inputs and outputs to a text file when a query is done. In order to follow the Chatham House Rule, the first party can be pseudonymized by hashing (or double hashing) the API key of the participant; then the hash can be used as a pseudonymous element in the session recording.

2.4. Blockchain Entry

The proposed method of this manuscript will produce a text file (data record) that includes the hash of the API key, the text produced by the user query, and the response of the LLM. As a final step, the data record will be added as an entry on a blockchain. By using this method, the records can be kept in a pseudonymous manner, since there is no direct link to the user's identity, but only a hash of the API and a public key linked to the blockchain entry. The method works best with a minimal fee for each entry and every data record can be audited at a later time. There is a need to carefully consider the financial aspect of the usability, scalability, and viability of the proposed method. A blockchain-based on minimal fees is required if the proposed method is to remain viable. To direct each ChatGPT query and associated response through API calls stored in a text file, a combination of Python libraries, such as OpenAI, requests, and JSON, must be employed. Queries can involve a single or multiple entries, all of which should be stored as a single transaction on the relevant blockchain.

There are two methods of data storage in a blockchain: on-chain storage and off-chain storage. On-chain storage stores data directly in the blockchain, resulting in a more expensive process due to the fees involved. Off-chain storage, on the other hand, only stores the metadata in the blockchain, resulting in less of a cost. The determination of which storage method to use depends on the duration demanded by the auditing process.

In the above graph (Figure 1), in the first stage, we describe the pseudonymization procedure aiming to comply with Chatham House Rule (first Entry). Then, after the query procedure starts, the output provided can be either code or text (second Entry). Both are archived in a text file (output) and the entries are recorded on a DLT medium, such as a Blockchain (third Entry). As a proof of concept, we provide a pseudo-code in Python which presents the steps necessary to perform the recording of the prompts into a text file which will then be sent as a data record on the blockchain. As noted in the TEA principle, the User is the first party, the LLM is the second party, and the Blockchain is the third party. The pseudocode in Python with detailed comments and code in JavaScript that provides the recording in the Blockchain can be found in Appendix A.

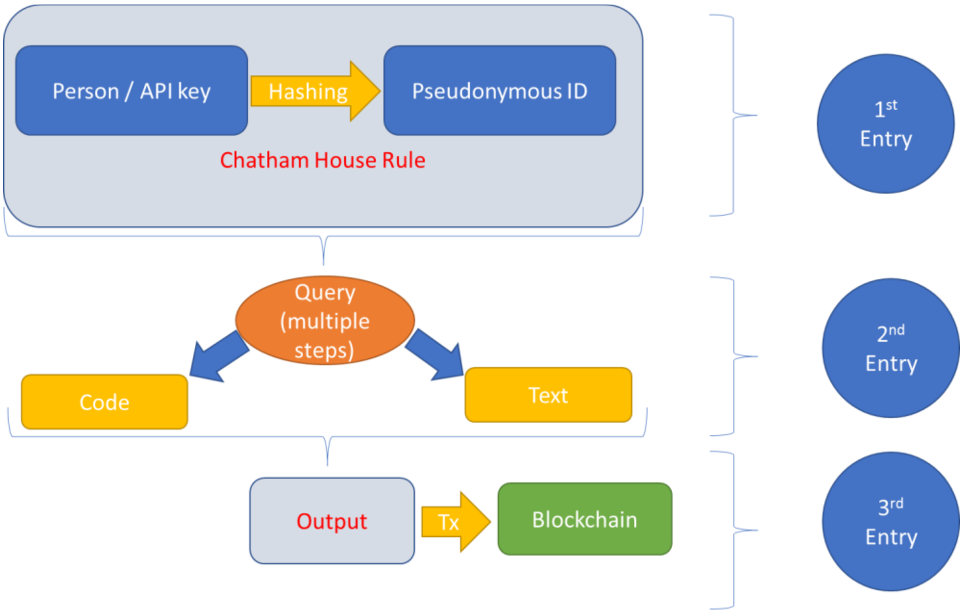

Figure 1. Graph of the Triple-Entry accounting for LLM.

3. Results

In the above section, we addressed the nature of LLMs and Transformer-based technologies together with some possible problems and ethical concerns that can occur by their usage. The problem of securing the participants' identity was addressed by employing the "Chatham House rule". We then proceeded to the methodology of gathering the information needed, whether this is plain text or produced code, from the LLM input and output. Finally, the last step was to create a Triple-Entry Accounting record on a Distributed Ledger Technology (DLT) medium.

