

On the Convergence of Artificial Intelligence and Distributed Ledger Technology: A Scoping Review and Future Research Agenda

KONSTANTIN D. PANDL^{ID}, SCOTT THIEBES^{ID}, MANUEL SCHMIDT-KRAEPELIN^{ID},
AND ALI SUNYAEV^{ID}

Institute of Applied Informatics and Formal Description Methods, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

Corresponding author: Ali Sunyaev (sunyaev@kit.edu)

This work was supported by the Karlsruhe Institute of Technology through the KIT-Publication Fund.

ABSTRACT Developments in artificial intelligence (AI) and distributed ledger technology (DLT) currently lead to lively debates in academia and practice. AI processes data to perform tasks that were previously thought possible only for humans. DLT has the potential to create consensus over data among a group of participants in untrustworthy environments. In recent research, both technologies are used in similar and even the same systems. This can lead to a convergence of AI and DLT, which in the past, has paved the way for major innovations of other information technologies. Previous work highlights several potential benefits of a convergence of AI and DLT but only provides a limited theoretical framework to describe upcoming real-world integration cases of both technologies. In this research, we review and synthesize extant research on integrating AI with DLT and vice versa to rigorously develop a future research agenda on the convergence of both technologies. In terms of integrating AI with DLT, we identified research opportunities in the areas of secure DLT, automated referee and governance, and privacy-preserving personalization. With regard to integrating DLT with AI, we identified future research opportunities in the areas of decentralized computing for AI, secure data sharing and marketplaces, explainable AI, and coordinating devices. In doing so, this research provides a four-fold contribution. First, it is not constrained to blockchain but instead investigates the broader phenomenon of DLT. Second, it considers the reciprocal nature of a convergence of AI and DLT. Third, it bridges the gap between theory and practice by helping researchers active in AI or DLT to overcome current limitations in their field, and practitioners to develop systems along with the convergence of both technologies. Fourth, it demonstrates the feasibility of applying the convergence concept to research on AI and DLT.

INDEX TERMS Artificial intelligence, blockchain, convergence, distributed ledger technology, machine learning.

I. INTRODUCTION

Artificial intelligence (AI) and distributed ledger technology (DLT) are among today's most actively debated developments in information technology with potential for tremendous impact on individuals, organizations, and societies over the next decades. A 2018 report from the McKinsey Global Institute estimates that the application of AI in various industries could deliver an additional global economic output of around USD 13 trillion by 2030 [1]. Similarly, a study by the World Economic Forum predicts that by 2025, up to 10 %

of the world's GDP may be stored on a blockchain [2], which is the most commonly used concept of DLT today [3].

AI can perform complex tasks that were previously thought possible only for humans to perform. In some application domains, AI can nowadays already exceed human capabilities. Research in the health care domain has, for instance, shown that AI can analyze echocardiograms faster and more accurate than medical professionals [4]. Furthermore, advancements in AI are also expected to be key enablers of important upcoming innovations such as autonomous driving [5] or intelligent robots [6], to name but a few. DLT, on the other hand, can provide consensus over a shared ledger in untrustworthy networks containing, for example, unreachable or maliciously behaving nodes [3]. It became known

The associate editor coordinating the review of this manuscript and approving it for publication was Jiafeng Xie.

to the public due to the emergence of the cryptocurrency Bitcoin [7]. Following the success of Bitcoin, further DLT applications are emerging in application domains beyond finance that are often corresponding to those of AI. DLT may, for example, be used to manage access control for electronic health records [8], or to secure the IT systems of autonomous cars [9].

As one result of these developments, we now increasingly see the emergence of applications using both information technologies in close integration. Recent work, for example, uses deep reinforcement learning to explore attacks on blockchain incentive mechanisms [10]. In doing so, the AI system can detect new attack strategies and provide security insights, even for well-studied DLT protocols like Bitcoin. Another recent work uses a DLT-based platform to exchange data and computing resources in order to enable AI applications [11]. The platform gives data providers the opportunity to share their data while keeping it confidential and maintaining the right to manage data access. Data consumers can then train algorithms on the provided data and compensate data providers for use of their data.

The examples above demonstrate that the integration of AI and DLT yields great potential to advance the capabilities of both technologies, and, ultimately, to increase the positive impact of AI and DLT on individuals, organizations, and societies. Yet, in order to make meaningful contributions, researchers and practitioners alike will have to keep up with the latest developments in both fields as well as the most recent developments and innovations related to their integration. Owing to the fast pace and interdisciplinary nature of both research areas, assessing the current state of research on AI and DLT and especially their integration in its entirety is a difficult task. In attempt to provide guidance to researchers enticed by the integration of AI and DLT, previous research has either focused on partial aspects of the integration of AI and DLT such as the use of blockchain for AI (e.g., [12]) or on deriving conceptual ideas of how both technologies might be integrated with each other (e.g., [13]). Despite the invaluable contributions that these publications have made to the nascent stream of literature concerning the integration of AI and DLT, we presently lack in-depth knowledge about the current state of research on the integration of AI and DLT that does (a) not only focus on a specific DLT concept (i.e., blockchain), (b) considers the reciprocal integration of both technologies (as opposed to the one-way integration of, for example, DLT into AI), and (c) goes beyond a purely conceptual level. In particular, we still lack a comprehensive overview of the most pressing research challenges that must be overcome in order to unleash the full potential of integrating both technologies. With this research, we aim to address this apparent knowledge gap by asking and answering the following two research questions:

RQ 1: What is the current state of research on the technological integration of AI and DLT?

RQ 2: What are open research challenges on the technological integration of AI and DLT?

To address our research questions, we draw on the concept of convergence (see section II.C) and conduct a systematic literature review on the current state of research on the convergence (i.e., integration) of AI and DLT and develop a future research agenda. The contribution of our work is thereby four-fold. First, in contrast to extant research in this area, we also include non-blockchain distributed ledgers in our analysis of the literature. Prior work has highlighted several current shortcomings of blockchain in the application for AI [12]. Other DLT concepts (e.g., BlockDAG or TDAG), may be more promising for solving some of these shortcomings [3] and thus potentially better suited for certain AI applications. Second, we consider the reciprocal nature of convergence and investigate both perspectives: the usage of AI for DLT, and the usage of DLT for AI. To the best of our knowledge, no holistic review on the convergence of AI and DLT exists today, which considers the large variety of interaction cases from both perspectives. Third, we aim to bridge the gap between theory and practice by drawing theoretical conclusions from practical research, as well as outlining future potential for practical and theoretical research from the theory in these fields. Fourth, we apply the concept of convergence [14], [15] as a theoretical lens for our article. In doing so, we contribute to the understanding of how convergence may drive product innovations and create economic value in the information technology (IT) industry. In addition, we demonstrate how convergence can be applied as a lens to tackle research questions in interdisciplinary, innovative, and emerging research fields.

The remainder of this article is organized as follows: In section 2, we discuss related work on AI and DLT, and introduce the concept of convergence. Afterward, we describe our methods in section 3. In section 4, we analyze the current literature on AI usage for DLT, and in chapter 5, provide our future research agenda on AI for DLT. In section 6, we analyze the current literature on DLT usage for AI, and a corresponding research agenda in section 7. In section 8, we discuss our results, before we conclude this article in section 9.

II. RELATED WORK

A. ARTIFICIAL INTELLIGENCE

AI enables computers to execute tasks that are easy for people to perform but difficult to describe formally. Such tasks typically occur in complex or uncertain environments [16]. Despite ongoing debates in society about *Artificial General Intelligence*, which describes computer programs that can control themselves and solve tasks in a variety of different domains [17], most deployed AI-based systems solve tasks in narrow application domains, and are referred to as *Narrow Artificial Intelligence*. Several approaches to design such narrow AI-based systems exist. For example, knowledge bases have seen a lot of attention by researchers in the past. Nowadays, Machine Learning (ML) seems to be the most well-spread approach toward building AI-based systems [16].

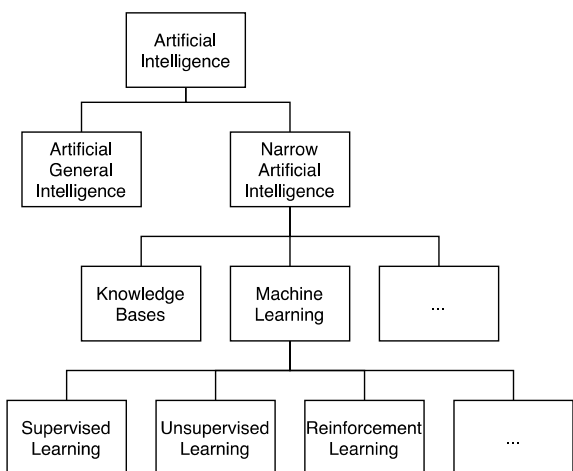


FIGURE 1. Overview of artificial intelligence.

ML-based systems consist of a model that represents a function between input data and output data. In most cases, ML models have to be trained. In this training phase, an optimization algorithm tweaks the model parameters in order to minimize a loss or maximize a reward. Depending on the application, different types of training exist. In the case of supervised machine learning, the input data and the corresponding output data are known during the training phase. In the case of unsupervised machine learning, only the input data is known but no output data. In a reinforcement learning setting, a learning agent executes actions that result in an immediate reward, whereas the agent’s goal is to maximize a future, cumulative reward. In general, the training phase can require large amounts of data and, thus, is often computationally intensive. This is especially the case for deep neural networks, which are complex ML models with many parameters that have paved the way for many of the recent advancements in ML [18]. In Figure 1, we present a high-level overview of different types of approaches toward designing AI-based systems.

The execution of a (trained) ML model is called inference. It is usually computationally less expensive than the initial training phase. Some models can only be described as black boxes, meaning that their inner functionalities are difficult to explain. Among many research streams, some cutting-edge research in ML aims to better explain the inner functioning of ML models in order to guarantee their robustness [19]. Other streams aim to increase ML systems’ capabilities [20], [21], or to ensure training data confidentiality when creating ML-based systems [16]. Besides these ML approaches introduced above, many variations exist. For example, for some ML algorithms, there is no explicit model building phase at all [22]. For a detailed overview of AI and ML, in particular, we refer to Russell and Norvig [23].

B. DISTRIBUTED LEDGER TECHNOLOGY

DLT enables the operation of a highly available, append-only, peer-to-peer database (i.e., the distributed ledger) in

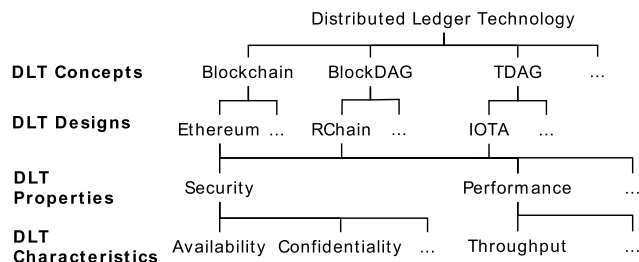


FIGURE 2. Overview of distributed ledger technology and its characteristics by Kannengießer et al. [3].

untrustworthy environments characterized by Byzantine failures, where separated storage devices (i.e., the nodes) maintain a local replication of the data stored on the ledger. A distributed ledger can either be deployed as a public ledger, if a new node can directly join a new network, or a private ledger, in case a node first needs permission to join the network. Another aspect to distinguish distributed ledgers are the write permissions: In the case of permissionless ledgers, all nodes have equal write permissions. In the case of permissioned ledgers, nodes first require granted permission to validate and commit new data to the ledger [3]. Today, several concepts for DLT exist with different characteristics, for example, regarding the transaction throughput or fault tolerance. The most widespread concept of DLT is blockchain [3], which became known to the public due to the emergence of the cryptocurrency Bitcoin [7]. Other concepts, for example, rely on directed acyclic graphs [3], [24], [25]. We provide an overview of DLT concepts with different characteristics in Figure 2. Nowadays, applications of DLT and especially blockchain beyond financial transactions are emerging, such as the management of medical data [26], autonomous driving [9], or decentralized games [27].

In a blockchain system such as Bitcoin [7], transactions of network participants are stored in blocks of a fixed size. Every block also includes a timestamp, and a hash of the previous block. Thus, the system can be regarded as a chain of blocks. Cryptographic techniques are used to ensure that only legitimate participants holding a cryptographic key can perform transactions, which are stored in the block. Bitcoin [7] was the first blockchain created and has mostly financial applications.

The main networks of the widely known Bitcoin and Ethereum [28] are public, unpermissioned blockchains. To secure the network, only a selected node can propose a new block that includes the cryptographically signed transactions. This node has to find a block candidate with a hash below a certain, network-defined threshold. As this hash calculation is impossible to reverse, a node has to make large computational efforts to find that new block, competing against other nodes. In return, the node that successfully found a block gets a reward in cryptocurrency payment. Since there are at any time many miners aiming to find the next block, the chances of finding the next block are relatively low for an individual miner. As a result, the variance of the mining payoff for an individual miner is relatively large. Therefore, most

miners nowadays mine together in so-called mining pools. If a mining pool finds a block, the pool collective gets the block reward and distributes it among its miners according to their share of hash calculations. This reduces the payoff variance for individual miners participating in the pool. Besides Bitcoin, several implementations of blockchains exist with different characteristics [3]. Some implementations provide high transaction confidentiality guarantees (e.g., Zcash [29]), or enable a universally usable, decentralized Turing complete computing platform (e.g., Ethereum [28]). In the latter case, programs can be stored and executed upon the blockchain system. These programs are referred to as smart contracts.

