



HAL
open science

Accountable Federated Machine Learning in Government: Engineering and Management Insights

Dian Balta, Mahdi Sellami, Peter Kuhn, Ulrich Schöpp, Matthias Buchinger,
Nathalie Baracaldo, Ali Anwar, Heiko Ludwig, Mathieu Sinn, Mark Purcell,
et al.

► **To cite this version:**

Dian Balta, Mahdi Sellami, Peter Kuhn, Ulrich Schöpp, Matthias Buchinger, et al.. Accountable Federated Machine Learning in Government: Engineering and Management Insights. 13th International Conference on Electronic Participation (ePart), Sep 2021, Granada, Spain. pp.125-138, 10.1007/978-3-030-82824-0_10 . hal-04014056

HAL Id: hal-04014056

<https://inria.hal.science/hal-04014056>

Submitted on 3 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

Accountable Federated Machine Learning in Government: Engineering and Management Insights

Dian Balta¹, Mahdi Sellami¹, Peter Kuhn^{1[0000-0001-6774-2904]}, Ulrich Schöpp¹, Matthias Buchinger¹, Nathalie Baracaldo², Ali Anwar², Heiko Ludwig², Mathieu Sinn², Mark Purcell², Bashar Altakrouri³

¹ fortiss GmbH, ² IBM Research, ³ IBM Cloud and Cognitive Software
balta@fortiss.org

Abstract. Machine learning offers promising capabilities to improve administrative procedures. At the same time, adequate training of models using traditional learning techniques requires the collection and storage of enough training data in a central place. Unfortunately, due to legislative and jurisdictional constraints, data in a central place is scarce and training a model becomes unfeasible. Against this backdrop, federated machine learning, a technique to collaboratively train models without transferring data to a centralized location, has been recently proposed. With each government entity keeping their data private, new applications that previously were impossible now can be a reality. In this paper, we demonstrate that accountability for the federated machine learning process becomes paramount to fully overcoming legislative and jurisdictional constraints. In particular, it ensures that all government entities' data are adequately included in the model and that evidence on fairness and reproducibility is curated towards trustworthiness. We also present an analysis framework suitable for governmental scenarios and illustrate its exemplary application for online citizen participation scenarios. We discuss our findings in terms of engineering and management implications: feasibility evaluation, general architecture, involved actors as well as verifiable claims for trustworthy machine learning.

Keywords: accountability, federated learning, framework, verifiable claims

1 Introduction

Machine learning (ML)—a concept that incorporates various characteristics of intelligent systems that follow particular goals, have a formal representation of knowledge and an automated logical inference [1]—is expected to represent a huge leap from data analysis to high quality and efficiency predictions, and increases the value of informed judgements [2].

At the same time, ML has to address general issues such as fairness, ethics, robustness and explainability of trained models (cf. e.g. [3, 4]), in order to address legal, economic and political uncertainty. The existing uncertainty regarding machine learning exposes government practice to **engineering and management challenges** (cf. e.g. [5])

and applications have yet to deliver sustainable and reproducible results in the government domain.

Three noteworthy challenges are **accountability** [6], **data sharing** (cf. e.g. [7, 8]) and **privacy preservation** (c.f. e.g. [9]). Accountability represents a relation, where a government is accountable for ML to the users of its ML-based services by providing transparency means as well as mechanisms that allow them to enforce control. For instance, users should be able to resolve responsibilities in case of data and model bias (cf. e.g. [10]), to request changes to existing policies (cf. e.g. [11]) as well as to install third-party audits (cf. e.g. [12]).

The second challenge, data sharing in government, emerges out of technical issues such as interoperability and heterogeneity of data infrastructure, but also legal constraints and organization resistance. For instance, in governments where jurisdiction is split along federal levels as well as department competencies, the lack of data sharing is motivated by a missing legal basis and corresponding administrative procedures. Moreover, a common information model is often needed (cf. e.g. [13, 14]). The third challenge, preserving privacy, is a cornerstone of digitization and machine learning, and is of crucial importance for governments. For instance, GDPR in the EU and CCPA as well as HIPAA in the USA are legal frameworks that introduce extensive requirements for government information systems.

The objective of this paper is to provide an *analysis framework* for addressing the described engineering and management challenges in government based on an approach named *Accountable Federated Machine Learning (AFML)*. In this context, accountability is focused on creating verifiable claims [15] towards trustworthy engineering of machine learning [12]. Federated Machine Learning (FML) is a novel approach to apply machine learning to generated knowledge based on shared models, but keep the data private at each participating party's side during the training process (cf. e.g. [16]). FML allows parties to collaborate on one of the main challenges of machine learning: quality and quantity of the training data.