The method provided is characterized by simplicity and it was created as a proof of concept by using the conjunction of the two technologies to achieve the best results at the lowest cost possible. The technology can be developed further and automated so that it can work in line with every LLM available freely or via micro-payments. Payments can also be done in blockchain tokens (e.g., stable coins) and can be used as payments for the usage of the LLM by the user as well. The usage of the hash on the user's API key maintains the pseudonymization of the person and certifies that only the information needed will be available during any possible audit, while when the action of the user is proven to produce a punishable act, then the auditor can lookup the person's identity via the respective hash and be de-anonymized to be held accountable.

It presents a notable fact, that the best practice to avert people from doing unlawful actions is to ensure that they know they are monitored for every action they perform. The parable of Plato's 'Ring of Gyges' (Woods 2010) describes the notion eloquently and provides proof that when transparency exists, the bad actors are kept in place. An extra measure of shifting from pseudonymization to anonymization of a person would be to automate the auditing process via an algorithmic procedure. In this way, the personal information of all the other participants involved can be ensured not to be accessible by any other physical person during the auditing procedure.

In this manuscript, we also examined a range of ethical concerns associated with the utilization of Language Models (LLMs) across various scenarios and contexts. Our analysis addresses the potential of LLMs, as AI-based systems, to significantly shape and influence human behavior, communication, and decision-making. Furthermore, we thoroughly discussed the challenges and risks inherent in LLMs, including cyberattacks, dissemination of misinformation, bias, violations of intellectual property and human rights, and exacerbation of societal inequalities. To address these pressing issues, we put forth a prospective solution that involves the implementation of an algorithm in the form of TEA and DLT technology and a smart contract mechanism, which can effectively ensure the accountability and transparency of LLMs and their users. By utilizing these measures, we can alleviate the negative impacts of LLMs on individuals and society. We provided the sample code in Appendix A that shows that such an algorithm can exist in the form of a smart contract which either can be stored on the same medium that the records exist (on-chain) or stored on a separate system (off-chain) that has access to the blockchain contents (oracle). Moreover, the method can be expanded to control the intellectual property of individuals whose works have been included in an LLM.

4. Discussion

LLMs are a new quantum technological leap in our world. There are many proponents of the technology and undoubtedly the contribution of LLMs to knowledge retrieval is extraordinary. However, this comes not without a price. AI Bias and AI Hallucination often decrease the quality of outputs of the models while at the same time, there are several voices coming from the AI research community that think that there will be a unique disparity in the world, mainly from the people who control the technology versus everybody else. It is true that right now there is no clear regulation for AI, while at the same time, there are certain cases where people who their works were involved in the LLM training process were never asked to participate (Davis 2023; Creamer 2023). A solution can be to use micropayments each time a work is accessed so that the creator can generate revenue, or grant a one-time fee each time a book is included in a model.

Given the global trend in finance, it is evident that the world is moving to a digitalized world. The age of Central Bank Digital Currencies (CBDC) is coming (Seth 2023) and it probably will be the norm within a decade. But what is fundamentally different from what we have now? If double-entry recordkeeping is maintained, there will be a time after some decades when the world will face the same problems it is facing today. A digital currency will not change human habits, and consequently will not solve the problems that the financial records have now. Some will argue that CBDCs will be backed by the trust and security of the government, but is this enough? We argue that an expansion of trust is needed if we want a more robust economic system than we have now. That expansion is dual and is intertwined between TEA and DLT public records (i.e., Blockchain). A system of trust needs to be fully auditable by everyone, not only the issuer. The answer to the question “who watches the watchers” must be “all of us”.

The development of Web3 passes through the technologies of TEA and DLT. As of today, about 45% of the connected devices to the Internet use come with IPv6 addresses while countries such as India lead the race for adoption by percentages over 65% (Google 2023). By the end of the decade, the IPv4 protocol will possibly be deprecated. The address of IPv6 is long enough to hold a Public Key and a Private Key and as such every IoT (Internet of Things) device will be able to communicate safely, securely, and peer-to-peer via micro-transactions that are worth a fraction of a cent. Those devices will be able to employ Artificial Intelligence agents like LLMs in the future, while several companies already are working on this pathway. By the tokenization of fully auditable digital currencies such as digital CHF (Centi 2023), the combination of TEA records and DLT will function as an “application” operating on IPv6, ultimately forming Web3, which can be termed as “The internet of value”. In the age of AI, the future of money will be “information”, and that information on Web3 will be accurate, pseudonymous, and fully auditable when needed. However, is it possible that those two technologies can help us to control Artificial Intelligence?