C. CONVERGENCE OF AI AND DLT

First described by Nathan Rosenberg in the 1960s, convergence describes a phenomenon in which two or more initially separate items move toward unity and become increasingly integrated with each other. Convergence typically occurs in four phases [15], [30]. The first phase, which is termed scientific convergence, occurs when distinct scientific fields begin citing each other, leading to an increase in cross-scientific research. It is followed by the technology convergence phase in which previously distinct technological fields increasingly overlap and new technology platforms arise. As technology convergence continues to blur existing technological and market boundaries, market convergence, the third phase, takes place, resulting in the emergence of new product-market combinations. In some cases, the convergence of technologies and markets may impact existing industries such that a fourth phase in form of the convergence of entire industries occurs. A typical and relatively recent example of convergence is the emergence of the smartphone, which nowadays combines initially separate technologies such as phones, cameras, and portable computing devices. It eventually created an entirely new market, completely transformed the mobile phone industry, and even overthrew the compact camera industry. Another example is Bitcoin [7], which was created by combining techniques from various computer science domains such as distributed systems, cryptography, security, and game theory.

Concerning a potential convergence of DLT and AI, we currently see the emergence of the first scientific publications that do not simply apply AI in the context of DLT (e.g., deep learning for prediction of the Bitcoin price) or vice versa, but instead discuss a deeper integration of both technologies. Extant literature that provides a consolidated overview of the integration possibilities of AI and DLT, however, is scarce. Dinh and Thai [13] provide a viewpoint with conceptual ideas of how an integration of AI and blockchain technology might look like. They outline possibilities for the reciprocal integration of AI and DLT, highlighting that AI can be used for blockchain, just as blockchain can be used for AI. Specifically, AI can support blockchain systems by increasing their security and scalability, by acting as an automated referee and governance mechanism, and for privacy-preserving personalized systems. Blockchain, on the

other hand, can serve AI systems by enabling decentralized computing, providing data sharing infrastructures, serving as a trail for explainable AI, or by coordinating untrusting devices. As the article is a viewpoint, the authors remain on a rather conceptual level and do not provide an in-depth review of extant literature. Salah *et al.* [12], on the other hand, provide a review and open research challenges on one of these two perspectives, the use of blockchain for AI. The authors identify enhanced data security, improved trust on robotic decisions, business process efficiency increases, collective decision making and decentralized intelligence as main drivers for the usage of blockchain for AI applications. Interestingly, the latter category has similarities with the application of privacy-preserving personalization, which Dinh and Thai [13] classify as an AI for blockchain perspective. Furthermore, Salah *et al.* [12] develop a taxonomy that provides an overview of different relevant technical characteristics, such as the consensus protocol or the blockchain type. Their review surveys research articles, as well as industry whitepapers. As such, it delivers valuable insights into how blockchain can be used for AI in vertical industry use case cases, such as banking, finance, or agriculture. The article concludes with several open research questions on technologies that are relevant in the context of blockchain for AI applications. These include topics of increasing blockchain systems' scalability and data confidentiality, interoperability, or quantum cryptography. Lastly, Karafiloski and Mishev [31] provide an overview of how blockchain can be used for storing and organizing large amounts of data. Although the authors do not consider AI use cases in much detail, the presented work is representative for some fields of the convergence of AI and DLT, such as data sharing.

In summary, prior research lacks an understanding how AI and DLT are reciprocally integrated in systems today, and how future research can advance the convergence of these technologies. Specifically, prior research focuses on blockchain, a specific DLT concept, and not on other potential DLT concepts that may be better suited in the context of AI [12], [13]. Furthermore, prior research only provides purely conceptual ideas how AI may be used for DLT, and does not evaluate today's technical feasibility of such systems [13].

III. METHODS

A. DATA COLLECTION

For the identification of articles addressing the convergence of AI and DLT, we systematically searched scientific databases. To cover a wide range of journal and conference publications, we queried IEEE Xplore, ACM Digital Library, AIS Electronic Library, Science Direct, Scopus, Ebsco Business Source, and ArXiv. The reason for including ArXiv preprints in our search is the fact that AI and DLT are fast moving fields of research, where new, potentially groundbreaking research results and insights may be available as preprints that are not finally published yet. Our search string required the publications to have a DLT-specific term and an

AI-specific term in either their title, abstract, or keywords. The search string, we employed was:

TIKEAB ((Blockchain OR „Distributed Ledger“ OR DLT OR Bitcoin OR Ethereum OR Cryptocurrency OR „Crypto currency“ OR „Crypto-currency“ OR „block chain“ OR „Smart contract“) AND (AI OR ML OR „Artificial Intelligence“ OR „Machine Learning“ OR „Deep Learning“ OR Clustering OR Classification OR „Neural Network“ OR „Big data“ OR „Data mining“ OR „Intelligent system“ OR „Statistical model“, OR „Statistic model“)).*

We excluded articles published before 2008, since the concept of Bitcoin emerged that year. For DLT, the search string included its most common concept of blockchain, as well as blockchain’s most frequent implementations (i.e., Bitcoin and Ethereum) and other technical terms. For AI, the search string included ML and the most common application forms of it. We searched on December 19th, 2019, this resulted in 2,411 unique articles.

In a first step, we first removed 140 articles published before 2008 from our set, as not all of the databases provided the option to include this as a search criterion. Afterward, we analyzed the titles, abstracts, and keywords for the remaining articles and removed 17 that were not written in English. For the remaining 2,254 articles, we not only analyzed the titles, abstracts, and keywords, but also their full texts. From this set, we first removed 431 articles that turned out to not be research articles. Most of these articles were either the title of a conference proceeding or book chapters. We, then, checked the remaining 1,823 articles for whether they actually covered both, AI and DLT. 1,544 turned out to not cover AI and DLT, the largest group among these consists of 141 articles from the medical field where DLT is an abbreviation for other terms such as *dose limiting toxicity* or *double lumen endobronchial tube*. From the remaining 279 articles, we excluded further 213 articles that did not cover the close integration of these technologies according to the concept of convergence [14], [15]. Out of these, the largest group of 92 articles covered AI-based cryptocurrency price prediction or trading. In the remaining set of 66 articles, we excluded another 34 articles that did not clearly answer why they used DLT. This resulted in a set of 32 articles that eventually turned out to be relevant for further analysis.

B. DATA ANALYSIS

Following the data collection, which resulted in the identification of 32 relevant articles, we categorized all 32 articles into groups. The groups were thereby derived from Dinh and Thai [13]’s viewpoint and adapted to our review, where necessary. Toward this end, we expanded the concept of blockchain to DLT, and reframed the focus and name of some categories to better suit the extant literature (e.g. *secure DLTs* instead of *secure and scalable DLTs*, and *coordination of devices* instead of *coordination of untrusting devices*). Table 1, below, provides an overview of the adapted coding scheme and the number of articles in each category. Note that we categorized some articles into multiple groups. Therefore, the sum of

TABLE 1. Classification of the identified articles that cover the integration of AI and DLT into groups.

Perspective	Group	# of articles
AI for DLT	Secure DLTs	7
	Automated referee and governance	3
	Privacy-preserving personalization	6
DLT for AI	Decentralized computing for AI	13
	Secure data sharing and marketplaces	14
	Explainable AI	6
	Coordinating devices	1
Both	Other	2

articles in Table 1 is higher than 32. We also added another type of perspective, *Both*, to the coding scheme. It consists of consolidating works, such as literature reviews or high-level articles covering both aspects of the convergence of AI and DLT. Several of these articles contain concepts from multiple groups, especially articles covering DLT for AI. In Table 6 in the appendix, we present an overview of the 32 analyzed articles and their classification into different groups. Our future research agenda is an extension from our review and draws from multiple sources. On the one hand, the findings, outlook and conclusion of the extant articles discussed in our review. On the other hand, our own assessment of further, recent developments in literature in the fields of AI and DLT.

IV. REVIEW ON AI FOR DLT

Drawing on the general distinction between AI for DLT and DLT for AI proposed by Dinh and Thai [13], this section describes our findings in terms of how extant research has applied AI for DLT. We identified three different groups of use contexts, which we detail on below, and summarize our findings in Table 2.

A. SECURE DLTS

1) DLT PROTOCOL SECURITY

Within this category, extant literature solely applies AI methods of reinforcement learning to explore and develop strategies in game-theoretic settings concerning the DLT protocol. Specifically, these articles analyze the fairness of mining activities. The learning agents get rewarded with a cryptocurrency-based miner reward. In the category of articles using reinforcement learning for DLT protocol security, two subcategories exist.

a: SELFISH MINING

Articles in this subcategory analyze selfish mining strategies for blockchain systems [10], [32]. The learning agent

TABLE 2. Overview of concepts in the literature on AI for DLT.

Group and categories	Summary of concepts	Exemplary sources
<i>Secure DLTs</i>		
DLT protocols	<ul style="list-style-type: none"> • Study of fairness of game-theoretic mining scenarios with reinforcement learning and implications on blockchain security • Some articles focus on selfish mining strategies • Other articles focus on mining pool operation 	<ul style="list-style-type: none"> • Hou, et al. [10], Wang, et al. [32], Haghighat and Shajari [35] • Hou, et al. [10], Wang, et al. [32] • Haghighat and Shajari [35]
Smart contracts	<ul style="list-style-type: none"> • Passive smart contract analysis using bytecode or metadata analysis • Active smart contract invalidation using reinforcement learning 	<ul style="list-style-type: none"> • Camino, et al. [36], Kim, et al. [37], Tann, et al. [38] • Zhang, et al. [39]
<i>Automated referee and governance</i>		
DLT protocols	<ul style="list-style-type: none"> • Automatically govern DLT protocol parameters or transactions 	<ul style="list-style-type: none"> • Lundbæk, et al. [40], Gladden [41]
Smart contracts	<ul style="list-style-type: none"> • Assist users in votes of changing SCs 	<ul style="list-style-type: none"> • Liu, et al. [42]
<i>Privacy-preserving personalization</i>		
Local computation	<ul style="list-style-type: none"> • DLT as an auditable data sharing infrastructure, computing ML model inference locally, federated learning for model updates 	<ul style="list-style-type: none"> • Zhao, et al. [43], Ramanan, et al. [44]
Trusted Execution Environments	<ul style="list-style-type: none"> • DLT as an auditable data sharing infrastructure, TEEs for confidential and computationally intensive SCs, ML training and inference in the SCs • Federated learning for increased data confidentiality • Differential privacy for increased data confidentiality 	<ul style="list-style-type: none"> • Hynes, et al. [11], Passerat-Palmbach, et al. [45] • Passerat-Palmbach, et al. [45] • Hynes, et al. [11]

thereby is a blockchain miner, which can strategically delay the publication of found blocks. By doing so, the selfish miner aims to waste mining power of honest miners that work on another fork of the blockchain [33]. Through generating these attacks on blockchain systems in a testing environment using reinforcement learning, researchers can find new insights on the security of these blockchain protocols. For example, Hou *et al.* [10] performed the simulation with protocols such as Bitcoin [7], Ethereum [28], or GHOST [34]. In their research, an agent learns mining strategies where it can gain disproportionately large mining rewards relative to its hash rate. While some of these adversarial mining strategies are known from theoretical studies of the blockchain protocols [33], AI has the potential to detect new, previously unknown attacks in complex scenarios, for example, in cases with multiple partially cooperating agents. Additionally, the simulation may contribute to a better understanding of the feasibility of these mining strategies [10].

b: POOL OPERATION

In research within this subcategory, the agent learns strategies for operating a mining pool. In our review, we identified only one article that fits this subcategory: Haghighat and Shajari [35] specifically analyze a block withholding game, where a pool operator can either decide to mine honestly, or attack another pool by mining (at least with a fraction of its power) for the other pool, but not actually submitting a solution, in case a block is found. As a consequence, the revenue of

the attacked pool drops, and it gets less attractive for other, honest miners to mine as a member of this attacked pool. The authors train an agent, which represents a pool operator, with reinforcement learning methods. It decides between mining honestly and attacking other pools at each step in the game. One insight from the results is that it may be more likely than widely assumed that one pool operator at a certain point in time was in control of more than 51 % of Bitcoin's network mining power [35].