With the analysis framework based on AFML, we address the question: *What engineering and managerial aspects should be considered along the process of introducing novel ML approaches in the government domain?* We discuss an argument that standardization artefacts (e.g. business processes, models, shared terminologies, software tools etc.) are required in the course of analysis. To support this argument, we present findings from a prototype setup of AFML within a use case of citizen participation in Germany and discuss their implications. We believe that our research should be of value to both researchers and practitioners, given the current progress in similar domains.

2 Theoretical Background

2.1 Federated Machine Learning

The term federated learning was recently introduced by McMahan et al. [16]: “*We term our approach Federated Learning, since the learning task is solved by a loose federation of participating devices (which we refer to as clients) which are coordinated by a*

central server. “. The real-world challenge addressed was primarily to learn from millions of devices (e.g. smart phones) by federating models [17]. Since then, interest in the research community evolved and included data-silos across organizations rather across single end-users. With respect to this development, a broader definition was introduced:

“Federated learning is a machine learning setting where multiple entities (clients) collaborate in solving a machine learning problem, under the coordination of a central server or service provider. Each client’s raw data is stored locally and not exchanged or transferred; instead, focused updates intended for immediate aggregation are used to achieve the learning objective.” [17]

In this paper, we chose to use the term federated machine learning (FML) instead of federated learning in order to avoid confusion among researchers outside the machine learning community. Thereby, we refer to the clients as FML **parties** and to the services provider as FML **aggregator**.

2.2 Accountability

Accountability can be generally defined as a *“a relationship between an actor and a forum, in which the actor has an obligation to explain and to justify his or her conduct, the forum can pose questions and pass judgment, and the actor may face consequences”*[18]. Accountability has been considered as a guiding principle for system design back in the 1960s (cf. e.g.[19]). In this paper, accountability is focused on creating verifiable claims [15] towards trustworthy FML [12]. In accordance with [20], trustworthiness from an engineering perspective can be represented as argument that aims at explaining the design of a system. Moreover, the argument should explain checks and tests performed during its development to ensure particular system properties. The argument should be organized around particular claims (or goals) and supporting evidence about the system. One can visualize it as a tree, broken down into claims and subclaims (the interior is the reasoning) with evidence at the leaves. To make the claims verifiable, they should be formalized and their evolvment (as well as of the supporting evidence) should allow for an audit trace (cf. e.g. [12, 21–23]).

From a management perspective, trustworthiness concerns assignable responsibilities, distributed obligations and corresponding rules and agreements (cf. e.g. [24]). Consequently, accountability can be presented in terms of: (i) compliance concerned with ensuring that activities are in accordance with prescribed and/or agreed norms (cf. e.g. [25–27]), (ii) control among the parties involved in a process that influence the conditions of the process (cf. e.g. [24]) and (iii) a regulatory framework that defines the requirements, goals, and completion criteria of a process in a satisfactory manner to the parties, who can build a consensus regarding their judgement on the completion of the process in a verifiable way (cf. e.g. [24, 25]).

2.3 Standardization

Standardization artefacts, representing among others a uniform set of agreements or specifications between the actors who develop and apply them, are a suitable approach

to study the applicability of a technology in government [28]. In order to analyze IT standardization artefacts in government, the following framework that consist of two dimension can be applied [29, 30]. The first dimension includes three levels of interoperability and the second dimension includes five functional views. The interoperability dimension is structured along three layers. First, business processes applied in delivering public service are found on the organizational layer. Second, exchange of information and data as well as to their meaning between parties involved is found on the semantic layer. Third, data structure and format, sending and receiving data based on communication protocols, electronic mechanisms to store data as well as software and hardware are situated on the technical/ syntactic layer.

The second dimension includes five functional views. The administration view includes predominantly non-technical standards. They affect personnel and process aspects as well as communication within or between public administrations. Second, the modeling view includes reference models and architectures, as well as modelling languages for each corresponding interoperability level. Third, standards that focus on the computation of data are included in the processing view. Fourth, corresponding standards for data and information exchange between different public administrations is handled in the communication and interaction view. Fifth, the security and privacy view contains standards that aim at addressing issues such as definition of access management policies, cryptography methods or requesting a required minimum of personal data.

