

COMMENTARY

The governance of federated learning: a decision framework for organisational archetypes

Tom Barbereau^{1,2} , Joaquin Delgado Fernandez³  and Sergio Potenciano Menci³ 

¹Dutch Organization for Applied Scientific Research (TNO), The Hague, The Netherlands

²Institute for Information Law (IViR), University of Amsterdam, Amsterdam, The Netherlands

³Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Esch-sur-Alzette, Luxembourg

Corresponding author: Tom Barbereau; Email: tom.barbereau@tno.nl

Received: 26 November 2024; **Revised:** 25 May 2025; **Accepted:** 14 June 2025

Keywords: archetypes; artificial intelligence; federated learning; Governance

Abstract

Federated learning (FL) is a machine learning technique that distributes model training to multiple clients while allowing clients to keep their data local. Although the technique allows one to break free from data silos keeping data local, to coordinate such distributed training, it requires an orchestrator, usually a central server. Consequently, organisational issues of governance might arise and hinder its adoption in both competitive and collaborative markets for data. In particular, the question of how to govern FL applications is recurring for practitioners. This research commentary addresses this important issue by inductively proposing a layered decision framework to derive organisational archetypes for FL's governance. The inductive approach is based on an expert workshop and post-workshop interviews with specialists and practitioners, as well as the consideration of real-world applications. Our proposed framework assumes decision-making occurs within a black box that contains three formal layers: data market, infrastructure, and ownership. Our framework allows us to map organisational archetypes ex-ante. We identify two key archetypes: consortia for collaborative markets and in-house deployment for competitive settings. We conclude by providing managerial implications and proposing research directions that are especially relevant to interdisciplinary and cross-sectional disciplines, including organisational and administrative science, information systems research, and engineering.

Policy Significance Statement

This commentary proposes a framework that identifies decision-making layers leading to an organisational archetype for governance in federated learning. What type of entity controls the server in such systems and orchestrates clients is a pending question that hinders adoption. The proposed framework allows for the ex-ante selection of the appropriate organisation for each setting. We apply the framework on the basis of real-world applications.

1. Introduction

The integration of artificial intelligence (AI) is set to bring about unprecedented changes to industries, public institutions, and civil society (Dwivedi et al., 2021). However, data of sufficient quality and quantity are not always available. It is costly to acquire, requires advanced capabilities to process, and is often restricted to organisational boundaries due to competition or regulation (Winter et al., 2014; Jordan

& Mitchell, 2015; Berente et al., 2021). These questions become especially relevant when data could leave the boundaries of a singular organisation and collaboration occurs, because it is here, where the “adoption of federated learning (FL) [...] is expected to have a catalytic impact towards precision” (Sheller et al., 2019).

Introduced in 2016, FL addresses concerns related to data sharing (Konečný et al., 2016). FL is a machine learning (ML) technique that allows for the training of a model across multiple decentralised devices or servers holding local data samples without exchanging them (so-called *clients*). In other words, data remain in its original location, within organisational boundaries or on a device. The model learns from the data locally and sends the updates back to (usually) a central server. That server aggregates these updates to improve the global model. This local learning and aggregation process is iterative until a performant model is trained and distributed. The processes can be improved by the addition of privacy-enhancing methods and technologies (Truong et al., 2021). In sum, FL allows organisations to collaborate in training a model without sharing data across their own boundaries. Given this feat, beyond the device-centric applications of Google (see Hard et al., 2018), FL finds promising application in (secondary) health data use in oncology (Dayan et al., 2021; Pati et al., 2022, in energy demand and short-term load forecasting (Fernández et al., 2022; Fernández et al., 2023), in governmental data sharing for predictions (Amard et al., 2023; Sprenkamp et al., 2024), in financial credit risk assessment and fraud prevention (Lee et al., 2023; Fernández et al., 2024), as well as numerous other domains.

However, whoever controls the central server in the training process wields considerable power. That is because the party acts as the *orchestrator* of the learning process and coordinates between the training clients (Pati et al., 2022; Bujotzek et al., 2024). It is also responsible for maintaining regulatory requirements, upholding data privacy, and system security during the entire learning process. The orchestrator must ensure that the learning process is fair and that the global model is not biased towards any particular type of data or device. It may also have purview over questions of commercialisation and ownership in the form of intellectual property (IP) rights. Broadly, in the study of AI more generally, these responsibilities fall under what is understood as questions of “governance” (Berente et al., 2021).

Good governance is essential for organisations dealing with technology (Weill & Ross, 2004). AI is no exception to it: governance is a necessary prerequisite to reap sustained benefits (Berente et al., 2021; Zhang, 2023). Governance is about aligning the affordances of a technology or data with organisational goals; and because this alignment is context-based on asset and organisation, governance is never a one-size-fits-all (Weber et al., 2009; Khatri & Brown, 2010). In the case of FL, and to be expected given the relatively low level of adoption, the knowledge base on its governance is scarce. Critical questions towards its adoption arise: how to convince clients to participate? Who should control the aggregation process? How to arrange ownership of the trained model? What type of organisation orchestrates?

This commentary engages with these questions. In particular, it considers what type of organisation orchestrates FL, and what *common* decision-making layers lead to the choice for that organisation. Given the scarcity of knowledge, we approach this question ex-ante and rely on knowledge gathered by hosting a workshop on the subject (Storvang et al., 2017) and discussing with experts (Mergel et al., 2019).¹ In total, we convened with 15 experts from the health, financial, and energy sectors to discuss questions related to decision-making in the governance of FL. The result of this process, presented in Section 3, yields a layered decision framework for the organisational governance of FL. In Section 4, we apply this framework to real-world applications—notably, Federated Tumor Segmentation in healthcare (Pati et al., 2022), where we identify the consortia archetype, and data aggregation by Google (Hard et al., 2018), where we identify the in-house archetype—and delineate their limitations. Finally, in Section 5, we discuss and propose future research directions that are relevant to interdisciplinary and cross-sectional disciplines. Specifically, we foresee researchers engaging with technological developments in FL and considering their implications for governance.

¹ Additionally, we ought to mention that all authors are actively working on applied FL projects.

2. Background

In FL, the peculiarity is that while training takes place at the edges, a central orchestrating entity is still present. The presence of that entity has technical and organisational implications.