2) SMART CONTRACT SECURITY

Articles within this category aim to protect users from interacting with insecure or malicious smart contracts. Two subcategories with different approaches exist in literature: First, the analysis and detection of such smart contracts. Second, the active interaction with such smart contracts in order to manipulate and invalidate them.

a: DETECTION

Three articles in this subcategory aim to detect security issues in smart contracts using data analysis techniques. As the smart contract opcode is usually published on the distributed ledger, two articles [37], [38] analyze the opcode with neural networks. Since a compiled opcode is challenging for humans to read and understand, Kim, *et al.* [37] use neural networks to estimate the functionality of a given smart contract (e.g., whether it is intended for a marketplace or for gaming). Tann, *et al.* [38], in contrast, aim at detecting security

vulnerabilities in smart contracts. To do so, they run long short-term memory neural networks on a sequence of opcodes. In a performance comparison with a symbolic analysis approach, the machine-learning-based analysis method can outperform the formal verification-based method with regards to classification accuracy and speed. Another article by Camino *et al.* [36] aims at detecting honeypot smart contracts. These types of smart contracts appear to contain free, withdrawable funds. However, once a user aims to invoke the honeypot contract to redeem free tokens, the honeypot actually does not release the funds, resulting in the user losing funds that they may have used to invoke the honeypot contract. Contrary to the other articles in this subcategory, Camino *et al.* [36] do not use opcode analysis techniques, but analyze the smart contract metadata (i.e., values derived from related transactions and fund flows) and off-chain data (i.e., the availability of a source code and the number of source code lines). In this way, data science techniques can classify more than 85 % of unseen smart contracts correctly. This approach especially works well in cases where there was a minimum of smart contract on-chain activity, and, as a result, on-chain metadata is available.

b: MANIPULATION

The only article identified within this subcategory [39] uses an active reinforcement learning approach to invalidate criminal smart contracts. Criminal smart contracts can be used for illegal activities, such as selling confidential information [46]. In this research, the agent learns to manipulate the contracts' data feed and thereby successfully invalidates a substantial share of the studied criminal smart contracts. While doing so, the system aims at invalidating these given smart contracts, but not at detecting possible criminal smart contracts in the first place.

B. AUTOMATED REFEREE AND GOVERNANCE

Another small group of articles uses AI for DLT-based system governance. Dinh and Thai [13] present the vision that people, devices, and smart contracts record transactions on a distributed ledger. An AI can then solve potential disputes of events happening on- or off-chain, and record the results on a distributed ledger. This automated arbitration could be data driven, unbiased, and, as a result, more consistent, justified, and accepted than arbitrations today [47]. In extant literature, articles toward this vision appear to be in an early stage and fall within two categories: AI-based DLT protocol governance, and AI-based smart contract governance.

1) DLT PROTOCOL GOVERNANCE

This category currently consists of two articles. Lundbæk *et al.* [40] propose an adjusted proof-of-work-based consensus mechanism. Only a subset of nodes participates in the required hash calculation and the DLT protocol uses ML to regularly update system governance parameters, such as the ideal number of miners or the level of mining difficulty. In this scenario, AI is tightly integrated with the DLT

protocol. However, the authors do not discuss in detail the functionality and security aspects of their ML-based governance. Gladden [41], on the other hand, discusses ethical aspects about cryptocurrencies governed by an AI system. The author claims that such a system can positively influence the ethos and values of societies. However, the article has a sociotechnical focus and does not describe the technical architecture of the system.

2) SMART CONTRACT GOVERNANCE

Only one article in our review fits this category. Liu *et al.* [42] propose to implement voting mechanisms for participants in smart contracts to alter smart contract parameters. As such, these voting mechanisms can govern smart contracts in complex situations and can alter the smart contracts towards an adaptive behavior. An ML-based system assists the users in voting, based on their past voting behavior. As a result, the ML-based system pre-chooses the users' selections and eases their tasks. In this scenario, the ML-based system runs outside of the actual DLT system.

C. PRIVACY-PRESERVING PERSONALIZATION

Many internet platforms nowadays collect data about their users and apply AI-based recommender systems to personalize their content. Examples include Facebook, Netflix, or Taobao in China. A recommendation by this AI-based system thereby is not only influenced by the individual user's data, but also by all the other users' data. Netflix, for example, can recommend new movies to users based on their individual and other users' watching behavior. However, this comes at the risk of impeding users' privacy. In several cases, private data from such platforms has been leaked to the public [48] or misused [49]. A group of articles envisions AI-based personalization for DLT-based data sharing platforms. This could, for example, include a social network built on a DLT infrastructure [13]. DLT can thereby serve as a transparent trail of data flows, and give users control over their data. Two categories of articles with different approaches toward designing such systems exist in literature. First, a category with articles that aim to use local computation. Second, a category with articles that use distributed ledgers based on hardware-assisted Trusted Execution Environments (TEEs).

1) LOCAL COMPUTATION

An ML model inference intended to personalize content on a platform requires data about the user to personalize its recommendation for them. If such a model inference is executed locally, there is no need for the user to share their data, while they can still get a personalized recommendation. However, if no user shares their data with the platform, it is challenging for a platform operator to train ML models using traditional methods. Therefore, articles in this category describe systems that use federated learning. With this distributed ML technique, an ML model gets trained locally on a user's device and only ML model parameter updates leave the device and are shared with the platform or potentially other

parties. Model parameter updates from many users are then aggregated into a new model. Federated learning is already successfully applied on a large scale on smartphones, for example, to predict the next word a user may want to type on their keyboard [50]. The distributed ledger provides an infrastructure to share data on an auditable and immutable ledger. Some articles describe systems that only store hashes of model gradient updates or hashes of aggregated models on the distributed ledger, in order to save storage space on the ledger and to preserve data confidentiality [51]. Other articles describe systems that store the complete gradient updates and aggregated models themselves on the ledger [52]. First applications are illustrated for the Internet of Things [43] or the taxi industry [44].

2) TRUSTED EXECUTION ENVIRONMENTS

The second category uses TEE-based DLTs. TEEs are typically located on a CPU and provide a special enclave for secure computation. In this enclave, applications are executed such that other applications executed on the same CPU but outside the enclave cannot interfere with the state or control flow of the application shielded by the TEE. As a result, the enclave appears as a black box to the outside and secures the application running inside. Furthermore, the hardware of the TEE can generate a proof to the application running inside the TEE to attest that the software is executing on this specific, trusted hardware. This feature is called remote attestation. Extant research uses DLT protocols to coordinate TEEs, thus enabling confidential smart contracts that execute inside the TEE. Since the TEE ensures data confidentiality, this category does not necessarily require local computation. First articles aim at providing solutions for the medical industry [11], [45]. To further increase data security, some articles describe systems that combine TEEs with federated learning [45] or differential privacy [11]. Simplified, differential privacy is a mechanism that adds randomly distributed noise to a set of data points. This protects the information privacy of users that provide data. As the random noise has zero mean (or a predefined mean), an aggregator can still draw meaningful conclusions from the aggregated data set without threatening anyone's information privacy.

V. FUTURE RESEARCH AGENDA ON AI FOR DLT

In this section, we present our analysis of future research opportunities on the advancement of DLT using AI. Again, we use the categories proposed by Dinh and Thai [13] to structure the future research agenda. We summarize our findings in Table 3.

A. SECURE DLTs

Dinh and Thai [13] have drawn a futuristic picture with far-reaching real-time analysis and decision possibilities of an AI as part of the DLT system. Current AI-based systems, however, do not provide the required security and robustness guarantees necessary to govern a DLT system. Precisely, while AI-based systems can detect software vulnerabilities,

they cannot (yet) guarantee that all available security vulnerabilities have been identified. If an AI detects no vulnerability, this does not necessarily mean that there are none [10]. This open research problem of AI robustness is also highly relevant in other AI application domains, such as autonomous driving, where first results may be transferable to software and DLT systems' security [53]. In general, methods from the field of explainable AI (XAI) [19] could help to better understand how reliable an AI-based decision is. We, therefore, see this as a fruitful field for future, foundational research. While AI is a promising technology to detect DLT security vulnerabilities, further research that aims to resolve these vulnerabilities and to build secure DLT-based systems is required in parallel [35]. Subsequently, we present future research opportunities in the two categories identified in our review in section IV. Furthermore, we see possible research opportunities in the convergence of these two categories, which we also present.

1) DLT PROTOCOL SECURITY

The security analysis of DLT protocols using AI is currently dominated by reinforcement learning. Extant literature analyzes game-theoretic incentive mechanisms in widely adopted protocols such as Bitcoin or Ethereum. Possible future research opportunities include the expansion of previous work with settings such as partially cooperating miners [10], or the analysis of strategies other than selfish mining [32]. In our view, beyond analyzing the fairness of mining settings and its DLT security implications, another additional interesting field of AI-based security analysis is the DLT source code itself. This is a highly relevant field in practice. For example, the deployed and widely-used cryptocurrency Zcash, which uses a novel form of zero knowledge proof cryptography, has been subject to a vulnerability that would have allowed an attacker to generate an infinite amount of cryptocurrency tokens [54], thus potentially inflating the assets of other users. Extant literature already uses AI-based analysis methods for software bug identification in general [55], and smart contract opcode analysis in particular [37], [38], thereby often outperforming formal verification-based methods [38]. We, therefore, see a promising future research avenue in further developing and applying such AI methods for analyzing DLT protocol code security.

2) SMART CONTRACT SECURITY

Extant literature in this subcategory mostly aims to protect users from insecure [37], [38] or malicious [36], [39] smart contracts. Practical experience, however, has shown that developers struggle to develop secure smart contracts in the first place [56]. Therefore, an interesting perspective for future research would be in the field of using AI to assist developers in developing secure smart contracts. This could include the development of early security warning tools for developers that automatically check developed code for security vulnerabilities. Recent work in other fields of software engineering suggests that such AI-based systems may be feasible [55], [57].

TABLE 3. Overview of future research opportunities in the field of AI for DLT.

Group and categories	Future research opportunities	Exemplary sources
<i>Secure DLTs</i>		
DLT protocols	<ul style="list-style-type: none"> • Advance reinforcement learning research for game-theoretic analysis • Use supervised learning methods for protocol software quality evaluation 	<ul style="list-style-type: none"> • Hou, et al. [10], Wang, et al. [32], Haghghat and Shajari [35] • Kim, et al. [37], Tann, et al. [38], Pradel and Sen [55]
Smart contracts	<ul style="list-style-type: none"> • Assist developers in developing secure SCs • Apply reinforcement learning for game-theoretic analysis of complex smart contracts • Refine supervised learning methods for SC security analysis (e.g., with regards to robustness) 	<ul style="list-style-type: none"> • Camino, et al. [36], Kim, et al. [37], Tann, et al. [38], Zhang, et al. [39] • Zhang, et al. [39] • Camino, et al. [36], Kim, et al. [37], Tann, et al. [38]
Interaction of DLT protocols and smart contracts	<ul style="list-style-type: none"> • Explore game-theoretic interplay between complex smart contracts and DLT protocols by means of reinforcement learning 	<ul style="list-style-type: none"> • Hou, et al. [10], Daian, et al. [58]
<i>Automated referee and governance</i>		
DLT protocols	<ul style="list-style-type: none"> • Advance AI robustness, evaluate system security 	<ul style="list-style-type: none"> • Lundbæk, et al. [40]
Smart contracts	<ul style="list-style-type: none"> • Advance AI robustness, extend systems and evaluate security 	<ul style="list-style-type: none"> • Liu, et al. [42]
<i>Privacy-preserving personalization</i>		
-	<ul style="list-style-type: none"> • Development of data and ML model sharing platforms • Application of practical data privacy technologies, such as TEEs combined with federated learning and differential privacy • Deployment and evaluation within trusting consortia 	<ul style="list-style-type: none"> • Hynes, et al. [11], Passerat-Palmbach, et al. [45], Zhang, et al. [59] • Hynes, et al. [11], Passerat-Palmbach, et al. [45] • Passerat-Palmbach, et al. [45]

Beyond this suggestion to expand the usage perspective of future research, we see further research opportunities in refining and applying the methods identified in our review. Reinforcement learning, which has already shown promising results for DLT protocol security analysis in simulation settings [10], [32], [35], appears as a promising method to analyze the security of game-theoretically complex smart contracts [39] in simulations. This could include, for example, decentralized token exchanges.

The other identified approach is to analyze security vulnerabilities in smart contracts by supervised learning. This appears as a promising field of AI application, as AI-based methods have outperformed classical, formal verification-based methods [38]. Future research could expand this work by considering further classes of smart contracts [36] and further details, such as different smart contract compiler versions [37]. Furthermore, future research that seeks to support software developers could aim at not only detecting security vulnerabilities in a smart contract, but also at localizing this vulnerability (e.g., by expressing the information which portions of the bytecode cause the vulnerabilities [38]) or even at providing suggestions to fix it [60], [61].

3) INTERACTION OF DLT PROTOCOLS WITH SMART CONTRACTS AND SECURITY IMPLICATIONS

In addition, we see a convergence of these two categories—AI for DLT protocol security, and AI for smart contract

security—as a promising avenue for future research. Recent research has investigated the extent to which game-theoretical aspects in DLT protocols and smart contracts influence each other and can lead to unfair conditions for regular users of the overall system [58]. For example, miners who may participate in the trading of assets on blockchain-based decentralized exchanges have the sovereignty to decide which transaction is included in a block and which one is not. This can incentivize other users to pay higher transaction costs for these exchange transactions than they would without the miner participation. Prior research considers this as an unfair setting [58]. The outlined scenario is only one example of the complex interplay of smart contracts and DLT protocol incentives, which is also relevant in decentralized applications other than token exchanges [62], [63]. As it appears to be an even more complex topic than the standalone security of DLT protocols or smart contracts, an AI-based system could provide valuable insights. Prior research uses reinforcement learning for game-theoretic analyses of DLT protocols and smart contracts separately, therefore, reinforcement learning also appears as a promising method to analyze the interplay between both. The results could be used to support the development of fair and secure DLT-based systems.