3 Research Approach

We follow a qualitative analysis approach to explorative research. We aim at developing a descriptive artefact (analysis framework) that can be categorized as a theory for analyzing [31]. Our research approach is rooted in the paradigm of pragmatism [32]. We studied the findings through an argumentative-deductive analysis [33], which comprised theoretically-founded concept development and prototype development. We conducted a hermeneutic literature review [30] to study the theoretical foundations. Thereby, we started to develop our understanding of the concepts of accountability [19, 34, 35] and FML [16, 17, 36, 37] and derived implications from a standardization perspective [28, 29].

For our prototype, we used data and a list of challenges from a research project on online citizen participation [28], where missing data showed the potential of data sharing while at the same time data privacy concerns hampered any progress on sharing. Online citizen participation can be described as a form of participation that is based on the usage of information and communication technology in societal democratic and consultative processes focused on citizens [38].

The envisioned application of machine learning with a particular relevance for online citizen participation is a methodology named natural language processing (NLP). NLP has been already applied in government practice (e.g. [39]) as well as in online citizen participation (e.g. [40]). With respect to methodological development and practical tool

availability of NLP, we decided to create a prototype architecture that involves *classification of ideas sampled during online citizen participation* and focuses on the following challenge: *How to apply FML for NLP classification tasks that are traditionally performed in a centralized manner?*

The data used for NLP in the prototype were collected during several citizen participation sessions in a German city. The dataset is a collection of citizen ideas—describing their problems, concerns and suggestions in German text form— including other information like title and category. Since the data originated from heterogeneous sources, several preprocessing steps were conducted to obtain a consolidated dataset: a collection of 3903 ideas with 8 categories, with stratified sampling into a training (90% - 3,512 ideas) and a test (10% - 391 ideas) set, and ultimately the training set was randomly split into 3 slices à ~1170 ideas each, simulating the three participating cities.

The prototype implements the use case as follows. Each city is represented by a party in the FML. The parties agree that an aggregator (as an independent actor responsible for the provision of the required technology) will manage the FML process and generate a global model based on local model updates from each party. Moreover, they agree on the set of test data that will be used to benchmark each updated version of the model. Each party can train locally the models in numerous rounds and store the required data locally. Technical components such as messaging and routing of the FML communication, storage of encrypted models as well as accountability-related rules and log data (verifiable claims and evidence about training data set, number of rounds, model updates, reached model quality, configuration of each party etc.) are stored on a cloud infrastructure. For running and orchestrating the FML process, we used the IBM FL framework [36]. The component for accountability rules and data (cf. for details <https://git.fortiss.org/evidentia>) is built atop a logic database named Datalog [41] and Hyperledger Fabric [42].

4 AFML Engineering

4.1 Feasibility Evaluation for FML

In order to structure FML approaches in a comprehensive manner, we provide the following overview (cf. Table 1) based on substantial extant research [37, 43, 44]. The overview has nine dimensions with corresponding characteristics. It is rather selective and does not claim to be fully exhaustive, since its sole purpose is to provide a starting point for evaluation the feasibility of engineering FML in the government context (for additional details cf. e.g. [45]).

The dimension *data partitioning* is characterized by the type of learning implied in terms of samples (data to learn from) and feature space (individual measurable property or characteristic of the data) to consider when labeling data. *Horizontal* implies that the data samples differ between the FML parties, but the feature space is the same. For example, departments with the same jurisdiction at different federal levels or at different cities share knowledge about user preferences about consumption of a public service. *Vertical* implies that the features differ between the FML parties, but the data

space is the same. For example, departments with different jurisdictions share knowledge about relevant user attributes to create a more detailed user profile. *Hybrid* approaches concerns difference in both samples and features and are rather interesting in academic research, but have limited practical implications.

Table 1. Feasibility Evaluation of FML.

Dimension	Characteristics		
data partitioning	horizontal	vertical	hybrid
ML model	linear model	decision tree	neural network
training data input	featured		raw
training data output	structured		unstructured
data federation	cross-silo		cross-device
privacy preservation	differential privacy		cryptographic techniques
network topology	centralized		decentralized
federation need	economic incentive		regulation
technology grade	research		industry

The *machine learning model* dimension describes different types of approaches that can be applied in FML. It is important the profound understanding of the mechanics of these approaches is present at each party’s side, since the parties should agree on one approach and be able to justify potential attacks and malfunctions (cf. e.g. [17]).