In technical terms, the orchestrator is responsible for the central server, or rather the training process, which aggregates updates from decentralised clients that locally process data, and then distributes the combined model back to these clients. Thus, the orchestrator will receive the model trained by the clients (or the delta) from the previous iteration and aggregate it before transmitting it back to the clients (McMahan et al., 2017).² Consequently, the orchestrator plays a crucial role, as it ensures the integration of diverse, local insights into a global model without directly sharing or aggregating data. Doing so requires data curation, standardisation of reference labels, and the formalisation of workflows, among others³. Furthermore, the orchestrator is also responsible for the aggregation mechanism, is in control of the update schedule, and handles security and privacy aspects (Yin et al., 2021). Thereby, it holds the capability to affect the accuracy, bias, and overall integrity of the system. The organisation governing the central server wields significant power (Pati et al., 2022; Bujotzek et al., 2024).

From an organisational perspective, in the case that FL is collaborative and outside of the boundaries of a single organisation (as is the case in, e.g., applications in healthcare under a consortium; Pati et al., 2022), the “structure” acting as the orchestrator must be trustworthy. This organisation—beyond technical tasks described previously—is also tasked with aligning participating clients to a set of agreed-upon rules and contracts, settling disputes, and ensuring the financial sustainability of the collaboration. These points each consider the “locus of accountability”—that is, here, some type of organisation, the one “who makes the decision” (Khatri & Brown, 2010). All participants must trust this accountable organisation *before* any collaboration and sharing of assets occur (Berente et al., 2021; Pati et al., 2022; Bujotzek et al., 2024). This is essential in order to create an environment where all parties are willing to collaborate in the short term, and benefit from shared insights or financial gains in the long term. An alternative case is when FL is deployed inside the boundaries of an accountable organisation, in which case governance is done in-house (see, e.g., Hard et al., 2018).

At last, it is noteworthy that the organisational structure for governance must be designed according to the context in which FL is trained. Technology governance is never a one-size-fits-all and depends on the asset plus organisational structure at hand (Brown & Grant, 2005; Weber et al., 2009; Khatri & Brown, 2010). In the case of training FL across devices by a single organisation, governance typically happens in-house (see McMahan et al., 2017; Hard et al., 2018). When between organisations, in a non-competitive and research-centric market for data, a consortium might be the right fit (see Mateus et al., 2024). Alternatively, a single party—such as a university or non-profit—could be trusted to act altruistically as a third party and reap shared rewards (see Pati et al., 2022).

Taking past understandings on designing data governance and decision domains (Khatri & Brown, 2010), we formulate a set of non-exhaustive questions related to and informing model governance (see Table 1).⁴ They are categorised in terms of decision domains appropriated from governance literature. Formally, Weill and Ross (2004) propose that IT governance design includes five major decision domains—principles, architecture, infrastructure, application needs, and investment and prioritisation. Khatri and Brown (2010) adapt these to data governance to consider principles, quality, metadata, access, and the life cycle. We adhere to and adapt the latter as we consider data to be the core asset.

The practical orientation of this article compels us to consider the framing of *governance* by Pati et al. (2022). Their study presents results from training an FL model for tumour detection involving data from 71 sites across 6 continents. Here, governance is referred to as (1) the definition of the problem statement and (2) the coordination with the collaborating sites. In consideration of this framing, the scope of this commentary is on the latter: the coordination (i.e., model governance) with collaborators. The goal of this

² Usually, this aggregation is a weighted arithmetic mean or a simple arithmetic mean.

³ By this, we understand the standardisation of model outputs, labels the models are trained against, and standardisation of the weights across the clients.

⁴ The technical and organisational implications are accounted for as part of the formulated questions.

Table 1. *Decision domains and questions tailored to FL governance, based on Khatri and Brown (2010)*

Governance domains	Domain decisions ^a
Principles	What are the system technicalities (e.g., architecture, rounds, PETs ^b , data structure, communication channels)? How does the regulatory environment(s) influence system design? (e.g., Is the data sensitive?) If any, how is client participation (financially) compensated over time? If any, what is the strategy for re-sharing (i.e., the terms of licensing and ownership)? What happens when the learning threshold is reached?
Quality	Does the FL process require a data quality check? How does the quality check process look? Where do the quality checks occur (i.e., orchestrator, clients, or both)?
Metadata	What information (e.g., versioning and identifiers) is in the metadata? What information does the metadata require for accountability? For compliance?
Access	Does anyone else, besides the clients, have model access? For how long do clients have model access? On what terms? If any, what is the access pricing strategy (i.e., subscription, openness, and freemium)?
Life cycle	Who stores the model? Where is the model stored? What is the long-term model retention strategy?

^aThe questions formulated by domain are non-exhaustive.

^bPrivacy enhancing techniques (see Yin et al., 2021).

commentary is to advance understanding of the *common layers leading to an organisational archetype for the governance of FL*. After delineating assumptions, describing our inductive approach, we propose the framework and its layers. At last, we apply the model to real-world cases.

3. Towards a framework of common layers

3.1. Assumptions

The development of information systems is a process that is constructivist (Weigl et al., 2023). Decision-making about technology in organisations involves actors that are involved in a “set of games” (Mintzberg, 1983) and “processes of getting commitment” (Keen, 1981). It is to be expected that actors who contributed more data to a model, a crucial piece to the model, or who simply have a more dominant market share can (ab-)use their “positional advantage” (Hård, 1993) in negotiations over technology. In particular, this is expected to be the case in competitive data markets (see Lee et al., 2023). Because the set of games involves actors whose motivations may be “hidden” (Grover et al., 1988), the assumption is that decisions over governance occur in a black box.

Conceptually, one may open this black box from a constructivist perspective. Given constraints in the technical domain, technology’s affordances are interpreted flexibility by different actors. These actors make decisions within some institutional context, based on the organisational status quo (operating procedures, culture, etc.) and environmental pressures (regulation, financing, etc.) (Weigl et al., 2023). Although we find value in the described perspective, given the practical orientation of this commentary, an alternative, inductive lens is put forward.