B. AUTOMATED REFEREE AND GOVERNANCE

In general, both identified categories—AI for smart contract and for DLT protocol governance—face the reality of

current AI's capabilities. To automatically make far-reaching decisions on a distributed ledger with a potential influence on financial or data confidentiality circumstances, most of today's AI's robustness and explainability guarantees are not strong enough. The robustness guarantees of many modern AI-based systems are weak [64] compared to the complexity of the actions that agents perform when interacting with DLT-based systems. Therefore, in our view, breakthroughs in the robustness, explainability, and ultimately security of AI-based systems are required before they can automatically govern DLT protocols and smart contracts on a large scale. However, in the case of human-in-the-loop based AI smart contract governance [42], an AI may assist a human decision maker which can intervene at any time in case an AI-based system faces and detects an irregular situation. This human-in-the-loop model, therefore, appears to provide good practicality with today's AI technology and appears to be a good starting point for future research [42].

C. PRIVACY-PRESERVING PERSONALIZATION

DLT-based platforms for data sharing are starting to get more attention in research but are not deployed and scaled in practice yet. As such, research on such platforms and their AI-based personalization is still in an early stage. Articles in the group of privacy-preserving personalization using AI for DLT often incorporate aspects from multiple other groups, such as *decentralized computing for AI*, or *secure data sharing and marketplaces for AI*. In our view, developments in the field of *privacy-preserving personalization* depend on advancements in the underlying secure computation and privacy-preserving data sharing technologies.

An interesting future research opportunity which is practical with today's technologies is the further development and deployment of DLT-based data sharing platforms and an AI deployment on top of these platforms. TEEs appear as a promising technology to realistically build such systems today, due to their relatively small computational overhead when compared with other secure computation technologies, as well as their ability to enforce policies [65]. TEEs can be combined with federated learning [45] or differential privacy techniques [11] for stronger data confidentiality guarantees. A reasonable avenue for deployment appears to be within consortia of (at least partially) trusting participants. This could, for example, include a consortium of hospitals aiming to deploy personalized, AI-based treatments while complying with strong privacy requirements [45].

VI. REVIEW ON DLT FOR AI

Drawing, again, on the general distinction between AI for DLT and DLT for AI proposed by Dinh and Thai [13], this section describes our findings in terms of how extant research has applied DLT for AI. We identified four different groups of use contexts and summarize our findings in Table 4.

A. DECENTRALIZED COMPUTING FOR AI

Due to the large amount of processed data, ML model training is often computationally expensive. Nowadays, graphics processing units (GPUs) are typically used to train ML models, as they provide more computational power for most ML training tasks. For complex optimization tasks, many GPUs may be necessary to train ML models in a reasonable amount of time [66]. However, many GPUs and CPUs in computers around the world are only slightly loaded or even unused. Previously, distributed computing has been applied to utilize these unused resources in communities at scale, for example, to perform protein folding calculations for medical research [67]. Articles within this group describe systems that use distributed computation methods to train AI models. DLT can serve as a trail to organize this decentralized computation, as well as to provide a ledger for rewards paid in cryptocurrency [13]. Looking at the literature, three different approaches for DLT-enabled distributed computing for AI exist: First, the computation within the DLT protocol. Second, the computation in smart contracts. Third, the computation outside the distributed ledger.

1) DLT PROTOCOL

Especially public, permissionless distributed ledgers, such as the Bitcoin or Ethereum main network, are nowadays secured by a proof-of-work-based consensus mechanism. It refers to the search for a number such that the hash of a blockchain block candidate, which includes that number, is below a certain threshold. This mechanism is required to limit the number of nodes that can propose a new block, and ultimately, to ensure blockchain security. By adjusting the threshold, the network can adjust the difficulty of this search problem, and adjust the average time to find a new block. The search problem is computationally expensive. However, the verification of a potential solution is computationally inexpensive.

The proof-of-work mechanism consumes large amounts of energy for protocols such as Bitcoin, which some authors see critical [68]. At the same time, the calculated hashes cannot be meaningfully used for purposes other than securing the network. Some authors therefore propose a proof-of-useful-work mechanism, in which miners train an ML model instead of finding a number for a certain block hash. The block which contains an ML model with the least test error is then accepted as the new block by the other nodes [69]. For such a system to work, several challenges have to be overcome. On the one hand, the computational difficulty of optimizing an ML model is hard to adjust [70]. This can also cause blockchain forks to occur on a very regular basis. On the other hand, the proof-of-useful-work mechanism for AI requires the secure provision of an ML model architecture, training data, and test data. Extant literature that describes such systems uses selected groups of participants and committees with reputation mechanisms to govern this data provision [69]–[71].

TABLE 4. Overview of concepts in the literature on DLT for AI.

Group and categories	Summary of concepts	Exemplary sources
<i>Decentralized computing for AI</i>		
DLT protocol	<ul style="list-style-type: none"> Mechanism for block candidate selection, training AI models instead of calculating hashes, separate group of task providers, reputation mechanisms to secure the system 	<ul style="list-style-type: none"> Truby [68], Chenli, et al. [69], Baldominos and Saez [71]
Smart contracts	<ul style="list-style-type: none"> Scaling smart contract capabilities (e.g., through TEEs), training AI models in these smart contracts while ensuring the computational integrity 	<ul style="list-style-type: none"> Hynes, et al. [11], Passerat-Palmbach, et al. [45], Zhang, et al. [59], Swan [72]
Off chain	<ul style="list-style-type: none"> Training ML models off-chain through federated learning Application of game-theoretic incentive mechanisms to ensure computational integrity 	<ul style="list-style-type: none"> Zhao, et al. [43], Ramanan, et al. [44], Preuveneers, et al. [52], Shayan, et al. [73] Lu, et al. [74]
<i>Secure data sharing and marketplaces for AI</i>		
Smart contracts	<ul style="list-style-type: none"> Manage access to data on decentralized file protocols through DLT Financial rewards and staking mechanisms to incentivize honest behavior 	<ul style="list-style-type: none"> Özyilmaz, et al. [75] Harris and Waggoner [76]
TEEs	<ul style="list-style-type: none"> Execute computationally intensive SCs in a TEE Manage a large share of the ML pipeline through the system 	<ul style="list-style-type: none"> Hynes, et al. [11], Passerat-Palmbach, et al. [45], Zhang, et al. [59] Hynes, et al. [11], Zhang, et al. [59]
Federated learning	<ul style="list-style-type: none"> Store model update plain data or hashes on the ledger, increase federated learning transparency and accountability 	<ul style="list-style-type: none"> Zhao, et al. [43], Ramanan, et al. [44], Preuveneers, et al. [52], Shayan, et al. [73]
<i>Explainable AI</i>		
-	<ul style="list-style-type: none"> DLT as a trail for model training and inference (e.g., for federated learning) TEEs to enforce data usage policies 	<ul style="list-style-type: none"> Hynes, et al. [11], Passerat-Palmbach, et al. [45], Sarpatwar, et al. [51], Preuveneers, et al. [52], Zhang, et al. [59] Hynes, et al. [11], Zhang, et al. [59]
<i>Coordination of devices</i>		
-	<ul style="list-style-type: none"> Distributed ledger for device coordination (e.g., through public keys and reputation management) Remote attestation of a TEE 	<ul style="list-style-type: none"> Kang, et al. [77] Cheng, et al. [78]

2) COMPUTATION IN SMART CONTRACTS

In some DLT networks, such as Ethereum, a virtual machine can execute Turing-complete smart contracts. As such, DLT could provide a substrate for ML model training [72]. However, traditional ledgers like Ethereum do not support intensive computations that would be necessary to perform ML model training [74]. To overcome this limitation, some articles propose to extend distributed ledger smart contract execution with TEEs such that these ledgers can support computationally more intensive smart contracts. This TEE can then train simple ML models [11], [59].

3) OFF-CHAIN COMPUTATION

Another subcategory of articles performs computations not on the distributed ledger, but off-chain. Unless using sophisticated cryptographic techniques [79], the integrity of the computation cannot be guaranteed. DLT-based federated learning systems, for example, train an ML model on the

edge [43], [44], [52], [73]. The main motivation of most articles focusing on DLT-based federated learning, however, is not the usage of decentralized edge computing resources per se. Instead, the main motivation is secure data sharing. Some of the DLT-based federated learning articles also propose financial rewards for both in combination, model training and data sharing [80], [81]. Besides federated learning, another article by Lu *et al.* [74] proposes a crowdsourced computation approach to offload heavy computation tasks from a blockchain, such as in ML model training. Multiple offloaded computing engines audit each other's work and game-theoretic incentive mechanisms are used to build a protocol. The authors also present a security analysis of their protocol.

B. SECURE DATA SHARING AND MARKETPLACES FOR AI

In addition to the increasingly available computing power, another fundamental reason for the recent advancement in

AI is the strong growth of available and digitized data. ML-based systems generally perform better the more data they are trained on, for example, with regards to classification accuracy [21], [82]. Some authors see recent developments in IT, in which few companies are in control of large amounts of personal data, critical and propose DLT-based data markets to democratize such data silos [13], [83]. As such, a group of articles proposes solutions based on DLT to build data sharing infrastructures, thus enabling the deployment of AI. These articles differ in their technical approaches toward designing such systems. Articles which incorporate advanced privacy-preserving mechanisms in their systems mostly focus on the health care industry. We suspect the reasons for this being strong confidentiality requirements for the handling of health data.

1) SMART CONTRACTS

Two identified articles use traditional smart contracts on Ethereum to build a data sharing infrastructure and marketplace. These articles differ in their focus, either on sharing training data itself [75], or ML models [76]. Especially the sharing of training data requires the storage and handling of large files. Therefore, Özyilmaz *et al.* [75] connect their system to the decentralized file protocol SWARM and manage data access rights through the blockchain. These marketplaces allow the payment of data providers with units of the cryptocurrency Ether. To incentivize participation with high quality data, Harris and Waggoner [76] propose staking mechanisms in which malicious participants sharing spam models lose their stake. Such systems can be applied in the Internet of Things industry in general [75].

2) TRUSTED EXECUTION ENVIRONMENTS

TEEs enable the computation of relatively intensive tasks, while preserving data confidentiality and integrity throughout the computation. Therefore, systems described in some articles execute computationally intensive smart contracts off-chain in a TEE [11], [45], [59], [84], [85]. These articles use the blockchain concept with different DLT designs [3]. Several articles use the Oasis blockchain [11], [84], [85] Hynes *et al.* [11], for example, present a privacy-preserving data market that provides solutions for a large share of the ML pipeline. Through smart contracts, data providers can define policies to share their data. These policies, for example, include the asking for a reward and differential privacy requirements. Data consumers can choose to fulfill these policies in order to train an ML model on the providers' data. Since a TEE ensures confidential computation, the training data does not get leaked and only data consumers can get access to ML model inference. The ML model itself is shielded inside a smart contract and inference executions count toward the provider's policies, which increases data provider's privacy against potential inference attacks. This type of attacks aims at executing the ML model in order to extract the underlying training data or the model itself [86].

Other articles, which describe a system with similar capabilities, use an Ethereum-based blockchain [45], [59]. As a special feature, a virtual machine in the enclave can allow the training of proprietary AI models while not comprising the training data security [59]. Passerat-Palmbach, *et al.* [45] incorporate federated learning into the TEE-enabled blockchain. This protects data providers from potential TEE side channel attacks and, therefore, further reduces data privacy risks.

3) FEDERATED LEARNING

A large subcategory of articles describes systems that use federated learning with distributed ledgers that are not TEE-enabled [43], [44], [51], [52], [81]. DLT thereby serves as a provenance record of data. This data can describe a large share of the ML pipeline —training data origin, training data, ML model modifications, or testing data [51]. In most systems described in these articles, the data itself is not stored on the blockchain, but hashes of the data [51]. In some cases, the systems use relatively simple ML models and store plain model updates on the blockchain [52]. As a result, every participating server can audit and compute the aggregated ML model weight updates. A blockchain-based solution can be well-suited for certain use cases and only provides a small performance overhead of ca. 5 % to 15 %, while enabling transparency and accountability [52]. Some systems go further and replace the traditionally centralized aggregator with a smart contract-based one [44].

As federated learning is potentially vulnerable to inference attacks [86], some articles describe systems that further use differential privacy techniques [43], [73] to increase data confidentiality guarantees. Further systems incorporate financial incentives for participants sharing model updates [43], [81].