The difficulty is evident. How do we control the LLMs so that the technology can offer its benefits, but without the drawbacks? We suggest that TEA technology together with the immutability of records of a public DLT medium is the key answer. TEA fundamentally challenges double-entry bookkeeping, which was designed to enable a professional class of accountants to manage a firm’s financial status. TEA and DLT combine computer science, accounting, governance, and cryptography to create an immutable, shared record of transactions in the fabric of trade that is attested to by three participants. The internal transaction bookkeeping of current accounting practice, which forms the basis of our entire economy, leaves considerable scope for malfeasance, fraud, and other inefficiencies. Moving to an externally legible and verifiable practice would help to reduce many of these losses and costs (Southurst 2023).

5. Conclusions

5.1. Implications

There is a great concern regarding the economic impact of LLMs and AI in general, in the labor market worldwide (Eloundou et al. 2023). Some of them are the possible generation of harmful content by LLMs, disinformation and influencing operations, weapon development, privacy matters, and others (Sajid 2023). Many organizations already aim to replace some of their employees with AI automated agents, while a lot of work including data analysis, economics, and legal affairs, to name a few, tends to be totally reformed, since AI algorithms present a groundbreaking advantage compared to human-only centered positions. However, it is not AI that will replace humans, but humans with AI that will replace humans without AI (Lakhani 2023).

The future of work will probably demand a new set of skills and competencies that can supplement and exploit the potential of AI systems. The World Economic Forum reports that some of the most sought-after skills for 2023 and beyond are critical thinking, creativity, emotional intelligence, complex problem-solving, and interpersonal communication (Masterson 2023). These skills are crucial for human workers to adapt to the changing nature of work and to cooperate effectively with AI agents. Hence, education and training systems need to be revised and harmonized with the emerging needs of the labor market, and workers need to be ready for lifelong learning and continuous upskilling. Our approach can be used to discourage and eventually eliminate these problems by employing TEA and DLT technologies, powered by micropayments without compromising the privacy of the participants.

5.2. Policy Implications

The usage of a transparent solution such as the one we proposed in this manuscript promotes ethical behavior and aims to deter malicious actions of the users. It offers a solution to secure the intellectual property of creators, that their work was included in a dataset by using TEA records. Micropayments, or one-time fees, can also ensure that they can be paid each time a model uses or includes their work. Nevertheless, the method is not a panacea. There are several ethical and legal challenges such as finding the gold spot between privacy and transparency, the responsibility and liability between auditors and users, and finally how accurate and reliable algorithmic auditing can be. Those problems need to be addressed by developing a clear and consistent policy and regulation regarding LLMs.

5.3. Future Research

We suggest that the simplicity of the method we proposed in this work is essential to its success. However, there needs to be an evaluation of the performance and usability in different domains. Moreover, a proper comparison with other existing methods and technologies needs to be addressed. Another main concern is how the user perceptions and attitudes will change in regard to their behavior, trust, and willingness to use an LLM controlled by TEA and DLT technologies.

There is also the need to investigate the economic and social impact of the method and how it may affect the quality of services and diversity of users whether they are individuals, organizations, or communities. Finally, there needs to be a development of new features that may improve the method in terms of privacy, pseudonymity, and transparency on par with the efficiency of the results. In that process, maybe new functions need to be added to enhance the method.

This study uses the innovation of the TEA invention to extend its usage into a battle that is not yet fully developed. TEA is the basis of the blockchain technology. The idea of using the Blockchain to embed smart contracts and sophisticated AI agents to use them under a controlled and auditable environment to extend their capabilities towards a formulation of an Artificial "General" Intelligence (AGI), was first presented a year ago (Sgantzos et al. 2022). LLMs are often misinterpreted as a form of AGI, while many researchers worldwide believe that it will be the beginning of constructing one in the future.