We assume that it is *never* possible to fully deconstruct the black box of decision-making (Winner, 1993). As each context is different, generalisation of any deconstruction would be limited (Weber et al., 2009). Given this assumption and the consequent limits to generalisation, based on our process of

information gathering, we instead provide a framework that considers three *common* layers—*quasi*, context factors—that lead to an organisational archetype for the governance of FL.

3.2. Inductive approach

The reflection and layered framework we propose is inductively developed by following principles of conceptual research (Jaakkola, 2020; Makowski, 2021). In September 2024, two of the three co-authors attended the 22nd International Conference on Intelligent Systems Applications to Power Systems in Budapest, Hungary. There, they discussed organisational archetypes for the governance of FL as part of a themed panel and brought back input for the creation of a framework. Workshops are commonly used as an inductive approach to research (Storvang et al., 2017). The workshop included academics and practitioners active in the field of FL. Some were operationally involved in the deployment of solutions; others were active in more fundamental research.

Between September and August 2024, we gathered post-workshop reflections and feedback with FL specialists from various industries (n.b. energy, finance, and defence). They are considered agents with implicit and factual knowledge about processes and decisions (Mergel et al., 2019). Among authors and with some of the specialists, we repeatedly met to review and iterate (Klein & Myers, 1999) until the final framework.

After the framework was developed, we performed an ex-ante application of it to a set of real-world applications identified using a modified version of an online open-access tool to do automated searches in different databases of academic literature (Gerloff, 2022). We primarily consulted Association for Computing Machinery (ACM), Institute of Electrical and Electronics Engineers (IEEE), and Scopus for technical literature, as well as PubMed, given that the main empirical results stem from there. We also consulted non-peer-reviewed academic works from arXiv, Medrxiv, and Biorxiv, although we do not refer to these specifically.

3.3. Layered framework

The result of this process of information gathering and iteration is the formulation of a conceptual framework of common layers to derive organisational archetypes for the governance of FL (Figure 1). The three common layers at play when deciding upon the locus of accountability and corresponding organisational archetype are (1) the nature of the data market (spectrum of collaborative–competitive), (2) the FL model’s architecture (spectrum of centralised–decentralised), and (3) the model’s ownership (spectrum of proprietary–shared). The result is archetypes.

While in no order of priority, there is a temporal element to these layers such that the initial consideration is done with regard to the market, followed by architectural design choices, and, at last, the definition of the ownership structure. We subsequently discuss each of the three common layers individually. Thereafter, we put forward a set of non-exhaustive questions to help decide on the archetype.




Figure 1. A layered framework to derive organisational archetypes for the governance of FL. The line represents the virtual decision-making process throughout each of the layers, resulting in a selected archetype.

3.3.1. Data market

First, stakeholders must evaluate the data market in which they are operating. There are four types of data markets: many-to-many, one-to-many, many-to-one, and one-to-one (Driessen et al., 2022). However, in stylised economic terms, we can simplify markets for data and place them on a spectrum between competitive and collaborative markets (Fernández, 2023). Competitive markets are characterised by organisations that guard their data closely to maintain a competitive edge. Legal restrictions and other market pressures also play a role in the limits of data sharing. Organisations will avoid sharing data that could potentially benefit their rivals, as their competitive advantage hinges on exclusive access to valuable data resources (Kearns & Lederer, 2004).

However, organisations have begun to shift towards collaborative business models and technologies to maximise returns. At once, data markets can be “inherently” collaborative such that organisations actively work together to share data, research on projects, and reap mutual benefits (Spiekermann, 2019). Examples include hospitals and research centres collaborating to create comprehensive datasets for better patient care.⁵ Other organisations, such as banks and insurance companies, will traditionally compete, although they will look to collaborate to address a shared problem and reap joint rewards (Lee et al., 2023).

3.3.2. Infrastructure

Second, stakeholders must decide upon the system architecture. FL presents two main architectures (Figure 2): centralised and decentralised (Yin et al., 2021). The former involves a set of clients and an orchestration agent (central server) for collaborative training of an ML model. The process consists of three phases as follows: preparation, training and learning, and use. Clients prepare their data, agree on the ML model to train, and select initialisation parameters. The orchestrator then selects a subset of clients to download, train, and share their model parameters. This process iterates until the model reaches the agreed objective, usually a certain performance level or learning threshold, although it could be a limit in rounds or training expenses. This architecture has trade-offs, including privacy concerns, a single point of failure (the server), communication overhead, scalability issues, security risks, and fairness challenges.

Decentralised FL operates differently and overcomes some of these trade-offs. It requires a protocol for peer-to-peer information sharing and a logic architecture for learning orchestration. The same three phases persist: prepare, train and learn, and use. By removing the dependency on a central coordination agent and distributing responsibility among clients, decentralised FL fosters a more resilient and efficient collaboration (Martínez Beltrán et al., 2023). It also removes the sharing of potentially compromising data with a coordination agent, avoids a single point of failure, and aims to improve fairness by shifting importance from the client’s data to every client, preventing clients with large datasets from dominating the learning process (Witt et al., 2023).

In certain scenarios (involving highly sensitive or private data), an alternative take on the centralised architecture emerges, such that the physical location of the training process changes. Under specific conditions, the federation may choose to aggregate the model at a trusted third party (such as a European body like eu-LISA or a market authority) that has no economic interest in the outcome but is responsible for ensuring the correct functioning of the system. For instance, in a fraud detection federation across banks (Yang et al., 2019), a financial market regulator could oversee the training process to ensure it is conducted properly, without having any interest in the model or its outputs.

Besides physical solutions to trust and ownership, researchers have implemented technological solutions. To uphold privacy, mechanisms like differential privacy add noise to obscure individual data contributions and mitigate reconstruction risks (Fu et al., 2024). Homomorphic encryption enables computations directly on encrypted data, preserving confidentiality during aggregation without sacrificing model utility (Jin et al., 2023). Similarly, secure multi-party computation

⁵ Admittedly, they are now more restricted in doing so, given privacy regulations (Price & Cohen, 2019).




Figure 2. Conceptual architectures for FL.

protocols allow for collaborative model training while ensuring no individual party's data are exposed (Kaminaga et al., 2023).