C. EXPLAINABLE AI

Complex ML models, such as deep neural networks, are nowadays often used in a black box manner [19]. This means that users or even system creators do not have the information how these models come to a certain prediction. However, obtaining this information can be desirable in certain cases, for example, to verify the system's robustness or to comply with legislation [19]. Dinh and Thai [13] outline DLT as a technology to increase AI-based systems' explainability. Their vision is that DLT provides an immutable trail to track the data flow of AI-based systems.

Looking into the extant literature on DLT for AI, the explainability of the black-box model itself is out of focus for most researchers. The DLT for XAI literature mainly covers data provenance or computational integrity aspects for model training or inference. Sarpatwar *et al.* [51] design a DLT-based federated learning system for trusted AI and present five requirements of blockchain for trusted AI. First, guarantees of AI model ownership and track of use are important. Second, the confidential sharing of AI assets as they are often created using sensitive data. Third, auditability regarding the AI training process. Fourth, the traceability of

the AI pipeline. Fifth, to keep a record in order to recognize a potential bias and model fairness issues.

In general, several extant articles aim to use DLT-based federated learning to ensure such aspects of AI model explainability and trustworthiness [51], [52], [76]. Other articles use TEEs to also cover the model inference with their system, and thus, provide stronger explainability and restrictions in order to track data flows for data providers [11], [45], [59]. Hynes *et al.* [11], for example, use XAI methods to compensate data providers relative to their influence their data had on an ML model's inference.

D. COORDINATION OF DEVICES

A final group of articles aims to coordinate devices with DLT. These devices, such as Internet of Things devices, generate data, which AI-based systems can analyze. A distributed ledger can serve as a registrar for these devices, and store different types of data that these devices generate. This includes, for example, metadata (such as hashes), or the generated data itself. In extant literature, articles that use DLT also connect devices to the distributed ledger. Using cryptographic principles, transactions on a distributed ledger (such as smart contracts) are signed, which ensures that only legitimate participants can perform certain transactions. Kang *et al.* [77] propose to use DLT for reputation management in a federated learning setting. Beyond identification through asymmetric cryptographic mechanisms on a distributed ledger, TEEs provide mechanisms to self-identify and remote-attest [78], such as physically unclonable functions [87].

VII. FUTURE RESEARCH AGENDA ON DLT FOR AI

In this section, we present our analysis on open, future research fields on DLT for AI. Again, we use the categories based on Dinh and Thai [13] to structure the future research agenda. We present an overview of our results in Table 5.

A. DECENTRALIZED COMPUTING FOR AI

As we identified three categories of approaches using DLT for decentralized computing for AI, the future research opportunities in these approaches differ. Subsequently, we present our identified research opportunities in all of these.

1) DLT PROTOCOL

First, extant research presents blockchain designs with proof-of-useful-work mechanisms that could be used for ML model training. Proof-of-useful-work itself has a long history with several blockchain-based cryptocurrency systems deployed [88], for example, for number theoretic research [89]. While this prior work is interesting from an academic point of view, the deployed cryptocurrencies so far had little practical relevance based on their market capitalization. Some blockchain designs, such as the public Ethereum mainnet, actually aim to not use consensus mechanisms based on proof-of-work in the future [90].

Despite that, we see promising avenues for future research on proof-of-useful-work for AI, as well as for actually

deploying and evaluating such systems. On the one hand, future research could analyze the economics of such systems and their utility for ML model requesters as well as regular DLT users. On the other hand, future research could analyze the practical security guarantees of such systems. Aspects include the security of incentive mechanisms, as well as the security of DLT transactions. This is particularly interesting, since an attacker with large computational power would not only benefit from rearranging transactions on the ledger, but also potentially from the useful work itself.

2) COMPUTATION IN SMART CONTRACTS

Extant research has aimed to compute smart contracts off-chain in TEEs. Out of several approaches to off-chain computation while ensuring integrity [79], TEEs are interesting due to the relatively high computational power they provide [91]. TEEs have seen prior research from both fields separately, AI and DLT. In the field of AI, extant research provides frameworks for ML model training and inference on TEEs [92]. It has shown that TEEs are in general capable of simple practical machine learning tasks, such as speech processing [93], or the processing of small images [91]. In the field of DLT, researchers have sought to increase the low performance of smart contracts using TEEs while ensuring confidentiality [78].

From our point of view, TEEs generally provide a promising trade-off between computational performance and confidentiality. Prior research that provides an ML pipeline with TEE-enabled DLTs does not evaluate its performance [11], [45], [59]. As such, a natural extension would be a practical performance evaluation. This could help researchers aiming to deploy such systems, for example, for privacy-preserving personalization. By knowing what amounts of data and what complexity of computation the system can handle, other researchers could select use cases and deploy and evaluate TEE-enabled DLTs for their use cases.

Another avenue for future research is the further development of computationally powerful TEEs, and to potentially even enable TEEs on GPUs [106], which are, for many ML model training tasks, better suited than CPUs. At the same time, research has identified security vulnerabilities in TEEs and potential security measures [94]. Therefore, future research could aim to develop secure TEEs and mechanisms to prevent attacks on TEEs, as well as study the feasibility and implications of TEEs on distributed ledgers. To prevent TEE security vulnerabilities, some researchers combine them with other privacy-enabling technologies, such as federated learning [45] or differential privacy [11], [91]. The further analysis of such combined methods with regards to data confidentiality guarantees, computational overhead, impact on machine learning quality, and practicality for use cases would be another promising avenue for further research.

In addition, TEEs with a strong integration into a blockchain protocol enable new mechanisms for block selection. These mechanisms use special functions of TEEs, such as the secure generation of random numbers, and are known

TABLE 5. Overview of future research opportunities in the field of DLT for AI.

Group and categories	Future research opportunities	Exemplary sources
<i>Decentralized computing for AI</i>		
DLT protocols	<ul style="list-style-type: none"> • Further develop, deploy, and evaluate systems, analyze incentive mechanisms, analyze system security 	<ul style="list-style-type: none"> • Truby [68], Chenli, et al. [69]
Smart contracts	<ul style="list-style-type: none"> • Develop systems enabling scalable SCs through TEEs • Improve TEE performance and security • Evaluate suitability of systems for practical use cases • Incorporate privacy mechanisms into computing, such as federated learning or differential privacy 	<ul style="list-style-type: none"> • Passerat-Palmbach, et al. [45], Zhang, et al. [59], Hynes, et al. [91] • Lee, et al. [94], Brasser, et al. [95] • Hynes, et al. [11], Passerat-Palmbach, et al. [45], Hynes, et al. [91], Kunkel, et al. [92], Milutinovic, et al. [96] • Passerat-Palmbach, et al. [45], Chen, et al. [97]
Off-chain computation	<ul style="list-style-type: none"> • Deploy DLT-based, federated learning especially within trusted consortia • Further research on game-theoretic incentives for secure computation • Foundational research on cryptographic technologies and DLT scalability in general 	<ul style="list-style-type: none"> • Passerat-Palmbach, et al. [45] • Lu, et al. [74] • Miller, et al. [98], Poon and Buterin [99], Dziembowski, et al. [100], Kalodner, et al. [101], Teutsch and Reitwießner [102]
<i>Secure data sharing and marketplace for AI</i>		
Smart contracts	<ul style="list-style-type: none"> • Increase system scalability and data confidentiality • Deployment and evaluation in practice, especially in permissioned ledgers within trusted consortia 	<ul style="list-style-type: none"> • Özyilmaz, et al. [75], Harris and Waggoner [76] • Passerat-Palmbach, et al. [45]
TEEs	<ul style="list-style-type: none"> • Deployment and evaluation of systems in practice • Security analysis of TEEs, and development of secure TEEs • Incorporation of incentive mechanisms (e.g., staking) 	<ul style="list-style-type: none"> • Hynes, et al. [11], Passerat-Palmbach, et al. [45] • Lee, et al. [94], Brasser, et al. [95] • Harris and Waggoner [76]
Federated learning	<ul style="list-style-type: none"> • Deployment and evaluation in trusted consortia • Security analysis and combination with other privacy techniques • Analysis of incentive mechanisms (e.g., staking) 	<ul style="list-style-type: none"> • Passerat-Palmbach, et al. [45] • Preuveneers, et al. [52], Geyer, et al. [103] • Harris and Waggoner [76]
<i>Explainable AI</i>		
-	<ul style="list-style-type: none"> • Deploy systems that use DLT as a trail for XAI and trustworthy AI • Evaluate the practical XAI benefits of such systems 	<ul style="list-style-type: none"> • Hynes, et al. [11], Sarpatwar, et al. [51], Preuveneers, et al. [52], Zhang, et al. [59], High-Level Expert Group on Artificial Intelligence [104] • Hynes, et al. [11], Passerat-Palmbach, et al. [45], Preuveneers, et al. [52]
<i>Coordination of devices</i>		
-	<ul style="list-style-type: none"> • Reputation management of DLT participants • Self-identifying and secure devices 	<ul style="list-style-type: none"> • Kang, et al. [77] • Cheng, et al. [78], Schulz, et al. [105]

as proof of luck [96] or proof-of-elapsed-time [97]. Future research could aim to further integrate such mechanisms into the DLT protocol and evaluate the practical benefits and system security implications.

3) OFF-CHAIN COMPUTATION

Most presented DLT-based federated learning protocols do not ensure the integrity of the model training calculations. As such, we see future research potential for such systems in at least partially trusting consortia. Federated learning is already successfully applied in systems with a large number of users [50], therefore, future applied research could provide further promising results when combined with DLT.

A further avenue for future research is the use of game-theoretic mechanisms to incentivize honest computation, for example, for ML model training [74]. In general, there is a lot of current research that aims to scale smart contract executions outside a distributed ledger while maintaining integrity [98]–[102]. Much of this research is at the conceptual level and does not yet deal in detail with the application of concrete computationally intensive applications, such as the training of ML models.

In this regard, the advancement of cryptographic technologies for secure computation is another avenue for future research. Homomorphic encryption, for example, enables data consumers to perform computations on encrypted data.

However, this technology and other technologies (such as secure multiparty computation or zero knowledge proofs) are not yet practicable for ML model computations [11], [78], [84], [91] due to the strong computational overhead. If future research achieves breakthroughs to enable intensive computations with these cryptographic technologies, they would provide an alternative to TEEs.

B. SECURE DATA SHARING AND MARKETPLACES FOR AI

Articles within this field already present systems that could be further developed and deployed for practical use. Future research could, therefore, evaluate the practicality and user acceptance of such systems. For this practical deployment, health care use cases may be particularly suitable, because the deployment of ML in this context is especially challenging due to high privacy and security requirements [11], [84], [85].

Beyond the deployment and evaluation, a further avenue for future research is the rigorous security analysis of such systems against different types of adversaries. Potential adversaries could, for example, aim to get rewards for sharing spam data, aim to extract training data by training a proprietary ML model, or seek to extract training data of a shared ML model through inference attacks or the shared ML model itself. Only some articles of in our body of literature study such attacks and security measures to protect the system [11], [52], [59]. Especially staking mechanisms have seen little attention from researchers so far [76]. Future research could possibly transfer insights from proof-of-stake-based DLT consensus protocols [107] toward staking in DLT-based data marketplaces.

We see differential privacy as a promising technique for ensuring users' privacy [43], [73], especially in large datasets with low dimensionality. In such cases, its negative effect on the aggregated model quality is small while improving the individual data sharer's privacy. Previous research on differential privacy for ML studies its theoretical implications on test data sets [16]. In our view, future research could build on these and study differential privacy's practicality in certain vertical use cases for both, classic ML (relevant in the context of DLT; e.g., through TEEs), and federated learning. Federated learning, for example, can be vulnerable to inference attacks [86]. Differential privacy can help impeding the practical feasibility of such attacks [103].

Articles describing DLT-based federated learning use DLT as an immutable trail and, in some cases, even for the data communication and storage. Furthermore, DLT serves as a ledger for the reward payment using cryptocurrencies. We, therefore, see plain federated learning approaches as particularly promising for applications in (at least partially) trusting consortia. TEE-based approaches, however, not only use DLT for the same reasons, but go further. Smart contracts on a distributed ledger can enforce policies for both, data providers and data consumers. From this point of view, the TEE-based approach is more powerful in terms of covering a large share of the ML pipeline, and ultimately, in ensuring system

security. Accordingly, it may have better utility for future research that aims to deploy DLT-based secure data sharing and marketplace systems.

Nevertheless, TEE implementations have a rich history of security vulnerabilities [108]. A detailed security analysis of the practical security of TEEs and their implications for TEE-based DLT systems, possibly with further security mechanisms in smart contracts could provide a valuable contribution to future research. Such an analysis could also include combined security measures in smart contracts, like federated learning [45], differential privacy [11], or secure multiparty computation [11]. We see such combinations of privacy techniques [65], [84] with TEEs as an interesting avenue for future research, although some of these cryptography-based privacy techniques need breakthroughs in practicability to be usable for ML [84].

C. EXPLAINABLE AI

In our view, DLT is well-suited as a trail for AI model metadata and hashes of data generated around the AI model training and inference phase. As such, DLT can help to enable XAI applications that aim to describe the black box behavior of AI models for both, off-chain, as well as on-chain [11] applications. This can increase an AI-based system's security and, ultimately, enable the deployment of trustworthy AI [51], [104].