The dimensions *training data input* and *output* as well as *data federation* are of particular relevance for FML in a government setup, when parties are organizations. This is the case, since *cross-device federation* typically addresses the challenge of applying FML to a large number of end devices (smart phones, IoT devices etc.), where challenges emerge out of scalability issues but the variety of data types is rather limited. In *cross-silo federation*, different organizations participate with different data. For instance, different public administrations have to agree on which data to use for building models and for which purpose. In this case, the involved parties should agree on the characteristics of data input and output and in case of heterogeneity, additional data pre- and post-processing should be introduced (cf. e.g. [44]).

Privacy preservation is typically addressed by either *differential privacy* or through *cryptographic methods* [12]. Basically, differential privacy aims at handling data in way that does not allow to reverse engineer any privacy-relevant information from the model or from queries based on the model. In practice this task is often challenging, due to the tension between robustness, fairness and privacy (cf. e.g. [17]). Cryptographic methods focus on securing that data stays private throughout the FML process, i.e. that computation is performed without revealing privacy related information. Practical and theoretical challenges include analyzing potential attacks as well as reaching the desired level of performance, since communication and processing effort is high.

The *network topology* could be centralized (e.g. aggregator) or decentralized (e.g. the model is aggregated among the parties in a peer-to-peer like setup). Practical implementations vary based on the type of data partitioning involved. A rather common topology is the centralized one combined with suitable privacy preservation.

Federation need represents a generic engineering dimension that addresses alignment with business and management. Economic incentives imply a clear efficiency or effectiveness need from the parties to collaboratively develop a model with corresponding privacy requirements. Needs that results from regulation are often derived from the legal framework and limitations regarding data sharing. Federation can be required from a mixture of needs, e.g. initial economic drivers and consideration of GDPR and jurisdiction legal norms among public administrations.

Technology grade implies the practical need to analyze and select a suitable technology stack. While research is active and ongoing, there are industry frameworks and service offerings available that take care of non-FML specific tasks such as authentication and identity management (cf. e.g. [17, 46]).

4.2 Architecture of AFML

Based on our prototype, we generalize the following architecture (cf. Figure 2) for AFML and extend the feasibility evaluation for FML with focus on practical implications. First, parties in the government domain should be capable of applying machine learning. AFML does not provide a remedy for missing machine learning pipelines and, in fact, only adjusts the way machine learning is applied. Moreover, data preprocessing is still a traditional challenge to be mastered, and in the context of AFML this might lead to additional effort when defending against attacks and malfunctions (cf. e.g. [17]).

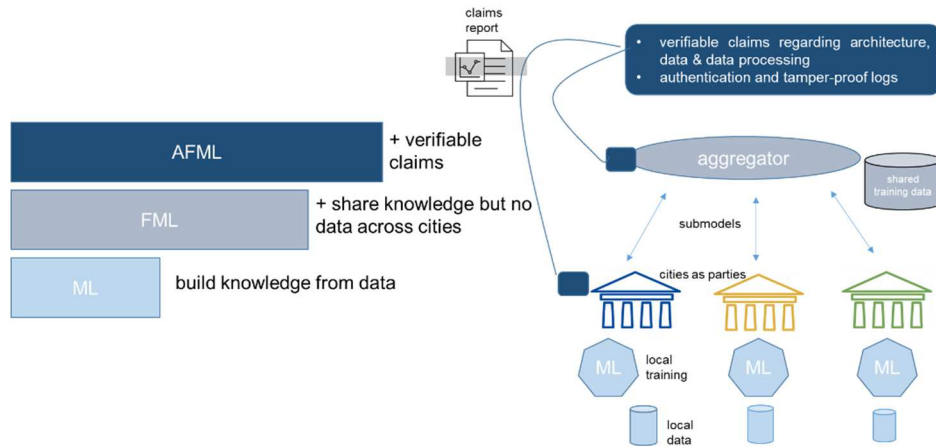


Fig. 1. A practical AFML architecture.

Second, in the context of government practice particular FML setups seem more feasible. We argue that *data partitioning would be rather horizontal, data federation would be cross-silo, network typology would be rather centralized and considered technologies should be industry or near-industry grade*. In the prototype developed, we gathered initial insights in support of our argument, since cities were willing to experiment but pointed out that a service provider should take over the FML setup and cross-device federation is currently lacking corresponding infrastructure and applications.

Based on our prototype and concept development, we also found that accountability of FML represents a potential enabler for accepting a novel technology. This is the case, since fairness, ethics and privacy are intensively discussed in the context of machine learning (cf. e.g. [3]) and, ultimately, trustworthiness of machine learning (cf. e.g. [12]), which is sometimes hampering public administrations to act.