To protect against adversarial behaviour, especially model poisoning, anomaly detection techniques have been developed to identify irregular or malicious updates (Vucovich et al., 2022). Trust and reputation-based systems can further reinforce robustness by weighting client updates according to their historical reliability and contribution quality (Rashid et al., 2025). Furthermore, research has proposed to integrate blockchain-based solutions, such as smart contracts (Cassano et al., 2024), and non-technical approaches utilising legal instruments like contracts and independent audits to ensure ethical, legal, and technical robustness in systems, thereby holding providers accountable (OECD, 2016).

Other emerging methods seek to enhance fairness and operational efficiency. Fairness-aware aggregation frameworks attempt to balance the influence of clients with heterogeneous data volumes, ensuring equitable model contributions (Ezzeldin et al., 2023), while model compression techniques reduce the communication burden associated with transmitting large updates, thereby broadening participation to include clients with limited resources (Yang et al., 2022). Collectively, these techniques enable the development of federated systems that are not only decentralised but also more secure, private, and equitable.

3.3.3. Ownership

Finally, the stakeholders need to address questions about ownership. In the context of AI, these questions are manifold. Post-GenAI, the body of works considering the ownership of outputs of an AI has grown (see Tzimas, 2021). Conceptually, we pragmatically consider ownership as who owns the IP rights to the

trained model and how the benefits derived are distributed (Berente et al., 2021). The benefits can be both financial in the commercialisation of the model or practical in its application.

The IP can be shared or in whole; correspondingly, ownership can be shared equally among parties, be distributed among parties, or be owned by one single party. In the former, parties can—for example, in the European Union—draw up a so-called Research and Development Agreement before development (European Commission, 2023). The agreement covers (1) IP ownership and access rights, (2) registration protocols, (3) exploitation strategy (in the form of licences), and (4) IP management in case of termination of the agreement. An alternative form of shared ownership can be formalised by the creation of a joint venture or consortium, where two or more independent organisations undertake a specific project together and share the IP accordingly. Here too, for both types of organisations, agreements typically cover the same four points (see also WIPO, 2005). By and large, IP rights define the ownership settings and access to the model. For example, whereas in-house governance will allow to fully leverage the IP rights, sharing IP as part of a consortium may put limits on commercialisation.

3.3.4. *Archetype*

Due to the limited number of scaled applications, the knowledge on resulting archetypes is scarce. However, we expect one to use our framework and for one archetype to emerge as a result of decisions taken along each layer of the framework. To assist practitioners in identifying the matching archetype, we propose a set of guiding questions along the layers (see Table 2). Effectively, the way one answers these questions will determine internal development choices, ultimately leading to a specific archetype. In the future, the patterns for which archetypes are more popular in which context may become clearer. We anticipate that context, application domains, verticals, country-specific regulations, and other variables to play a significant role in shaping them.

4. **Organisational archetypes**

4.1. *Application of the framework*

In Table 3, we consider a limited set of real-world applications through the lens of our framework. It is noteworthy that we only selected applications that are mature and have progressed beyond initial proofs of concept, prototypes, or small-scale pilots and are actively utilised in real-world environments, of which, according to the UK’s Department for Science, Innovation, and Technology (2024), there are relatively few of them. We expected this to be the case given the general, risk-averse sentiment towards FL adoption (Müller et al., 2024).

We note that the majority of applications are in the healthcare sector. This was to be expected (Rieke et al., 2020). In terms of archetypes, we observe that in collaborative data markets with decentralised

Table 2. *Guiding questions by layer of our proposed framework*

Layer	Guiding questions ^a
Data market	In what data market type am I participating? What is my relation towards the orchestration and competitors/collaborators (i.e., other clients)?
Infrastructure	What type of data do I have? What type of data do others have? Is my data subject to specific regulations? What regulations am I and other clients subject to?
Ownership	Do I want to maintain ownership of the model? If yes, what type of ownership? Do I expect short-term or long-term benefits?

^aThe questions formulated by layer are non-exhaustive. We see these as starting point for discussion.

Table 3. Real-world applications of FL through the lens of the layered framework

Real-world application	Type	TRL ^a	Layers			
			Data market ^b	Infrastructure ^c	Ownership	Archetype
FeTS (Pati et al., 2022)	Cross-silo	9	Coll.	Dec.	Shared	Consortium
Gboard (Hard et al., 2018)	Cross-device	9	Comp.	Cen.	Proprietary	In-house
NCDC (Mateus et al., 2024)	Cross-silo	9	Coll.	Dec.	Shared	Consortium
RACoon (Bujotzek et al., 2024)	Cross-silo	>7	Coll.	Cen.	Shared	Consortium

^aWe estimate the application's technology readiness levels (see NASA, 2023) on the basis of available documentation.

^bColl, collaboration; Comp, competition.

^cDec, decentralised; Cen, centralised.

infrastructure and shared IP ownership, the *consortium* appears as an evident organisational archetype. Popularised in the information technology industry at the turn of the century, consortia are a form of alliance between multiple organisations that collaborate on a shared objective while maintaining their independence (see Hawkins, 1999). Literature has shown that such collaborative structure mitigates the risks of monopolisation, enhances collective innovation, and promotes fair competition by aligning incentives toward mutual benefit rather than market dominance (Doz et al., 2000). In the context of FL, the archetype helps to balance power dynamics by ensuring that decision-making, resource allocation, and IP rights are distributed more equitably among stakeholders (Bujotzek et al., 2024; Mateus et al., 2024).

The second archetype we see emerging is labelled as “in-house.” Here, the application is Gboard, Google’s virtual keyboard, which utilises FL to improve predictive text and typing suggestions while maintaining user privacy. Instead of sending raw typing data to centralised servers, Gboard trains directly on users’ devices, aggregating only the necessary updates to improve the overall model (Hard et al., 2018). In terms of governance, the model principles and access (refer to Table 1) largely define and point towards a centralised, proprietary setting.

Beyond the two archetypes observed “in the wild,” we foresee other archetypes to emerge. One of these is the joint venture whereby independent organisations form a new organisation under which model training occurs or is coordinated, by which ownership is equitably shared and typically commercialised. This archetype is opposed to consortia, where it is typically not commercialised.⁶ The conditions for such a form of alliance are relationship and firm-specific (see Pateli & Lioukas, 2011). Another archetype we foresee to be of interest is to delegate orchestration to a third party. This could be, for example, a financial market authority or central bank (Lee et al., 2023). A more speculative, experimental archetype is to delegate orchestration to a decentralised autonomous organisation (Majeed et al., 2023).