In many DLT-based federated learning articles, only AI model hashes and metadata [51] or relatively small AI models [52] are stored on the distributed ledger. TEE-based systems, on the other hand, can securely write and read external memory while preserving data confidentiality [59]. Thus, we see TEE-based systems [11], [45], [59] as specifically suited for large, data intensive ML environments. Future research could aim to deploy such systems and evaluate the practical explainability of the AI model.

D. COORDINATION OF DEVICES

One of the core functions of DLT is to provide an immutable ledger where only cryptographically legitimate participants can perform transactions. The usage of DLT for coordinating devices and participants is therefore essential and basic at the same time. Reputation management of participants on a distributed ledger is one field that has only started to see attention and appears to be of potential for future research. Kang *et al.* [77] specifically mention a dynamic system with variable thresholds as a future research opportunity. With regard to self-identifying TEEs, prior research already uses DLT to coordinate devices [78], [109]. We consider the further dissemination and use of self-identifying functionalities through physically unclonable functions [105] as an interesting future research opportunity. Such self-identification functionalities could potentially be built in Internet of Things devices, such as sensors or actuators. However, mechanisms for hardware-based authentication have been identified as insecure in the past [110], [111]. Therefore, we see an analysis of their practical security guarantees as another open research problem.

VIII. DISCUSSION

A. PRINCIPAL FINDINGS

Both fields, AI and DLT, currently experience a lot of hype within research and practice. Even though our initial database query listed a few articles that did not clearly explain what AI or especially what DLT was used for, a substantial number of extant articles provided profound insights for our review and future research agenda.

One analysis aspect in our work is that we do not limit our work to blockchain as the only DLT concept, but also consider other concepts such as directed acyclic graphs. This decision was based on the expectation that other DLT concepts with distinct characteristics might be better suited for some AI applications than the concept of blockchain [3], [12]. Yet, only one out of 32 articles within our review considered DLT concepts other than blockchain in the context of AI [75]. This result is in line with previous research, which has called for more research on other DLT concepts [3]. Improving and deploying DLT concepts other than blockchain is, therefore, a highly interesting and relevant avenue for future research.

Overall, the framework provided by Dinh and Thai [13] served as a helpful tool to classify extant literature with regard to the convergence of AI and DLT. However, several articles cover aspects from multiple groups in the subsections of AI for DLT or DLT for AI. This is particularly the case for articles in our review that cover multiple subsections of DLT for AI or the subsection privacy-preserving personalization from the section AI for DLT. We have slightly modified the framework of Dinh and Thai [13] in two ways. First, extant literature applies AI for DLT in the context of secure distributed ledgers, but not for scalable distributed ledgers as well. A possible reason for this is the lack of AI's robustness and security guarantees, as discussed in our future research agenda. Second, we renamed the subsection *coordination of untrusting devices* to *coordination of devices*, because DLT is not the element that establishes trust in all of the articles. For example, TEEs can use physically unclonable functions for remote attestation [87].

In our view, some of the research fields are mature enough to transfer systems into practice and evaluate their influence and user acceptance. This includes, for example, the fields of security analysis of smart contracts or marketplace systems based on TEE-enabled DLTs. Research particularly focus on the application of these marketplaces in the health care industry [11], [45], [59], [84], possibly due to the presence of data lakes and strong data confidentiality requirements. Therefore, further research-based deployment and evaluation of DLT-based marketplaces for AI in health care settings may deliver a promising contribution to their real-world deployment. At the same time, other research fields, such as AI-based automated referee and governance for DLT protocols, appear to require substantial progress in fundamental research (e.g., robust and secure AI), before transferring scientific knowledge into practice and the establishing of real-world systems.

By applying convergence as a theoretical lens on extant literature, we were able to focus our research on innovative articles that closely integrate AI and DLT. Furthermore, we were able to exclude research that does not closely integrate AI and DLT, such as AI-based cryptocurrency price prediction or trading. Drawing on the definition of convergence [30], many articles on AI and DLT's convergence fall into the first phase with cross-scientific research on their integration. Some articles already pave the way for the second phase with new platforms arising that could, for example, accelerate health care research [11], [45], [84]. During our review, we noticed that the concept of convergence has received relatively little attention by IT researchers in the past. This came to our surprise, as convergence has been a main driver of IT innovations over the recent years [14]. Therefore, we consider convergence as a promising theoretical lens to explore interdisciplinary technological settings.

B. LIMITATIONS

Both fields, AI and DLT, are moving very fast and breakthroughs are regularly achieved. In this respect, we cannot rule out the possibility that in some subfields of the convergence of AI and DLT, future research will achieve innovative breakthroughs that may enable use cases not identified in our review or future research agenda. We have taken several steps to minimize the chances of this outcome. First, our analysis includes ArXiv preprints of research articles that may only soon be published in journals or presented at conferences. Second, we included insights from foundational research on the topics of AI, DLT, and secure computation into our future research agenda. In doing so, we sought to consider aspects that may not be covered in the reviewed body of literature but are nevertheless highly relevant in the individual fields (e.g., the practical security guarantees of TEEs). Third, we have also included articles that cover little researched technologies, often with little practicality today, but that may see breakthroughs in the future. This includes Artificial General Intelligence, DLT which is not based on the concept of blockchain, and cryptographic protocols for computationally intensive tasks.

IX. CONCLUSION

In this research, we investigated the current research status and future research opportunities on the convergence of AI and DLT. In order to assess the current state of convergence, we conducted a systematic literature review and analyzed extant literature through the lens of convergence. Our findings include several different ways that describe how AI can advance DLT applications, as well as how DLT can advance AI applications. In order to develop a future research agenda, we built on the structure of our literature review and linked the ongoing research with other research in the separate fields of AI and DLT, as well as our own view on future research opportunities. Our results reveal multiple future research opportunities in this interdisciplinary field for both, theory- as well as practice-oriented research.

TABLE 6. Overview of relevant and analyzed papers and their classification into different groups.

	Title	AI for DLT perspective			DLT for AI perspective				Other
		Secure DLTs	Automated referee and governance	Privacy-preserving personalization	Decentralized computing for AI	Secure data sharing and marketplaces	Explainable AI	Coordinating devices	
Baldominos and Saez [71]	Coin.AI: A Proof-of-Useful-Work Scheme for Blockchain-based Distributed Deep Learning				●				
Camino, et al. [36]	A Data Science Approach for HoneyPot Detection in Ethereum (<i>arXiv version 2</i>)	●							
Chen, et al. [80]	When Machine Learning Meets Blockchain: A Decentralized, Privacy-preserving and Secure Design				●				
Chenli, et al. [69]	Energy-recycling Blockchain with Proof-of-Deep-Learning				●				
Dinh and Thai [13]	AI and Blockchain: A Disruptive Integration								●
Gladden [41]	Cryptocurrency with a Conscience: Using Artificial Intelligence to Develop Money that Advances Human Ethical Values		●						
Haghighat and Shajari [35]	Block withholding game among bitcoin mining pools	●							
Harris and Waggoner [76]	Decentralized & Collaborative AI on Blockchain					●	●		
Hou, et al. [10]	SquirRL: Automating Attack Discovery on Blockchain Incentive Mechanisms with Deep Reinforcement Learning (<i>arXiv version 1</i>)	●							
Hynes, et al. [11]	A demonstration of sterling: A privacy-preserving data marketplace			●	●	●	●		
Johnson, et al. [85]	Building a Secure Biomedical Data Sharing Decentralized App: Tutorial					●			
Jones, et al. [84]	Privacy-Preserving Methods for Feature Engineering Using Blockchain: Review, Evaluation, and Proof of Concept					●			
Kim, et al. [81]	Blockchained On-Device Federated Learning				●	●			
Kim, et al. [37]	ScanAT: Identification of Bytecode-Only Smart Contracts with Multiple Attribute Tags	●							
Liu, et al. [42]	Strengthening Smart Contracts to Handle Unexpected Situations		●						
Lu, et al. [74]	On Enabling Machine Learning Tasks atop Public Blockchains: A Crowdsourcing Approach				●				
Lundbæk, et al. [40]	Proof of Kernel Work: A democratic low-energy consensus for distributed access-control protocols		●						
Montes and Goertzel [83]	Distributed, decentralized, and democratized artificial intelligence					●			

With our article, we contribute to the current state of research in four ways. First, we expand prior research, which did not consider DLT concepts other than blockchain in the

integration with AI. Second, we consider both perspectives, AI for DLT, and DLT for AI and the many different concepts of their integration. Third, we bridge the gap

TABLE 6. (Continued) Overview of relevant and analyzed papers and their classification into different groups.

	Title	AI for DLT perspective			DLT for AI perspective				Other
		Secure DLTs	Automated referee and governance	Privacy-preserving personalization	Decentralized computing for AI	Secure data sharing and marketplaces	Explainable AI	Coordinating devices	
Özyilmaz, et al. [75]	IDMoB: IoT Data Marketplace on Blockchain					●			
Passerat-Palmbach, et al. [45]	A blockchain-orchestrated Federated Learning architecture for healthcare consortia (<i>arXiv version 1</i>)			●		●	●		
Preuveneers, et al. [52]	Chained anomaly detection models for federated learning: An intrusion detection case study			●	●	●	●		
Ramanan, et al. [44]	BAFFLE: Blockchain based Aggregator Free Federated Learning (<i>arXiv version 1</i>)			●	●	●			
Salah, et al. [12]	Blockchain for AI: Review and Open Research Challenges								●
Sarpatawar, et al. [51]	Towards Enabling Trusted Artificial Intelligence via Blockchain			●		●	●		
Shayan, et al. [73]	Biscotti: A Ledger for Private and Secure Peer-to-Peer Machine Learning (<i>arXiv version 4</i>)				●	●			
Swan [72]	Blockchain Thinking: The Brain as a Decentralized Autonomous Corporation [Commentary]				●				
Tann, et al. [38]	Towards Safer Smart Contracts: A Sequence Learning Approach to Detecting Security Threats (<i>arXiv version 3</i>)	●							
Turesson, et al. [70]	Privacy-Preserving Blockchain Mining: Sybil-resistance by Proof-of-Useful-Work (<i>arXiv version 2</i>)				●				
Wang, et al. [32]	When Blockchain Meets AI: Optimal Mining Strategy Achieved by Machine Learning (<i>arXiv version 1</i>)	●							
Zhang, et al. [39]	A game-theoretic method based on Q-learning to invalidate criminal smart contracts	●							
Zhang, et al. [59]	Genie: A Secure, Transparent Sharing and Services Platform for Genetic and Health Data (<i>arXiv version 1</i>)				●	●	●	●	
Zhao, et al. [43]	Mobile Edge Computing, Blockchain and Reputation-based Crowdsourcing IoT Federated Learning: A Secure, Decentralized and Privacy-preserving System (<i>arXiv version 1</i>)			●	●	●			

between theory and practice by drawing theoretical conclusions from practical research and outlining future practical research opportunities from theory. Fourth, we describe how convergence creates innovation in an emerging field.

This article provides insights for researchers and practitioners interested in deepening their knowledge for interdisciplinary applications of any of the fields: AI, DLT, their convergence, or convergence in general. By

providing these insights with an overview on the upcoming convergence of DLT and AI, we contribute to the development of future innovations in this fast-paced field problem.

APPENDIX

A. RELEVANT AND ANALYZED PAPERS

See Table 6.

ACKNOWLEDGMENT

The authors would like to thank Mikael Beyene and Niclas Kannengießler for taking the time to discuss the structure of this article.