Regardless, LLMs as they are now, also present an enormous potential regarding the possibilities to advance science and technology, but also bring a threat of misuse. Another issue is the intellectual property of the material they were trained on. Our findings underscore the imperative need for heightened attention to the ethical dimensions surrounding Artificial Intelligence and LLMs. By proactively addressing these concerns, we can prevent or mitigate their potential adverse consequences. With the introduction of TEA as a bookkeeping record keeper and auditing mechanism on a publicly accessed Distributed Ledger, the problem is minimized, and ethical usage is maximized.

Author Contributions: Conceptualization, K.S.; methodology, K.S.; software, K.S. and M.A.H.; validation, K.S., P.T., S.S. and M.A.H.; writing—original draft preparation, K.S. except for the AI ethics part in Section 2 by P.T. and S.S. and Appendix A, K.S. and M.A.H.; writing—review and editing, K.S., P.T., S.S. and M.A.H.; 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.

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 Georgios Papageorgiou and Joseph Harvey Perling, for their helpful discussions and proofreading on early drafts of this work.

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

Abbreviations

AI	Artificial Intelligence
API	Advanced Programming Interface
CBDC	Central Bank Digital Currency
DLT	Distributed Ledger Technology
GDPR	General Data Protection Regulation
GPT	Generative Pre-trained Transformer
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IoT	Internet of Things
JSON	JavaScript Object Notation
LLM	Large Language Models
NLP	Natural Language Processing
RLHF	Reinforcement Learning from Human Feedback
TEA	Triple-Entry Accounting
Web3	World Wide Web version 3.0

Appendix A

Pseudocode in Python with comments:

```
import openai
import requests
import json
import hashlib

# Replace the ’XXXXXX’ with your API key from OpenAI.
# Link:  https://platform.openai.com/account/api-keys
api_key = ’XXXXXX’

# Transform/Encode the API key string as UTF-8
api_key_bytes = api_key.encode(’utf-8’)

# Here we generate a hash_object via the SHA-256 hash function
hash_object = hashlib.sha256(api_key_bytes)

# Transform the hex representation of the hash as a string
hash_hex_string = hash_object.hexdigest()

# Then store the hash string in a txt variable
hashed_api_key = hash_hex_string

# print(hashed_api_key)  # Comment out this line if you want to see the hashed
API key

# We need to authenticate with OpenAI using your API key
openai.api_key = api_key

# path to the text file (replace the given one)
f_path = "c:\AIauditor\file.txt"

# Open the file and read its contents
with open(f_path, "r") as f:
   contents = f.readlines()

# Repeat the process through the lines in the file
for line in contents:
   # Send a query to ChatGPT using OpenAI’s API & parameters
   response = openai.Completion.create(
      engine="davinci",  # or any other engine you prefer
      prompt = line,
      max_tokens = 1024,
      n = 1,
      stop = None,
      temperature = 0.5
   )

   # Get the response from the generated text via the API
   answer = response.choices[0].text.strip()

   # Send the query and answer through an API call
   data = {
      "Hashed API Key":  hashed_api_key,
      "query":  line.strip(),
      "answer":  answer
   }
   headers = {"Content-Type":  "application/json"}
   url = "https://your-api-endpoint.com" # replace with your API endpoint
   response = requests.post(url, data = json.dumps(data), headers = headers)
```

The following pseudo-code in Javascript can be used to store a text file in the Bitcoin SV blockchain:

```
<script type = "text/javascript">
  var textData = "Put your file contents as a string here";
  const descriptor = ’OP_RETURN <textData>’;
</script>
```

Note: AI Tool used to create some parts and clean up the code: Wizardcoder GGML—Link: https://huggingface.co/TheBloke/WizardCoder-15B-1.0-GGML/tree/main (accessed on 2 June 2023).

References

Andersen, Nicolai. 2016. Blockchain Technology—A Game-Changer in Accounting? Available online: https://www2.deloitte.com/content/dam/Deloitte/de/Documents/Innovation/Blockchain_A%20game-changer%20in%20accounting.pdf (accessed on 11 August 2023).

Boutin, Chad. 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 11 August 2023).

Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33: 1877–901. [CrossRef]

Burgess, Matt. 2023. ChatGPT Has a Big Privacy Problem. Available online: https://www.wired.com/story/italy-ban-chatgpt-privacy-gdpr/ (accessed on 11 August 2023).

Centi. 2023. Centi Franc, the CHF Stablecoin. Available online: https://centi.ch/centi-franc/ (accessed on 11 August 2023).