4.2. Limitations

The proposed framework is subject to limitations. For instance, when derived inductively, here using a workshop and reflections, there are limits to the generalisability of the conceptual contribution due to the context-specific nature of the data collection (Storvang et al., 2017; Jaakkola, 2020; Makowski, 2021). We addressed this limitation by application to some real-world cases, an established method for ex-ante validation, and by delineating boundary conditions in the form of assumptions (Busse et al., 2016).

Additionally, the subjective interpretations of the researchers and the potential for bias in data collection and analysis further constrain the applicability of the findings (Klein & Myers, 1999). We addressed this limitation through additional conversations and iterations with experts. Still, while our

⁶ For a comparison between a consortium and an equity joint venture, see Costa et al. (2017).

frameworks can provide conceptually rich understandings, their validity and reliability across different settings and populations remains limited, necessitating further validation and adaptation to enhance generalisability. As FL models mature, we see this as a future research direction.

5. Discussion

While writing this commentary, and in the discussions with respective experts working on FL, we encountered three distinct points that are worthwhile to discuss in the context of governance. All three are of organisational relevance; for all three, supportive literature is scarce. Therefore, we view each as an avenue for future research.

First, because the value of contributions in FL decreases as the number of participants increases (what we describe as the *learning threshold*), collaboration is at odds with utility and individual return. One key challenge to FL is that the model training can reach a ceiling of accuracy (Ω) after a certain number of participants (n) contributed. In other words, as more participants join and contribute, the model becomes more accurate, and for each new user joining at the time ($n+1$), the accuracy peaks and the utility of the contribution decrease (see Figure 3). This may indeed lead to unintended outcomes. It may be that when Ω is reached, the sustainability of the collaboration is in jeopardy; those who contributed first may argue that the value of those who contributed later is lesser (Lee et al., 2023). The orchestrator must define clear participation and ownership (reward) structures.

The second point is that who initiates the collaboration is typically in control. Checks and balances within the system are not only a question of contracts and agreements; it is also about who initiates the use case in the first place. Similar to the power plays in the formation of consortia, the initiator of an FL collaboration largely defines what kind of governance structure to set up and what the relationship is between participants. It may be that, in a commercial example, the initiator opts to create a vendor lock-in and abuse their market position (“winner-takes-all”). To do so, it would reduce the amount of power end users have by, among others, reducing or eliminating the orchestrator as a whole. Valuable lessons may be learned from literature on consortia formation and the power dynamics at play (Doz et al., 2000; Sakakibara, 2002).

At last, in our conversations and the workshop, experts flagged private companies’ inherent aversion towards the adoption of FL due to their awareness of associated risk (see also, Müller et al., 2024). On the one hand, the adoption of FL demands certain organisational capabilities—it requires necessary (technical) skills, governance structures, and communication channels between participants (Bujotzek et al., 2024). On the other hand, companies adopting FL are—at worst—exposing themselves to market risks and cybersecurity risks in the form of data leakages, predatory behaviour, and so forth. At best, they




Figure 3. Stylised visualisation of model accuracy versus contributors. The intersection between Ω and n represents the learning threshold.

need to trust an orchestrator. Either way, experts pointed out that FL may be in the limbo of a chicken-and-egg problem. This is sensible given that distributed technologies commonly face this problem (Drasch et al., 2020).

6. Outlook

As with any emerging technology, FL faces the dynamics of typical technology adoption life cycles—be it Gartner’s hype cycle or others. While deterministic by nature (Pollock & Williams, 2010), these frameworks can provide a useful lens to estimate the stage of development for FL.

The academic corpus surrounding FL appears to be entering a stage of consolidation, marked by an increasing number of early adopters seeking to understand, explore, and advance the technology. These efforts are projected into the first real-world implementations in the healthcare sector (Pati et al., 2022; Bujotzek et al., 2024; Mateus et al., 2024). These examples illustrate how real-world applications are progressing beyond purely technical aspects, in finding answers to organisational questions. Their success relies on interdisciplinary and cross-domain collaboration (van Drumpt et al., 2024).

By contrast, applications outside of the healthcare sector remain underdeveloped, often lagging in the depth and breadth of available research. This disparity suggests that while research in other sectors may be maturing, real-world adoption is still in its infancy, requiring more innovators and organisations to bridge the gap. Consequently, as also seen in the healthcare sector, we recommend to decision-makers and policy-makers that interdisciplinary and cross-sectional collaboration is key for FL’s adoption. The uptake of use cases requires more than technical expertise.

Research opportunities will continue to emerge as both industry and academia explore the implementation of FL across sectors. As collaboration markets emerge and evolve, it is essential to examine the technical advancements in FL and the governance concepts that will accompany them. For instance, we envision that organisational frameworks must be developed to address long-term and short-term adoption strategies. Short-term strategies are particularly relevant in the context of FL, as the emergence of short-term federations may necessitate agile organisational structures that enhance governance’s dynamism and account for such contingencies (see Sambamurthy & Zmud, 1999). In line with these developments, a clear research opportunity building from our theoretical framework could lead to the creation of a decision tree to guide organisations in the selection process of an organisational archetype from a practical perspective. We foresee managers and C-level executives within organisations to be the main counterparties. Moreover, we also foresee the necessity of understanding the trade-offs; in other words, the benefits and drawbacks of certain decisions when in the context of an FL organisation.

Additional research opportunities in governance arise as the global “building boom” of high-performance computers and data centres is kept alive (Weise, 2025). Specifically, beyond those owned and operated by Big Tech, we see these facilities as providing on-demand computational resources and being non-competitive, trustworthy orchestrators for the training of FL. Still, they also require governance frameworks to establish the rules based on the type of archetype they will interact with. However, given the preliminary nature of our approach and the lack of post-data analysis, a significant question remains uncertain: *will there be a multitude of different archetypes, or will they converge?*

Data availability statement. The authors confirm that all data generated or analysed during this study are included in this published article.

Author contribution. Conceptualisation: T.B., J.D.F., and S.P.M. Formal analysis: T.B., J.D.F., and S.P.M. Investigation: T.B., J.D.F., and S.P.M. Methodology: T.B., J.D.F., and S.P.M. Writing—original draft: T.B., J.D.F., and S.P.M. Writing – review and editing: T.B., J.D.F., and S.P.M. All authors approved the final submitted draft.