From an engineering point of view, this challenge could be addressed with **verifiable claims**. Such claims should be defined regarding the architecture, data and data processing of FML. Examples include the framework configuration, data samples and features, model configuration and lineage along the training rounds, and, of course, continuous integration of changes (cf. e.g. [47]). Moreover, verifiable claims should be built upon corresponding authentication and ID management, as well as tamper-proof logs (cf. e.g. [21–23, 48]). Consequently, the additional feasibility requirement would be to introduce accountability and to make it accessible to a broad number of stakeholders through **claims report**.

5 AFML Management

Managing accountability for FML comprises transparent assignment and ownership of responsibilities based on rules and agreements about the expected results and obligations, facilitating to judge whether all parties have fulfilled their responsibilities. Moreover, it might comprise mechanisms to impose sanctions if obligations are not fulfilled, thereby enabling to distribute FML goals across multiple organizations.

5.1 Actors in AFML

In order to resolve particular management challenges, it is important to study which actors are involved in FML besides parties and aggregators. Based on our analysis, we distinguish between the following actors [47]: supplier (also referred to as producer), deployer (consumer), aggregator, party and auditor (cf. Figure 2).

The supplier is the entity which owns the process of training a FML model. The supplier is responsible for prescribing global training parameters, the model architecture, the FML protocol and fusion algorithm, and specific data handlers. The supplier owns the trained FML model and provides it (e.g. through software licensing) to the deployer.

The deployer is the entity which controls the usage, risks and benefits of the trained FML model. The aggregator is responsible for aggregating the FML model updates provided by the parties, adhering to relevant supplier’s prescriptions, and for making the trained FML model accessible to the supplier. The parties are responsible for providing the aggregator with FML model updates, adhering to relevant supplier’s prescriptions. The auditor is an independent (potentially accredited) body which verifies and/or certifies that the deployment of the FML model (and/or the FML model itself) adheres to technical standards and/or applicable governance, risk & compliance (GRC) obligations.

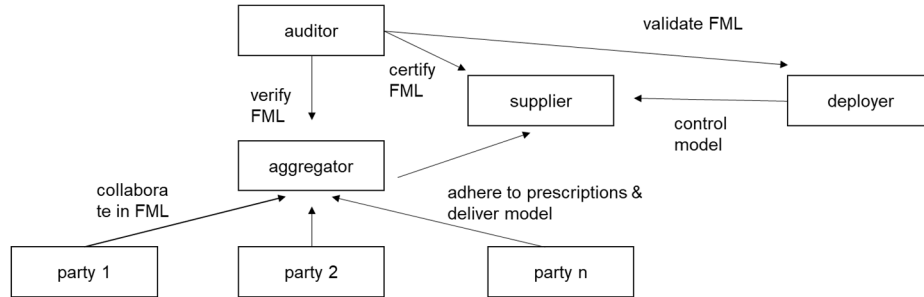


Fig. 2. Overview of actors in AFML.

5.2 Trust between parties

Another aspect of addressing management challenges of AFML is the mapping of rules and agreements between the parties according to their responsibilities and obligations. Since the parties are different organization, from a management perspective FML activities takes place in setup where different organizations carry out activities jointly to achieve a common business objective. In particular, compliance challenges of linking regulation with governance, business processes and their execution across organizations pose threats to coping with responsibilities and obligations [27]. In order to cope with the latter, a corresponding level of trust is required. Achieving trust is challenging in the case of FML due to potentially unequal power relations (e.g. who provides the data, who is interested in the model) and the possibility of protecting the interests of individual organizations (e.g. what is a “good model”) by manipulating and controlling collaborative learning processes (e.g. attacks and/ or malfunction).

As a remedy, accountability in management of FML would include introducing **verifiable claims** that link (i) compliance with prescribed and/or agreed norms (cf. e.g. [25–27]), control mechanism for the parties involved in a process (cf. e.g. [24, 49]) as well as (iii) a regulatory framework. Such claims should represent the basis for operationalizing trust among the actors in a satisfactory manner. In particular, this would result in a non-repudiable consensus regarding the actors’ judgement on the completion of the process in a verifiable way (cf. e.g. [24, 25]).

Claim reports could represent a feasible artefact that operationalizes accountability. The purpose of such reports (sometimes referred to as factsheets [47, 50]) is to provide transparency and instill trust into ML services. They are to be “completed by AI service providers for examination by consumers” and shall contain “sections on all relevant attributes of an AI service”, in particular “how the service was created, trained, and deployed” [47]. Yet, it is still an open issue what level of detail and which claims are most suitable for a report to adhere to a required level of trust.