Chatham House. 2022. The Chatham House Rule. Available online: https://www.chathamhouse.org/about-us/chatham-house-rule (accessed on 2 February 2023).

Creamer, Ella. 2023. Authors File a Lawsuit against OpenAI for Unlawfully ‘Ingesting’ Their Books. Available online: https://www.theguardian.com/books/2023/jul/05/authors-file-a-lawsuit-against-openai-for-unlawfully-ingesting-their-books (accessed on 11 August 2023).

Davis, Wes. 2023. Sarah Silverman Is Suing OpenAI 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 11 August 2023).

Dennett, Daniel C. 1998. *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge: The MIT Press.

Dilmegani, Cem. 2022. Bias in AI: What It Is, Types, Examples & 6 Ways to Fix It in 2023. Available online: https://research.aimultiple.com/ai-bias/ (accessed on 1 July 2023).

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv* arXiv:2303.10130. [CrossRef]

Ethayarajh, Kawin, Yejin Choi, and Swabha Swayamdipta. 2023. Stanford Human Preferences Dataset. Available online: https://huggingface.co/datasets/stanfordnlp/SHP (accessed on 10 August 2023).

Europol. 2023. ChatGPT—The Impact of Large Language Models on Law Enforcement, a Tech Watch Flash Report from the Europol Innovation Lab, Publications Office of the European Union, Luxembourg. Available online: https://www.europol.europa.eu/publications-events/publications/chatgpt-impact-of-large-language-models-law-enforcement (accessed on 11 August 2023).

Google. 2023. IPv6 Statistics, IPv6 Adoption. Chart. Available online: https://www.google.com/intl/en/ipv6/statistics.html (accessed on 24 August 2023).

Grigg, Ian. 2005. Triple Entry Accounting. Systemics Inc. Available online: https://iang.org/papers/triple_entry.html (accessed on 2 January 2023).

Gunkel, David J. 2012. *The Machine Question: Critical Perspectives on AI, Robots, and Ethics*. Cambridge: The MIT Press.

Huggingface. 2023. Byte-Pair Encoding Tokenization. Available online: https://huggingface.co/course/chapter6/5?fw=pt (accessed on 10 August 2023).

Ibañez, Juan Ignacio, Chris N. Bayer, Paolo Tasca, and Jiahua Xu. 2021. Triple-Entry Accounting, Blockchain and Next of Kin: Towards a Standardization of Ledger Terminology. Available online: https://arxiv.org/ftp/arxiv/papers/2101/2101.02632.pdf (accessed on 24 August 2023).

Lakhani, Karim. 2023. AI Won’t Replace Humans—But Humans with AI Will Replace Humans without AI. Available online: https://hbr.org/2023/08/ai-wont-replace-humans-but-humans-with-ai-will-replace-humans-without-ai (accessed on 11 August 2023).

Lehrke, Zoya. 2023. The Father of Accounting: Luca Pacioli. Available online: https://bench.co/blog/accounting/luca-pacioli/ (accessed on 11 August 2023).

Manyika, James, Jake Silberg, and Brittany Presten. 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 24 August 2023).

Marr, Bernard. 2023. ChatGPT: What Are Hallucinations and Why Are They a Problem for AI Systems. Available online: https://bernardmarr.com/chatgpt-what-are-hallucinations-and-why-are-they-a-problem-for-ai-systems/ (accessed on 24 August 2023).

Masterson, Victoria. 2023. Future of Jobs 2023: These Are the Most In-Demand Skills Now-and Beyond. Available online: https://www.weforum.org/agenda/2023/05/future-of-jobs-2023-skills/ (accessed on 11 August 2023).

McGuffie, Kris, and Alex Newhouse. 2020. The Radicalization Risks of GPT-3 and Advanced Neural Language Models. *arXiv* arXiv:2009.06807. [CrossRef]

Mearian, Lucas. 2023. What Are LLMs, and How Are They Used in Generative AI? Available online: https://www.computerworld.com/article/3697649/what-are-large-language-models-and-how-are-they-used-in-generative-ai.html (accessed on 10 June 2023).

Mensier, Anthony. 2023. LLMs Could Become Weapons of Mass Disinformation. Towards Data Science. Available online: https://towardsdatascience.com/llms-weapons-of-mass-disinformation-4def0dc3dc7 (accessed on 11 August 2023).