Funding statement. T.B. is supported by TNO under the Early Research Program “Next Generation Cryptography.” J.D.F. and S.P.M. are supported by the Luxembourg National Research Fund (FNR) and PayPal, PEARL grant reference 13,342,933/Gilbert Fridgen, and by FNR grant reference HPC BRIDGES/2022_Phase2/17886330/DELPHI. For the purpose of open access and in

fulfilling the obligations arising from the grant agreement, the authors have applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

Competing interests. The authors declare none.

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Amard A, Fernandez JD, Barbereau T and Fridgen G (2023) Federated learning in migration forecasting. In *Proceedings of the 44th International Conference on System Sciences*. https://aisel.aisnet.org/treos_icis2023/23/
- Berente N, Gu B, Recker J and Santhanam R (2021) Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450. <https://aisel.aisnet.org/misq/vol45/iss3/16/>
- Brown AE and Grant GG (2005) Framing the frameworks: A review of IT governance research. *Communications of the AIS* 15(1), 38. <https://doi.org/10.17705/1CAIS.01538>.
- Bujotzek MR, Akünal Ü, Denner S, Neher P, Zenk M, Frodl E, Jaiswal A, Kim M, Krekieleh NR, Nickel M, Ruppel R, Both M, Döllinger F, Opitz M, Persigehl T, Kleesiek J, Penzkofer T, Maier-Hein K, Bucher A and Braren R (2024) Real-world federated learning in radiology: Hurdles to overcome and benefits to gain. *Journal of the American Medical Informatics Association* 32(1), 193–205. <https://doi.org/10.1093/jamia/ocae259>.
- Busse C, Kach AP and Wagner SM (2016) Boundary conditions: What they are, how to explore them, why we need them, and when to consider them. *Organizational Research Methods* 20(4), 574–609. <https://doi.org/10.1177/1094428116641191>.
- Cassano L, D'Abramo J, Munir S and Ferretti S (2024) Trust and Resilience in Federated Learning Through Smart Contracts Enabled Decentralized Systems. Available at <https://arxiv.org/abs/2407.06862>
- Costa e, Silva S and Oliveira SM (2017) Partner selection in international joint ventures: A framework for the analysis of factors relevant to the selection of partners. *The Marketing Review* 17(2), 199–215. <https://doi.org/10.1362/146934717X14909733966182>.
- Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, Liu A, Costa AB, Wood BJ, Tsai C-S, Wang C-H, Hsu C-N, Lee CK, Ruan P, Xu D, Wu D, Huang E, Kitamura FC, Lacey G, et al. (2021) Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine* 27, 1735–1743. <https://doi.org/10.1038/s41591-021-01506-3>.
- Department for Science, Innovation and Technology (2024) *Repository of Privacy Enhancing Technologies (PETs) Use Cases*. United Kingdom Government. Available at <https://www.gov.uk/guidance/repository-of-privacy-enhancing-technologies-pets-use-cases>
- Doz YL, Olk PM and Ring PS (2000) Formation processes of R&D consortia: Which path to take? Where does it lead? *Strategic Management Journal* 21(3), 239–266. [https://doi.org/10.1002/\(SICI\)1097-0266\(200003\)21:3<239::AIDSMJ97>3.0.CO;2-K](https://doi.org/10.1002/(SICI)1097-0266(200003)21:3<239::AIDSMJ97>3.0.CO;2-K).
- Drasch BJ, Fridgen G, Manner-Romberg T, Nolting FM and Radszuwill S (2020) The token's secret: The two-faced financial incentive of the token economy. *Electronic Markets* 30(3), 557–567. <https://doi.org/10.1007/s12525-020-00412-9>.
- Diessen SW, Monsieur G and Van Den W-J (2022) Data market design: A systematic literature review. *IEEE Access* 10, 33123–33153. <https://doi.org/10.1109/ACCESS.2022.3161478>.
- Dwivedi YK, Hughes L, Ismagilova E, Aarts G, Coombs C, Crick T, Duan Y, Dwivedi R, Edwards J, Eirug A, Galanos V, Ilavarasan PV, Janssen M, Jones P, Kar AK, Kizgin H, Kronemann B, Lal B, Lucini B, et al. (2021) Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management* 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>.
- European Commission (2023) *Commission Regulation (EU) 2023/1066 of 1 June 2023 on the Application of Article 101(3) of the Treaty on the Functioning of the European Union to Certain Categories of Research and Development Agreements*. Official Journal of the European Union. Available at <http://data.europa.eu/eli/reg/2023/1066/oj>
- Ezzeldin YH, Yan S, He C, Ferrara E and Avestimehr S (2023) Fairfed: Enabling group fairness in federated learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 37(6), 7494–7502.
- Fernández JD (2023) *Breaking Data Silos with Federated Learning* [Doctoral dissertation, University of Luxembourg]. <https://hdl.handle.net/10993/57042>
- Fernández JD, Barbereau T, Baim R and Rieger A (2024) Opportunities and applications of federated learning in the financial services industry. In Fridgen G, Guggenberger T, Sedlmair J, Urbach N (eds), *Decentralization Technologies. Financial Innovation and Technology*. Springer, Cham. https://doi.org/10.1007/978-3-031-66047-4_11
- Fernández JD, Potenciano Menci S, Lee CM, Rieger A and Fridgen G (2022) Privacy-preserving federated learning for residential short-term load forecasting. *Applied Energy* 326, 119915. <https://doi.org/10.1016/j.apenergy.2022.119915>.
- Fernández JD, Potenciano Menci S and Pavic I (2023) Towards a peer-to-peer residential short-term load forecasting with federated learning. In *2023 IEEE Belgrade PowerTech*. pp. 1–6. <https://doi.org/10.1109/PowerTech55446.2023.10202782>.
- Fu J, Hong Y, Ling X, Wang L, Ran X, Sun Z, Wang WH, Chen Z and Cao Y (2024) Differentially private federated learning: A systematic review. *arXiv preprint arXiv:2405.08299*.
- Gerloff C (2022) *Christiangerloff/set-you-free: First step* (Version v0.1.0). Zenodo. <https://doi.org/10.5281/zenodo.6907681>
- Grover V, Lederer AL and Sabherwal R (1988) Recognizing the politics of mis. *Information & Management* 14(3), 145–156. [https://doi.org/10.1016/0378-7206\(88\)90005-5](https://doi.org/10.1016/0378-7206(88)90005-5).