6 Discussion and Conclusion

In this paper, we formulated the need to face challenges of accountability, data sharing and privacy preservation along the course of machine learning in a government context.

We address this need by introducing a novel approach named federated machine learning, which relaxes limitations of data sharing and privacy constraints. Moreover, we address the formulated need by introducing accountability from an engineering and management perspective towards generating verifiable claims for FML.

Through an argumentative-deductive analysis of literature and a prototype of AFML for online citizen participation, we explored AFML from an engineering and a management perspective. The engineering perspective includes feasibility evaluation of FML and adds an accountability perspective based on a corresponding architecture for practical applications in government. The management perspective includes an analysis of actors involved in AFML and means to establish trustworthiness between them.

Based on this analysis framework, we approach the question of introducing AFML in the government domain based on the following overview of standardization artefacts (cf. Table 2, [7, 28]). Exemplary artefacts in bold show that substantial progress has already been made or industry-ready solutions already exist, which is promising for exploring improvements of existing approaches (cf. e.g. [39]). Regarding the other exemplary artefacts, the status is either an open research problem or the current solutions are for application in a research setup.

Table 2. Implications for streamlining AFML in government.

	Administration	Modeling	Processing	Communication & Interaction	Security & Privacy
Organizational / Semantic	governance of incentives vs. regulations	lifecycle blueprint	FML training integration	enterprise infrastructure integration	compliance
Technical / Syntactic	claim report semantics	trust semantics	explainability	data & model metadata	guarantees for attacks and threats
	evidence granularity & tamper-proof guarantees	common accountability criteria	toolchain	tool & model interoperability	cryptology & differential privacy, ID mgmt

Our research has a number of limitations. First, the engineering analysis is rather general and omits details that might be of relevance for a thorough feasibility evaluation in practice, especially from a methodological perspective given the fact that we developed the prototype and used it as a basis for interpretation. Second, the presented architecture is focused on a cross-silo data federation. Emerging developments (e.g. smart city, edge computing) might pose the need for a cross-device FML in government, which is out of our research scope. Third, the management analysis was solely based on argumentation and deduction from relevant literature as well as prototype development, due to the novelty of FML and the limited access to suitable interviewees in the

government domain for sampling of primary empirical data. Fourth, a limitation for ML in general is a possible security and privacy breach by reverse engineering a model, which might leak the data.

We believe that future research should build on our findings and address the describing limitations of our research. We strongly encourage researchers to explore potential use cases and to derive engineering and management requirements for AFML. We also believe that practitioners can directly benefit from the presented findings and apply them as a basis for exploring novel FML techniques to overcome traditional challenges.

Acknowledgements

This research was partially funded by the Bavarian Ministry of Ministry of Economic Affairs, Regional Development and Energy in the context of the project BayernCloud (funding code 'AZ: 20-13-3410.1-01A-2017').

We thank our reviewers for their careful reading and their constructive remarks.

References

1. Russell, S.J., Norvig, P.: Artificial intelligence: a modern approach. Malaysia; Pearson Education Limited, (2016).
2. Agrawal, A., Gans, J., Goldfarb, A.: Prediction Machines: The simple economics of artificial intelligence. Harvard Business Press (2018).
3. Winfield, A.F., Michael, K., Pitt, J., Evers, V.: Machine ethics: The design and governance of ethical AI and autonomous systems [scanning the issue]. *Proceedings of the IEEE*. 107, 509–517 (2019).
4. Dwivedi, Y.K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A.: Artificial Intelligence (AI): Multi-disciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*. 101994 (2019).
5. Sun, T.Q., Medaglia, R.: Mapping the challenges of Artificial Intelligence in the public sector: Evidence from public healthcare. *Government Information Quarterly*. (2018).
6. Council, A.U.P.P.: Statement on algorithmic transparency and accountability. *Commun. ACM*. (2017).
7. Scholl, H.J., Klischewski, R.: E-government integration and interoperability: framing the research agenda. *International Journal of Public Administration*. 30, 889–920 (2007).
8. Wang, F.: Understanding the dynamic mechanism of interagency government data sharing. *Government Information Quarterly*. 35, 536–546 (2018).
9. Janssen, M., van den Hoven, J.: Big and Open Linked Data (BOLD) in government: A challenge to transparency and privacy? *Government Information Quarterly*. 32, 363–368 (2015). <https://doi.org/10.1016/j.giq.2015.11.007>.