REFERENCES

- [1] J. Bughin, J. Seong, J. Manyika, M. Chui, and R. Joshi. *Notes From the AI Frontier: Modeling the Impact of AI on the World Economy*. Accessed: Mar. 4, 2020. [Online]. Available: <https://www.mckinsey.com/~media/McKinsey/Featured%20Insights/Artificial%20Intelligence/Notes%20from%20the%20frontier%20Modeling%20the%20impact%20of%20AI%20on%20the%20world%20economy/MGI-Notes-from-the-AI-frontier-Modeling-the-impact-of-AI-on-the-world-economy-September-2018.ashx>
- [2] World Economic Forum. *Building Block(chain)s for a Better Planet*. Accessed: Mar. 4, 2020. [Online]. Available: http://www3.weforum.org/docs/WEF_Building-Blockchains.pdf
- [3] N. Kannengießler, S. Lins, T. Dehling, and A. Sunyaev, "Mind the gap: Trade-offs between distributed ledger technology characteristics," *ACM Comput. Surv.*, to be published. [Online]. Available: <https://arxiv.org/abs/1906.00861v1>
- [4] A. Madani, R. Arnaout, M. Mofrad, and R. Arnaout, "Fast and accurate view classification of echocardiograms using deep learning," *NPJ Digit. Med.*, vol. 1, no. 1, p. 6, Dec. 2018.
- [5] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Müller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba, "End to end learning for self-driving cars," 2016, *arXiv:1604.07316v1*. [Online]. Available: <https://arxiv.org/abs/1604.07316v1>
- [6] N. Sünderhauf, O. Brock, W. Scheirer, R. Hadsell, D. Fox, J. Leitner, B. Uperoft, P. Abbeel, W. Burgard, M. Milford, and P. Corke, "The limits and potentials of deep learning for robotics," *Int. J. Robot. Res.*, vol. 37, nos. 4–5, pp. 405–420, Apr. 2018.
- [7] S. Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. Accessed: Mar. 4, 2020. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [8] G. G. Dagher, J. Mohler, M. Milojkovic, and P. B. Marella, "Ancile: Privacy-preserving framework for access control and interoperability of electronic health records using blockchain technology," *Sustain. Cities Soc.*, vol. 39, pp. 283–297, May 2018.
- [9] L. Davi, D. Hatebur, M. Heisel, and R. Wirtz, "Combining safety and security in autonomous cars using blockchain technologies," in *Proc. Int. Conf. Comput. Saf., Rel., Secur.* Cham, Switzerland: Springer, 2019, pp. 223–234.
- [10] C. Hou, M. Zhou, Y. Ji, P. Daian, F. Tramer, G. Fanti, and A. Juels, "Squirrel: Automating attack discovery on blockchain incentive mechanisms with deep reinforcement learning," 2019, *arXiv:1912.01798v1*. [Online]. Available: <http://arxiv.org/abs/1912.01798v1>
- [11] N. Hynes, D. Dao, D. Yan, R. Cheng, and D. Song, "A demonstration of sterling: A privacy-preserving data marketplace," *Proc. VLDB Endowment*, vol. 11, no. 12, pp. 2086–2089, Aug. 2018.
- [12] K. Salah, M. H. U. Rehman, N. Nizamuddin, and A. Al-Fuqaha, "Blockchain for AI: Review and open research challenges," *IEEE Access*, vol. 7, pp. 10127–10149, 2019.
- [13] T. N. Dinh and M. T. Thai, "AI and blockchain: A disruptive integration," *Computer*, vol. 51, no. 9, pp. 48–53, Sep. 2018.
- [14] S. Jeong, J.-C. Kim, and J. Y. Choi, "Technology convergence: What developmental stage are we in?" *Scientometrics*, vol. 104, no. 3, pp. 841–871, Sep. 2015.
- [15] G. Duysters and J. Hagedoorn, "Technological convergence in the IT industry: The role of strategic technology alliances and technological competencies," *Int. J. Econ. Bus.*, vol. 5, no. 3, pp. 355–368, Nov. 1998.
- [16] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Comput. Secur. (CCS)*, 2016, pp. 308–318.
- [17] B. Goertzel and C. Pennachin, *Artificial General Intelligence*. Berlin, Germany: Springer-Verlag, 2007.
- [18] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [19] W. Samek and K.-R. Müller, "Towards explainable artificial intelligence," in *Explainable AI: Interpreting, Explaining Visualizing Deep Learning*. Cham, Switzerland: Springer, 2019, pp. 5–22.
- [20] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.
- [21] D. Mahajan, "Exploring the limits of weakly supervised pretraining," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 181–196.
- [22] B. W. Silverman and M. C. Jones, "E. Fix and J.L. Hodges (1951): An important contribution to nonparametric discriminant analysis and density estimation: Commentary on Fix and Hodges (1951)," *Int. Stat. Rev.*, vol. 57, no. 3, pp. 233–238, 1989.
- [23] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ, USA: Prentice-Hall, 2016.
- [24] L. G. M. E. Eykholt and J. Denman. *RChain Architecture Documentation*. Accessed: Mar. 4, 2020. [Online]. Available: <https://buildmedia.readthedocs.org/media/pdf/rchain-architecture/stable/rchain-architecture.pdf>
- [25] S. Popov. *The Tangle, IOTA Whitepaper*. Accessed: Mar. 4, 2020. [Online]. Available: https://assets.ctfassets.net/r1dr6vzfxhev/2t4uxvsIqk0EUau6g2sw0g/45eae33637ca92f85dd9f4a3a218e1ec/iota1_4_3.pdf
- [26] S. Thiebes, N. Kannengießler, M. Schmidt-Kraepelin, and A. Sunyaev, "Beyond data markets: Opportunities and challenges for distributed ledger technology in genomics," in *Proc. 53rd Hawaii Int. Conf. Syst. Sci.*, 2020, pp. 1–10.
- [27] T. Min, H. Wang, Y. Guo, and W. Cai, "Blockchain games: A survey," 2019, *arXiv:1906.05558v1*. [Online]. Available: <http://arxiv.org/abs/1906.05558v1>
- [28] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum Project Yellow Paper*, vol. 151, pp. 1–32, Apr. 2014.
- [29] E. B. Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from bitcoin," in *Proc. IEEE Symp. Secur. Privacy*, May 2014, pp. 459–474.
- [30] C.-S. Curran, S. Bröring, and J. Leker, "Anticipating converging industries using publicly available data," *Technological Forecasting Social Change*, vol. 77, no. 3, pp. 385–395, Mar. 2010.
- [31] E. Karafiloski and A. Mishev, "Blockchain solutions for big data challenges: A literature review," in *Proc. 17th Int. Conf. Smart Technol. (EUROCON)*, Jul. 2017, pp. 763–768.
- [32] T. Wang, S. Chang Liew, and S. Zhang, "When blockchain meets AI: Optimal mining strategy achieved by machine learning," 2019, *arXiv:1911.12942v1*. [Online]. Available: <http://arxiv.org/abs/1911.12942v1>
- [33] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.* Berlin, Germany: Springer, 2014, pp. 436–454.
- [34] Y. Sompolinsky and A. Zohar, "Secure high-rate transaction processing in bitcoin," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.* Berlin, Germany: Springer, 2015, pp. 507–527.
- [35] A. Troghhi Haghghat and M. Shajari, "Block withholding game among bitcoin mining pools," *Future Gener. Comput. Syst.*, vol. 97, pp. 482–491, Aug. 2019.
- [36] R. Camino, C. Ferreira Torres, M. Baden, and R. State, "A data science approach for honeypot detection in ethereum," 2019, *arXiv:1910.01449v2*. [Online]. Available: <http://arxiv.org/abs/1910.01449v2>
- [37] Y. Kim, D. Pak, and J. Lee, "ScanAT: Identification of bytecode-only smart contracts with multiple attribute tags," *IEEE Access*, vol. 7, pp. 98669–98683, 2019.
- [38] W. J.-W. Tann, X. J. Han, S. S. Gupta, and Y.-S. Ong, "Towards safer smart contracts: A sequence learning approach to detecting security threats," 2018, *arXiv:1811.06632v3*. [Online]. Available: <http://arxiv.org/abs/1811.06632v3>
- [39] L. Zhang, Y. Wang, F. Li, Y. Hu, and M. H. Au, "A game-theoretic method based on Q-learning to invalidate criminal smart contracts," *Inf. Sci.*, vol. 498, pp. 144–153, Sep. 2019.
- [40] L.-N. Lundbæk, D. Janes Beutel, M. Huth, S. Jackson, L. Kirk, and R. Steiner, "Proof of kernel work: A democratic low-energy consensus for distributed access-control protocols," *Roy. Soc. Open Sci.*, vol. 5, no. 8, Aug. 2018, Art. no. 180422.

- [41] M. E. Gladden, "Cryptocurrency with a conscience: Using artificial intelligence to develop money that advances human ethical values," *Annales. Etyka W yciu Gospodarczym*, vol. 18, no. 4, pp. 85–98, 2015.
- [42] S. Liu, F. Mohsin, L. Xia, and O. Seneviratne, "Strengthening smart contracts to handle unexpected situations," in *Proc. IEEE Int. Conf. Decentralized Appl. Infrastruct. (DAPPCON)*, Apr. 2019, pp. 182–187.
- [43] Y. Zhao, J. Zhao, L. Jiang, R. Tan, and D. Niyato, "Mobile edge computing, blockchain and reputation-based crowdsourcing IoT federated learning: A secure, decentralized and privacy-preserving system," 2019, *arXiv:1906.10893v1*. [Online]. Available: <http://arxiv.org/abs/1906.10893v1>
- [44] P. Ramanan, K. Nakayama, and R. Sharma, "BAFFLE: Blockchain based aggregator free federated learning," 2019, *arXiv:1909.07452v1*. [Online]. Available: <http://arxiv.org/abs/1909.07452v1>
- [45] J. Passerat-Palmbach, T. Farnan, R. Miller, M. S. Gross, H. Leigh Flannery, and B. Gleim, "A blockchain-orchestrated federated learning architecture for healthcare consortia," 2019, *arXiv:1910.12603v1*. [Online]. Available: <http://arxiv.org/abs/1910.12603v1>
- [46] A. Juels, A. Kosba, and E. Shi, "The ring of gyges: Investigating the future of criminal smart contracts," presented at the ACM SIGSAC Conf. Comput. Commun. Secur., Vienna, Austria, Oct. 2016, doi: 10.1145/2976749.2978362.
- [47] J. Kang, M. Bennett, D. Carbedo, and P. Casey, "Implicit bias in the courtroom," *UCLA Law Rev.*, vol. 59, p. 1124, Jun. 2012.
- [48] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2008, pp. 111–125.
- [49] J. Isaak and M. J. Hanna, "User data privacy: Facebook, Cambridge analytica, and privacy protection," *Computer*, vol. 51, no. 8, pp. 56–59, Aug. 2018.
- [50] T. Yang, G. Andrew, H. Eichner, H. Sun, W. Li, N. Kong, D. Ramage, and F. Beaufays, "Applied federated learning: Improving Google keyboard query suggestions," 2018, *arXiv:1812.02903v1*. [Online]. Available: <http://arxiv.org/abs/1812.02903v1>
- [51] K. Sarpatwar, R. Vaculin, H. Min, G. Su, T. Heath, G. Ganapavarapu, and D. Dillenberger, "Towards enabling trusted artificial intelligence via blockchain," in *Policy-Based Autonomic Data Governance*. Cham, Switzerland: Springer, 2019, pp. 137–153.
- [52] D. Preuveneers, V. Rimmer, I. Tsiengenopoulos, J. Spooren, W. Joosen, and E. Ilie-Zudor, "Chained anomaly detection models for federated learning: An intrusion detection case study," *Appl. Sci.*, vol. 8, no. 12, p. 2663, 2018.
- [53] X. Huang, M. Kwiatkowska, S. Wang, and M. Wu, "Safety verification of deep neural networks," in *Proc. Int. Conf. Comput. Aided Verification*. Springer, 2017, pp. 3–29.
- [54] B. W. Josh Swihart and S. Bowe, *Zcash Counterfeiting Vulnerability Successfully Remediated*. Accessed: Mar. 4, 2020. [Online]. Available: <https://electriccoin.co/blog/zcash-counterfeiting-vulnerability-successfully-remediated/>
- [55] M. Pradel and K. Sen, "DeepBugs: A learning approach to name-based bug detection," *ACM Program. Lang.*, vol. 2, pp. 1–25, Oct. 2018.
- [56] M. Bartoletti and L. Pompianu, "An empirical analysis of smart contracts: Platforms, applications, and design patterns," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.* Cham, Switzerland: Springer, 2017, pp. 494–509.
- [57] M. Allamanis, E. T. Barr, P. Devanbu, and C. Sutton, "A survey of machine learning for big code and naturalness," *ACM Comput. Surv.*, vol. 51, no. 4, pp. 1–37, Jul. 2018.
- [58] P. Daian, "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2020, pp. 566–583.
- [59] S. Zhangy, A. Kim, D. Liu, S. C. Nuckchady, L. Huangy, A. Masurkary, J. Zhangy, L. P. Karnatiz, L. Martinezx, T. Hardjono, M. Kellis, and Z. Zhang, "Genie: A secure, transparent sharing and services platform for genetic and health data," 2018, *arXiv:1811.01431v1*. [Online]. Available: <http://arxiv.org/abs/1811.01431v1>
- [60] V. Raychev, M. Vechev, and E. Yahav, "Code completion with statistical language models," *ACM SIGPLAN Notices*, vol. 49, no. 6, pp. 419–428, Jun. 2014.
- [61] A. Hindle, E. T. Barr, Z. Su, M. Gabel, and P. Devanbu, "On the naturalness of software," in *Proc. 34th Int. Conf. Softw. Eng. (ICSE)*, Jun. 2012, pp. 837–847.
- [62] S. Eskandari, S. Moosavi, and J. Clark, "SoK: Transparent dishonesty: Front-running attacks on blockchain," in *Financial Cryptography and Data Security*. Cham, Switzerland: Springer, 2020, pp. 170–189.
- [63] A. Judmayer, "Pay-to-win: Incentive attacks on proof-of-work cryptocurrencies," *Cryptol. ePrint Arch.*, 2019. [Online]. Available: <https://eprint.iacr.org/2019/775/20200217:161450>
- [64] H. Bae, J. Jang, D. Jung, H. Jang, H. Ha, and S. Yoon, "Security and privacy issues in deep learning," 2018, *arXiv:1807.11655v3*. [Online]. Available: <http://arxiv.org/abs/1807.11655v3>
- [65] UN Global Working Group on Big Data. *UN Handbook on Privacy-Preserving Computation Techniques*. Accessed: Mar. 4, 2020. [Online]. Available: <http://publications.officialstatistics.org/handbooks/privacy-preserving-techniques-handbook/UN%20Handbook%20for%20Privacy-Preserving%20Techniques.pdf>
- [66] H. Cui, H. Zhang, G. R. Ganger, P. B. Gibbons, and E. P. Xing, "GeepS: Scalable deep learning on distributed GPUs with a GPU-specialized parameter server," in *Proc. 11th Eur. Conf. Comput. Syst. (EuroSys)*, 2016, pp. 1–16.
- [67] A. L. Beberg, D. L. Ensign, G. Jayachandran, S. Khaliq, and V. S. Pande, "Folding@home: Lessons from eight years of volunteer distributed computing," in *Proc. IEEE Int. Symp. Parallel Distrib. Process.*, May 2009, pp. 1–8.
- [68] J. Truby, "Decarbonizing bitcoin: Law and policy choices for reducing the energy consumption of blockchain technologies and digital currencies," *Energy Res. Social Sci.*, vol. 44, pp. 399–410, Oct. 2018.
- [69] C. Chenli, B. Li, Y. Shi, and T. Jung, "Energy-recycling blockchain with proof-of-deep-learning," 2019, *arXiv:1902.03912v1*. [Online]. Available: <http://arxiv.org/abs/1902.03912v1>
- [70] H. Turesson, M. Laskowski, A. Roatis, and H. M. Kim, "Privacy-preserving blockchain mining: Sybil-resistance by proof-of-useful-work," 2019, *arXiv:1907.08744v2*. [Online]. Available: <http://arxiv.org/abs/1907.08744v2>
- [71] A. Baldominos and Y. Saez, "Coin.AI: A proof-of-useful-work scheme for blockchain-based distributed deep learning," *Entropy*, vol. 21, no. 8, p. 723, 2019.
- [72] M. Swan, "Blockchain thinking: The brain as a decentralized autonomous corporation [commentary]," *IEEE Technol. Soc. Mag.*, vol. 34, no. 4, pp. 41–52, Dec. 2015.
- [73] M. Shayan, C. Fung, C. J. M. Yoon, and I. Beschastnikh, "Biscotti: A ledger for private and secure peer-to-peer machine learning," 2018, *arXiv:1811.09904v4*. [Online]. Available: <http://arxiv.org/abs/1811.09904v4>
- [74] Y. Lu, Q. Tang, and G. Wang, "On enabling machine learning tasks atop public blockchains: A crowdsourcing approach," in *Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW)*, Nov. 2018, pp. 81–88.
- [75] K. R. Özyilmaz, M. Do'an, and A. Yurdakul, "IDMoB: IoT data marketplace on blockchain," in *Proc. Crypto Valley Conf. Blockchain Technol. (CVCBT)*, Jun. 2018, pp. 11–19.
- [76] J. D. Harris and B. Waggoner, "Decentralized and collaborative AI on blockchain," in *Proc. IEEE Int. Conf. Blockchain (Blockchain)*, Jul. 2019, pp. 368–375.
- [77] J. Kang, Z. Xiong, D. Niyato, Y. Zou, Y. Zhang, and M. Guizani, "Reliable federated learning for mobile networks," 2019, *arXiv:1910.06837v1*. [Online]. Available: <http://arxiv.org/abs/1910.06837v1>
- [78] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. Johnson, A. Juels, A. Miller, and D. Song, "Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contracts," in *Proc. IEEE Eur. Symp. Secur. Privacy (EuroS&P)*, Jun. 2019, pp. 185–200.
- [79] J. Eberhardt and J. Heiss, "Off-chaining models and approaches to off-chain computations," in *Proc. 2nd Workshop Scalable Resilient Infrastruct. Distrib. Ledgers (SERIAL)*, 2018, pp. 7–12.
- [80] X. Chen, J. Ji, C. Luo, W. Liao, and P. Li, "When machine learning meets blockchain: A decentralized, privacy-preserving and secure design," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2018, pp. 1178–1187.
- [81] H. Kim, J. Park, M. Bennis, and S.-L. Kim, "Blockchain on-device federated learning," *IEEE Commun. Lett.*, to be published.
- [82] M. Banko and E. Brill, "Scaling to very very large corpora for natural language disambiguation," in *Proc. 39th Annu. Meeting Assoc. Comput. Linguistics-ACL*, 2001, pp. 26–33.
- [83] G. A. Montes and B. Goertzel, "Distributed, decentralized, and democratized artificial intelligence," *Technol. Forecasting Social Change*, vol. 141, pp. 354–358, Apr. 2019.