- Hård M** (1993) Beyond harmony and consensus: A social conflict approach to technology. *Science, Technology, & Human Values* 18 (3), 408–432. <https://doi.org/10.1177/016224399301800402>.
- Hard A, Rao K, Mathews R, Ramaswamy S, Beaufays F, Augenstein S, Eichner H, Kiddon C and Ramage D** (2018) Federated Learning for Mobile Keyboard Prediction. *arXiv*: <https://doi.org/10.48550/arXiv.1811.03604>
- Hawkins R** (1999) The rise of consortia in the information and communication technology industries: Emerging implications for policy. *Telecommunications Policy* 23(2), 159–173. [https://doi.org/10.1016/S0308-5961\(98\)00085-8](https://doi.org/10.1016/S0308-5961(98)00085-8).
- Jaakkola E** (2020) Designing conceptual articles: Four approaches. *AMS review* 10(1), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>.
- Jin W, Yao Y, Han S, Joe-Wong C, Ravi S, Avestimehr S and He C** (2023) Fedml-hc: An efficient homomorphic encryption-based privacy-preserving federated learning system. *arXiv preprint arXiv:2303.10837*.
- Jordan MI and Mitchell TM** (2015) Machine learning: Trends, perspectives, and prospects. *Science* 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>.
- Kaminaga H, Awaysheh FM, Alawadi S and Kamm L** (2023) Mpcfl: Towards multi-party computation for secure federated learning aggregation. In *Proceedings of the 16th IEEE/ACM International Conference on Utility and Cloud Computing (UCC '23)*. Association for Computing Machinery, New York, NY, USA, 19, 1–10. <https://doi.org/10.1145/3603166.3632144>.
- Kearns GS and Lederer AL** (2004) The impact of industry contextual factors on IT focus and the use of IT for competitive advantage. *Information & Management* 41(7), 899–919. <https://doi.org/10.1016/j.im.2003.08.018>.
- Keen PGW** (1981) Information systems and organizational change. *Communications of the ACM* 24(1), 24–33. <https://doi.org/10.1145/358527.358543>.
- Khatri V and Brown CV** (2010) Designing data governance. *Communications of the ACM* 53(1), 148–152. <https://doi.org/10.1145/1629175.1629210>.
- Klein H and Myers M** (1999) A set of principles for conducting and evaluating interpretive field studies in information systems. *MIS Quarterly* 23(1), 67–93. <https://doi.org/10.2307/249410>.
- Konečný J, McMahan B, Ramage D and Richtárik P** (2016) Federated optimization: Distributed machine learning for on-device intelligence. <https://arxiv.org/abs/1610.05492>.
- Lee CM, Fernández JD, Potenciano Menci S, Rieger A and Fridgen G** (2023) Federated learning for credit risk assessment. In *Proceedings of the 56th Hawaii International Conference on System Sciences*. pp. 386–395. <https://hdl.handle.net/10125/102676>
- Majeed U, Hassan SS, Han Z and Hong CS** (2023) DAO-FL: Enabling Decentralized Input and Output Verification in Federated Learning with Decentralized Autonomous Organizations. *Authorea Preprints*. <https://doi.org/10.36227/techrxiv.24546502.v1>
- Makowski PT** (2021) Optimizing concepts: Conceptual engineering in the field of management—The case of routines research. *Academy of Management Review* 46 (4), 702–724. <https://doi.org/10.5465/amr.2019.0252?journalCode=amr>.
- Martínez Beltrán ET, Pérez MQ, Sánchez PMS, Bernal SL, Bovet G, Pérez MG, Pérez GM and Celdrán AH** (2023) Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges. *IEEE Communications Surveys & Tutorials* 25(4), 2983–3013. <https://doi.org/10.1109/COMST.2023.3315746>.
- Mateus P, Moonen J, Beran M, Jaarsma E, van der Landen SM, Heuvelink J, Birhanu M, Harms AGJ, Bron E, Wolters FJ, Cats D, Mei H, Oomens J, Jansen W, Schram MT, Dekker A and Bermejo I** (2024) Data harmonization and federated learning for multi-cohort dementia research using the OMOP common data model: A Netherlands consortium of dementia cohorts case study. *Journal of Biomedical Informatics* 155, 104661. <https://doi.org/10.1016/j.jbi.2024.104661>.
- McMahan B, Moore E, Ramage D, Hampson S and y Arcas BA** (2017) Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 54, pp. 1273–1282. Fort Lauderdale, Florida, USA. JMLR: W&CP. <https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf>
- Mergel I, Edelmann N and Haug N** (2019) Defining digital transformation: Results from expert interviews. *Government Information Quarterly* 36(4), 101385. <https://doi.org/10.1016/j.giq.2019.06.002>.
- Mintzberg H** (1983) *Power in and around Organizations*. Prentice Hall: Hoboken, New Jersey, U.S. <https://doi.org/10.1177/017084068400500419>
- Müller T, Zahn M and Matthes F** (2024) Revealing the impacting factors for the adoption of federated machine learning in organizations. In *Proceedings of the 57th Hawaii International Conference on System Sciences*. <https://doi.org/10.24251/HICSS.2023.881>
- NASA**. (2023) *Technology Readiness Levels—NASA*. Available at <https://www.nasa.gov/directorates/somd/space-communications-navigation-program/technology-readiness-levels>
- OECD** (2016) Recommendation of the council for development co-operation actors on managing the risk of corruption. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0431>
- Patel A and Lioukas S** (2011) The choice of governance mode in ICT alliances: A property rights approach. *Information & Management* 48(1), 69–77. <https://doi.org/10.1016/j.im.2011.01.002>.
- Pati S, Baid U, Edwards B, Sheller M, Wang S-H, Reina GA, Foley P, Gruzdev A, Karkada D, Davatzikos C, Sako C, Ghodasara S, Bilello M, Mohan S, Vollmuth P, Brugnara G, Preetha CJ, Sahm F, Maier-Hein K, et al.** (2022) Federated learning enables big data for rare cancer boundary detection. *Nature Communications* 13(7346), 1–17. <https://doi.org/10.1038/s41467-022-33407-5>.
- Pollock N and Williams R** (2010) The business of expectations: How promissory organizations shape technology and innovation. *Social Studies of Science* 40(4), 525–548. <https://doi.org/10.1177/0306312710362275>.