Menzli, Amal. 2023. Tokenization in NLP: Types, Challenges, Examples, Tools. Available online: https://neptune.ai/blog/tokenization-in-nlp (accessed on 11 August 2023).

Mollman, Steve. 2023. OpenAI Is Getting Trolled for Its Name after Refusing to Be Open about Its A.I. Available online: https://fortune.com/2023/03/17/sam-altman-rivals-rip-openai-name-not-open-artificial-intelligence-gpt-4/ (accessed on 20 August 2023).

OpenAI. 2023. Data Usage Policies. Available online: https://platform.openai.com/docs/data-usage-policies (accessed on 24 August 2023).

Sajid, Haziqa. 2023. 8 Ethical Considerations of Large Language Models (LLM) Like GPT-4. Available online: https://www.unite.ai/8-ethical-considerations-of-large-language-models-llm-like-gpt-4/ (accessed on 11 August 2023).

Schulman, John, Barret Zoph, Christina Kim, Jacob Hilton, Jacob Menick, Jiayi Weng, Juan Felipe Ceron Uribe, Liam Fedus, Luke Metz, Michael Pokorny, and et al. 2022. Introducing ChatGPT. Available online: https://openai.com/blog/chatgpt (accessed on 11 August 2023).

Seth. 2023. What Is a Central Bank Digital Currency (CBDC)? Available online: https://www.investopedia.com/terms/c/central-bank-digital-currency-cbdc.asp (accessed on 11 August 2023).

Sgantzos, Konstantinos, Ian Grigg, and Mohamed Al Hemairy. 2022. Multiple Neighborhood Cellular Automata as a Mechanism for Creating an AGI on a Blockchain. *Journal of Risk and Financial Management* 15: 360. [CrossRef]

Southurst, Jon. 2023. Triple Entry Accounting Matters. Available online: https://coingeek.com/why-triple-entry-accounting-matters-for-business-and-society-interview-with-ian-grigg/ (accessed on 11 August 2023).

Stelios, Spyridon. 2023. Artificial Intelligence or Artificial Morality? In *Technology, Users and Uses: Ethics and Human Interaction through Technology and AI*. Edited by J. Casas-Roma, Joan Conesa and Santi Caballé. Cambridge: Ethics International Press.

Surana, Gourav, and Shurveer S. Bhanawat. 2021. A Journey of Triple-Entry Accounting. *The Chartered Accountant*. Available online: https://ssrn.com/abstract=3751294 (accessed on 27 December 2022).

Tarquini, Luca. 2016. *Il Falco e il Topo Manualetto di Gestione Aziendale*. Morrisville: Lulu.com.

Thompson, Alan D. 2023. Microsoft Bing Chat (Sydney/GPT-4). Available online: https://lifearchitect.ai/bing-chat/ (accessed on 20 March 2023).

Tyra, Jason M. 2023. Triple Entry Bookkeeping with Bitcoin. Available online: https://bitcoinmagazine.com/business/triple-entry-bookkeeping-bitcoin-1392069656 (accessed on 11 August 2023).

Tzavaras, Panagiotis, and Spyridon Stelios. 2022. "Digital virtues?" Aristotelian Leadership in the Fourth Industrial Revolution. *Journal of Advanced Research in Leadership* 1: 1–8. [CrossRef]

Van der Wal, Oskar, Dominik Bachmann, Alina Leidinger, Leendert van Maanen, Willem Zuidema, and Katrin Schulz. 2023. Undesirable Biases in NLP: Averting a Crisis of Measurement. Available online: https://arxiv.org/pdf/2211.13709v2.pdf (accessed on 11 August 2023).

Wikipedia. 2023. Lemma: Commoncrawl. Available online: https://en.wikipedia.org/wiki/GPT-3 (accessed on 10 February 2023).

Woods, Cathal. 2010. Glaukon's Challenge (Republic 2). Available online: https://ssrn.com/abstract=1661519 (accessed on 4 August 2022).

Wright, Craig S. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. Available online: https://ssrn.com/abstract=3440802 (accessed on 2 January 2023).

Zacharakos, Alexis. 2023. How Hackers Can Abuse ChatGPT to Create Malware. Available online: https://www.techtarget.com/searchsecurity/news/365531559/How-hackers-can-abuse-ChatGPT-to-create-malware (accessed on 4 August 2022).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.