- [84] M. Jones, M. Johnson, M. Shervey, J. T. Dudley, and N. Zimmerman, "Privacy-preserving methods for feature engineering using blockchain: Review, evaluation, and proof of concept," *J. Med. Internet Res.*, vol. 21, no. 8, 2019, Art. no. e13600.
- [85] M. Johnson, M. Jones, M. Shervey, J. T. Dudley, and N. Zimmerman, "Building a secure biomedical data sharing decentralized app (DApp): Tutorial," *J. Med. Internet Res.*, vol. 21, no. 10, 2019, Art. no. e13601.
- [86] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting unintended feature leakage in collaborative learning," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2019, pp. 691–706.
- [87] V. Costan and S. Devadas, "Intel SGX explained," *IACR Cryptol. ePrint Arch.*, vol. 2016, no. 86, pp. 1–118, 2016. [Online]. Available: <https://eprint.iacr.org/2016/086/20170221:054353>
- [88] F. Tschorsch and B. Scheuermann, "Bitcoin and beyond: A technical survey on decentralized digital currencies," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 3, pp. 2084–2123, 3rd Quart., 2016.
- [89] S. King, *Primecoin: Cryptocurrency With Prime Number Proof-of-Work*. Accessed: Mar. 4, 2020. [Online]. Available: <https://primecoin.io/bin/primecoin-paper.pdf>
- [90] F. Saleh, "Blockchain without waste: Proof-of-stake," *SSRN Electron. J.*, 2019. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3183935
- [91] N. Hynes, R. Cheng, and D. Song, "Efficient deep learning on multi-source private data," 2018, *arXiv:1807.06689v1*. [Online]. Available: <http://arxiv.org/abs/1807.06689v1>
- [92] R. Kunkel, D. Le Quoc, F. Gregor, S. Arnautov, P. Bhatotia, and C. Fetzer, "TensorSCONE: A secure TensorFlow framework using intel SGX," 2019, *arXiv:1902.04413v1*. [Online]. Available: <http://arxiv.org/abs/1902.04413v1>
- [93] F. Brasser, T. Frassetto, K. Riedhammer, A.-R. Sadeghi, T. Schneider, and C. Weinert, "VoiceGuard: Secure and private speech processing," in *Proc. InterSpeech*, Sep. 2018, pp. 1303–1307.
- [94] D. Lee, D. Kohlbrenner, S. Shinde, D. Song, and K. Asanović, "Keystone: An open framework for architecting TEEs," 2019, *arXiv:1907.10119v2*. [Online]. Available: <http://arxiv.org/abs/1907.10119v2>
- [95] F. Brasser, S. Capkun, A. Dmitrienko, T. Frassetto, K. Kostiainen, and A.-R. Sadeghi, "DR.SGX: Automated and adjustable side-channel protection for SGX using data location randomization," in *Proc. 35th Annu. Comput. Secur. Appl. Conf. (ACSAC)*, 2019, pp. 788–800.
- [96] M. Milutinovic, W. He, H. Wu, and M. Kanwal, "Proof of luck: An efficient blockchain consensus protocol," in *Proc. 1st Workshop Syst. Softw. Trusted Execution (SysTEX)*, 2016, pp. 1–6.
- [97] L. Chen, L. Xu, N. Shah, Z. Gao, Y. Lu, and W. Shi, "On security analysis of proof-of-elapsed-time (PoET)," in *Proc. Int. Symp. Stabilization, Saf., Secur. Distrib. Syst.* Cham, Switzerland: Springer, 2017, pp. 282–297.
- [98] A. Miller, I. Bentov, S. Bakshi, R. Kumaresan, and P. McCorry, "Sprites and state channels: Payment networks that go faster than lightning," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.* Cham, Switzerland: Springer, 2019, pp. 508–526.
- [99] J. Poon and V. Buterin, *Plasma: Scalable Autonomous Smart Contracts*. Accessed: Mar. 4, 2020. [Online]. Available: <https://www.plasma.io/plasma.pdf>
- [100] S. Dziembowski, S. Faust, and K. Hostáková, "General state channel networks," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Jan. 2018, pp. 949–966.
- [101] H. Kalodner, S. Goldfeder, X. Chen, S. M. Weinberg, and E. W. Felten, "Arbitrum: Scalable, private smart contracts," in *27th USENIX Secur. Symp. (USENIX Security)*, 2018, pp. 1353–1370.
- [102] J. Teutsch and C. Reitwießner, "A scalable verification solution for blockchains," 2019, *arXiv:1908.04756v1*. [Online]. Available: <http://arxiv.org/abs/1908.04756v1>
- [103] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," 2017, *arXiv:1712.07557v2*. [Online]. Available: <http://arxiv.org/abs/1712.07557v2>
- [104] High-Level Expert Group on Artificial Intelligence. *Ethics Guidelines for Trustworthy AI*. Accessed: Mar. 4, 2020. [Online]. Available: https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419
- [105] S. Schulz, A.-R. Sadeghi, and C. Wachsmann, "Short paper: Lightweight remote attestation using physical functions," in *Proc. 4th ACM Conf. Wireless Netw. Secur. (WiSec)*, 2011, pp. 109–114.
- [106] S. Volos, K. Vaswani, and R. Bruno, "Graviton: Trusted execution environments on GPUs," in *Proc. 13th USENIX Symp. Oper. Syst. Design Implement. (OSDI)*, 2018, pp. 681–696.
- [107] S. Bano, A. Sonnino, M. Al-Bassam, S. Azouvi, P. McCorry, S. Meiklejohn, and G. Danezis, "SoK: Consensus in the age of blockchains," in *Proc. 1st ACM Conf. Adv. Financial Technol. (AFT)*, 2019, pp. 183–198.
- [108] J. Van Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx, "Foreshadow: Extracting the keys to the Intel SGX kingdom with transient out-of-order execution," in *Proc. 27th USENIX Secur. Symp. (USENIX Security)*, 2018, pp. 991–1008.
- [109] P. Das, "FastKitten: Practical smart contracts on bitcoin," in *Proc. 28th USENIX Secur. Symp. (USENIX Security)*, 2019, pp. 801–818.
- [110] A. Vijayakumar and S. Kundu, "A novel modeling attack resistant PUF design based on non-linear voltage transfer characteristics," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, 2015, pp. 653–658.
- [111] F. Ganji, *On the Learnability of Physically Unclonable Functions*. Cham, Switzerland: Springer, 2018.



KONSTANTIN D. PANDL studied electrical engineering and information technology at the Karlsruhe Institute of Technology (KIT), Germany, Purdue University, USA, and Tongji University, China, and graduated with the master's degree. He is currently pursuing the Ph.D. degree with the Institute of Applied Informatics and Formal Description Methods, KIT. He also gained experience in industry at Siemens' Venture Unit Next47, San Francisco Bay Area, and Kearney, Germany.

His research interests include machine learning, system security, and distributed systems. His previous research appeared at the IEEE International Conference on Intelligent Transportation Systems.



SCOTT THIEBESS received the degree in information systems from the University of Cologne. He is currently pursuing the Ph.D. degree with the Institute of Applied Informatics and Formal Description Methods, Karlsruhe Institute of Technology, Germany. His research interests include emerging technologies in health care, information security and privacy, and gamification. His recent research interests include on the implications of applying distributed ledger technology and artificial intelligence in health care settings. His research appeared in journals, including the *Journal of Medical Internet Research*, *BMC Bioinformatics*, and the *European Journal of Human Genetics*.



MANUEL SCHMIDT-KRAEPELIN received the degree in information systems from the University of Cologne. He is currently pursuing the Ph.D. degree with the Institute of Applied Informatics and Formal Description Methods, Karlsruhe Institute of Technology, Germany. He is also a Research Associate with the Institute of Applied Informatics and Formal Description Methods, Karlsruhe Institute of Technology. His research interests include gamification and emerging technologies in health care. His research appeared in leading scientific conferences, including ICIS, ECIS, and HICSS.



ALI SUNYAEV is currently a Professor of computer science with the Karlsruhe Institute of Technology, Germany. His research interests include trustworthy Internet technologies and complex health IT applications. His research work accounts for the multifaceted use contexts of digital technologies with research on human behavior affecting Internet-based systems and vice versa. His research appeared in journals, including *ACM CSUR*, *JIT*, *JMIS*, the *IEEE TRANSACTIONS ON*

CLOUD COMPUTING, *Communications of the ACM*, and others. His research work has been appreciated numerous times and is featured in a variety of media outlets.

...