- Price WN and Cohen IG (2019) Privacy in the age of medical big data. *Nature Medicine* 25, 37–43. <https://doi.org/10.1038/s41591-018-0272-7>.
- Rashid MM, Xiang Y, Uddin MP, Tang J, Sood K and Gao L (2025) Trustworthy and fair federated learning via reputation-based consensus and adaptive incentives. *IEEE Transactions on Information Forensics and Security* 20, 2868–2882.
- Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, Bakas S, Galtier MN, Landman BA, Maier-Hein K, Ourselin S, Sheller M, Summers RM, Trask A, Xu D, Baust M and Cardoso MJ (2020) The future of digital health with federated learning. *npj Digital Medicine* 3(119), 1–7. <https://doi.org/10.1038/s41746-020-00323-1>.
- Sakakibara M (2002) Formation of R&D consortia: Industry and company effects. *Strategic Management Journal* 23(11), 1033–1050. <https://doi.org/10.1002/smj.272>.
- Sambamurthy V and Zmud RW (1999) Arrangements for information technology governance: A theory of multiple contingencies. *MIS Quarterly* 23(2), 261–290. <https://doi.org/10.2307/249464>.
- Sheller M, Edwards B, Anthony Reina G, Martin J and Bakas S (2019) Federated learning in neuro-oncology for multi-institutional collaborations without sharing patient data. *Neuro-Oncology*, 21(Supplement 6), vi176–vi177. <https://doi.org/10.1093/neuonc/noz175.737>
- Spiekermann M (2019) Data marketplaces: Trends and monetisation of data goods. *Intereconomics* 54(4), 208–216. <https://doi.org/10.1007/s10272-019-0826-z>.
- Sprenkamp K, Fernández JD, Eckhardt S and Zavolokina L (2024) Overcoming intergovernmental data sharing challenges with federated learning. *Data & Policy* 6, e27. <https://doi.org/10.1017/dap.2024.19>.
- Storvang P, Mortensen B and Clarke AH (2017) Using workshops in business research: A framework to diagnose, plan, facilitate and analyze workshops. In *Collaborative Research Design*. Springer: Singapore, pp. 155–174. https://doi.org/10.1007/978-981-10-5008-4_7
- Truong N, Sun K, Wang S, Guitton F and Guo Y (2021) Privacy preservation in federated learning: An insightful survey from the GDPR perspective. *Computers & Security* 110, 102402. <https://doi.org/10.1016/j.cose.2021.102402>.
- Tzimas T (2021) *Legal and Ethical Challenges of Artificial Intelligence from an International Law Perspective*. Springer International Publishing: Switzerland. <https://doi.org/10.1007/978-3-030-78585-7>
- van Drumpt S, Timan T, Talie S, Veugen T and van de Burgwal L (2024) Digital transitions in healthcare: The need for transdisciplinary research to overcome barriers of privacy enhancing technologies uptake. *Health and Technology* 14(4), 709–723. <https://doi.org/10.1007/s12553-024-00850-x>.
- Vucovich M, Tarcar A, Rebelo P, Gade N, Porwal R, Rahman A, Redino C, Choi K, Nandakumar D, Schiller R, et al. (2022) Anomaly detection via federated learning. *arXiv preprint arXiv:2210.06614*.
- Weber K, Otto B and Österle H (2009) One size does not fit all—a contingency approach to data governance. *Journal of Data and Information Quality* 1(1), 1–27. <https://doi.org/10.1145/1515693.1515696>.
- Weigl L, Barbereau T and Fridgen G (2023) The construction of self-sovereign identity: Extending the interpretive flexibility of technology towards institutions. *Government Information Quarterly* 40(4), 101873. <https://doi.org/10.1016/j.giq.2023.101873>.
- Weill P and Ross JW (2004) *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results*. Harvard Business Press: New York City, New York, U.S.
- Weise K (2025) Big Tech keeps its A.I. Data center spending boom alive. *New York Times*. Available at <https://www.nytimes.com/2025/02/08/technology/deepseek-data-centers-ai.html>
- Winner L (1993) Upon opening the black box and finding it empty: Social constructivism and the philosophy of technology. *Science, Technology, & Human Values* 18(3), 362–378. <https://doi.org/10.1177/016224399301800306>.
- Winter S, Berente N, Howison J and Butler B (2014) Beyond the organizational ‘container’: Conceptualizing 21st century sociotechnicalwork. *Information and Organization* 24(4), 250–269. <https://doi.org/10.1016/j.infoandorg.2014.10.003>.
- WIPO (2005) *Exchanging Value—Negotiating Technology Licensing Agreements: A Training Manual*. World Intellectual Property Organisation. Available at https://www.wipo.int/edocs/pubdocs/en/licensing/906/wipo_pub_906.pdf
- Witt L, Heyer M, Toyoda K, Samek W and Li D (2023) Decentral and incentivized federated learning frameworks: A systematic literature review. *IEEE Internet of Things Journal* 10(4), 3642–3663. <https://doi.org/10.1109/JIOT.2022.3231363>.
- Yang T-J, Xiao Y, Motta G, Beaufays F, Mathews R and Chen M (2022) Online model compression for federated learning with large models. *arXiv preprint arXiv:2205.03494*.
- Yang W, Zhang Y, Ye K, Li L and Xu C-Z (2019) Ffd: A federated learning based method for credit card fraud detection. In *BigData: 8th International Congress*. pp. 18–32. https://doi.org/10.1007/978-3-030-23551-2_2
- Yin X, Zhu Y and Hu J (2021) A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. *ACM Computing Surveys* 54(6), 1–36. <https://doi.org/10.1145/3460427>.
- Zhang S (2023) Good governance essential for enterprises deploying AI. *MIT Technology Review*. Available at <https://www.technologyreview.com/2023/07/18/1075972/good-governance-essential-for-enterprises-deploying-ai>