10. Howard, A., Borenstein, J.: The ugly truth about ourselves and our robot creations: the problem of bias and social inequity. *Science and engineering ethics*. 24, 1521–1536 (2018).
11. Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., Floridi, L.: Artificial intelligence and the ‘good society’: the US, EU, and UK approach. *Science and engineering ethics*. 24, 505–528 (2018).
12. Brundage, M., Avin, S., Wang, J., Belfield, H., Krueger, G., Hadfield, G., Khlaaf, H., Yang, J., Toner, H., Fong, R.: Toward trustworthy AI development: mechanisms for supporting verifiable claims. *arXiv preprint arXiv:2004.07213*. (2020).
13. Scholta, H., Niemann, M., Halsbenning, S., Räckers, M., Becker, J.: Fast and Federal—Policies for Next-Generation Federalism in Germany. (2019).
14. Scholta, H., Balta, D., Räckers, M., Becker, J., Krcmar, H.: Standardization of forms in governments. *Business & Information Systems Engineering*. 62, 535–560 (2020).
15. Küsters, R., Truderung, T., Vogt, A.: Accountability: definition and relationship to verifiability. In: *Proceedings of the 17th ACM conference on Computer and communications security*. pp. 526–535. ACM (2010).
16. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: *Artificial Intelligence and Statistics*. pp. 1273–1282. PMLR (2017).
17. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R.: Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*. (2019).
18. Bovens, M.: Analysing and assessing accountability: A conceptual framework 1. *European law journal*. 13, 447–468 (2007).
19. Eriksén, S.: Designing for accountability. In: *Proceedings of the second Nordic conference on Human-computer interaction*. pp. 177–186. ACM (2002).
20. Bloomfield, R., Rushby, J.: Assurance 2.0: A Manifesto. *arXiv preprint arXiv:2004.10474*. (2020).
21. Cleland-Huang, J., Gotel, O., Zisman, A.: *Software and systems traceability*. Springer (2012).
22. Cleland-Huang, J., Gotel, O.C., Huffman Hayes, J., Mäder, P., Zisman, A.: Software traceability: trends and future directions. In: *Future of Software Engineering Proceedings*. pp. 55–69 (2014).
23. Gotel, O., Cleland-Huang, J., Hayes, J.H., Zisman, A., Egyed, A., Grünbacher, P., Dekhtyar, A., Antoniol, G., Maletic, J., Mäder, P.: Traceability fundamentals. In: *Software and systems traceability*. pp. 3–22. Springer (2012).
24. Baldoni, M., Baroglio, C., Micalizio, R., Tedeschi, S.: Accountability and Responsibility in Business Processes via Agent Technology. In: *Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion (RCRA 2018)*. pp. 1–18. CEUR-WS (2018).
25. Yao, J., Chen, S., Levy, D.: Accountability-based compliance control of collaborative business processes in cloud systems. In: *Security, Privacy and Trust in Cloud Systems*. pp. 345–374. Springer (2014).

26. Sadiq, S., Governatori, G., Namiri, K.: Modeling control objectives for business process compliance. In: International conference on business process management. pp. 149–164. Springer (2007).
27. Hashmi, M., Governatori, G., Lam, H.-P., Wynn, M.T.: Are we done with business process compliance: state of the art and challenges ahead. *Knowledge and Information Systems*. 57, 79–133 (2018).
28. Balta, D., Kuhn, P., Sellami, M., Kulus, D., Lieven, C., Kremar, H.: How to streamline AI application in government? A case study on citizen participation in Germany. In: International Conference on Electronic Government. pp. 233–247. Springer (2019).
29. Balta, D.: Effective Management of Standardizing in E-Government. *Corporate Standardization Management and Innovation*. 149–175 (2019). <https://doi.org/10.4018/978-1-5225-9008-8.ch008>.
30. Balta, D., Kremar, H.: Managing Standardization in eGovernment: A Coordination Theory based Analysis Framework. In: Parycek, P., Glassey, O., Janssen, M., Scholl, H.J., Tambouris, E., Kalampokis, E., and Virkar, S. (eds.) *Electronic Government*. pp. 60–72. Springer International Publishing (2018).
31. Gregor, S.: The nature of theory in information systems. *MIS quarterly*. 611–642 (2006).
32. Goldkuhl, G.: Pragmatism vs interpretivism in qualitative information systems research. *European journal of information systems*. 21, 135–146 (2012).
33. Wilde, T., Hess, T.: Forschungsmethoden der Wirtschaftsinformatik. *Wirtsch. Inform.* 49, 280–287 (2007). <https://doi.org/10.1007/s11576-007-0064-z>.
34. Nissenbaum, H.: Computing and accountability, <https://link.gale-group.com/apps/doc/A15020194/AONE?sid=lms>, last accessed 2019/10/06.
35. Beckers, K., Landthaler, J., Matthes, F., Pretschner, A., Walzl, B.: Data accountability in socio-technical systems. In: *Enterprise, Business-Process and Information Systems Modeling*. pp. 335–348. Springer (2016).
36. Ludwig, H., Baracaldo, N., Thomas, G., Zhou, Y., Anwar, A., Rajamoni, S., Ong, Y., Radhakrishnan, J., Verma, A., Sinn, M., Purcell, M., Rawat, A., Minh, T., Holohan, N., Chakraborty, S., Whitherspoon, S., Steuer, D., Wynter, L., Hassan, H., Laguna, S., Yurochkin, M., Agarwal, M., Chuba, E., Abay, A.: IBM Federated Learning: an Enterprise Framework White Paper V0.1. arXiv:2007.10987 [cs]. (2020).
37. Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., He, B.: A survey on federated learning systems: vision, hype and reality for data privacy and protection. arXiv preprint arXiv:1907.09693. (2019).
38. Susha, I., Grönlund, A.: eParticipation research: Systematizing the field. *Government Information Quarterly*. 29, 373–382 (2012).
39. Androustopoulos, A., Karacapilidis, N., Loukis, E., Charalabidis, Y.: Transforming the communication between citizens and government through AI-guided chatbots. *Government Information Quarterly*. (2018).
40. Maragoudakis, M., Loukis, E., Charalabidis, Y.: A review of opinion mining methods for analyzing citizens’ contributions in public policy debate. In: International Conference on Electronic Participation. pp. 298–313. Springer (2011).

41. Greco, S., Molinaro, C.: Datalog and logic databases. *Synthesis Lectures on Data Management*. 7, 1–169 (2015).
42. Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., Manevich, Y.: Hyperledger fabric: a distributed operating system for permissioned blockchains. In: *Proceedings of the thirteenth EuroSys conference*. pp. 1–15 (2018).
43. Mothukuri, V., Parizi, R.M., Pouriyeh, S., Huang, Y., Dehghantanha, A., Srivastava, G.: A survey on security and privacy of federated learning. *Future Generation Computer Systems*. 115, 619–640 (2021).
44. Verma, D., White, G., de Mel, G.: Federated AI for the enterprise: A web services based implementation. In: *2019 IEEE International Conference on Web Services (ICWS)*. pp. 20–27. IEEE (2019).
45. Song, L., Wu, H., Ruan, W., Han, W.: SoK: Training Machine Learning Models over Multiple Sources with Privacy Preservation. *arXiv preprint arXiv:2012.03386*. (2020).
46. Torzadehmahani, R., Nasirigerdeh, R., Blumenthal, D.B., Kacprowski, T., List, M., Matschinske, J., Späth, J., Wenke, N.K., Bihari, B., Frisch, T.: Privacy-preserving Artificial Intelligence Techniques in Biomedicine. *arXiv preprint arXiv:2007.11621*. (2020).
47. Arnold, M., Bellamy, R.K., Hind, M., Houde, S., Mehta, S., Mojsilović, A., Nair, R., Ramamurthy, K.N., Olteanu, A., Piorkowski, D.: FactSheets: Increasing trust in AI services through supplier’s declarations of conformity. *IBM Journal of Research and Development*. 63, 6–1 (2019).
48. Spanoudakis, G., Zisman, A.: Software traceability: a roadmap. In: *Handbook Of Software Engineering And Knowledge Engineering: Vol 3: Recent Advances*. pp. 395–428. World Scientific (2005).
49. Baldoni, M., Baroglio, C., May, K.M., Micalizio, R., Tedeschi, S.: Computational accountability. In: *Deep Understanding and Reasoning: A Challenge for Next-generation Intelligent Agents, URANIA 2016*. pp. 56–62. *CEUR Workshop Proceedings* (2016).
50. Piorkowski, D., González, D., Richards, J., Houde, S.: Towards evaluating and eliciting high-quality documentation for intelligent systems. *arXiv:2011.08774 [cs]*. (2020).