



OPEN ACCESS

EDITED BY

Vinod Kumar Chauhan,
University of Strathclyde, United
Kingdom

REVIEWED BY

Elahe Khatibi,
University of California, Irvine, United
States
Marios Tyrovolas,
Technological Educational Institute of
Epirus, Greece

*CORRESPONDENCE

Fabio Stella
✉ fabio.stella@unimib.it

RECEIVED 02 April 2026

REVISED 15 June 2026

ACCEPTED 15 June 2026

PUBLISHED 17 July 2026

CITATION

Rocchi N, Scutari M, Zanga A,
Nagarajan R and Stella F (2026)
Federated causal discovery in medicine:
trends, opportunities, and challenges.
Front. Digit. Health 8:1846020.
doi: 10.3389/fdgth.2026.1846020

COPYRIGHT

© 2026 Rocchi, Scutari, Zanga,
Nagarajan and Stella. This is an open-
access article distributed under the
terms of the [Creative Commons
Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use,
distribution or reproduction in other
forums is permitted, provided the
original author(s) and the copyright
owner(s) are credited and that the
original publication in this journal is
cited, in accordance with accepted
academic practice. No use, distribution
or reproduction is permitted which does
not comply with these terms.

Federated causal discovery in medicine: trends, opportunities, and challenges

Niccolò Rocchi^{1,2}, Marco Scutari³, Alessio Zanga¹,
Radha Nagarajan⁴ and Fabio Stella^{1*}

¹Department of Informatics, Systems and Communication (DISCo), University of Milano-Bicocca, Milan, Italy, ²Department of Epidemiology and Data Science, Fondazione IRCCS Istituto Nazionale dei Tumori, Milan, Italy, ³Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland, ⁴Rady Children's Health, Orange, CA, United States

Exponential growth and continued digitisation have accelerated the adoption of data-driven and evidence-based approaches in medicine. This includes deciphering associations, including potential causal associations, from multivariate observational biomedical data under certain implicit assumptions. Such an evidence approach marks the shift from classical hypothesis testing to discovery and hypothesis generation. Widespread adoption of common data models has especially accelerated collaborative approaches in medicine while transitioning from centralised to federated architectures that facilitate discovery without the explicit sharing of sensitive medical data. This perspective provides an overview of causal discovery from observational data with a focus on federated learning. Specifically, it outlines the trends, opportunities, and challenges of federated causal discovery in medicine. While medical research has traditionally relied on the hierarchy of evidence generated from the evidence pyramid, the ability of federated causal discovery to facilitate evidence generation collaboratively from heterogeneous sources is expected to enhance the generalizability and transportability of findings while addressing sample size considerations—a critical aspect for its successful and widespread adoption in medicine.

KEYWORDS

causal graph, federated architecture, federated learning, healthcare, medicine

1 Introduction

During the last decade, an exponential growth in *multivariate* and *multimodal* biomedical data has taken place due to several factors, including the widespread adoption of digital technology, continued digitisation efforts across a multitude of source systems, and high-throughput assays (1). While multivariate data typically consists of *relational* and *structured* data that can be represented in a tabular format (e.g., numerical and categorical features), multimodal data includes *non-relational* and *unstructured* data (e.g., texts and images). Multivariate data are typically stored in *data warehouses* for querying and downstream analytics. *Centralized* and *federated* warehouse architectures have been proposed for the same (2). A brief description of these architectures, along with their pros and cons, is shown in [Table 1](#).

There is an increasing interest in harnessing information from multivariate and multimodal observational data to improve patient outcomes in an evidence-based manner. This includes modelling associations, including causal associations, from multivariate observational biomedical data, complementing traditional evidence

TABLE 1 Pros and cons of centralised and federated architectures.

Architecture	Pros	Cons
Centralised	Centralised governance with representation from the participating institutions.	A centralised resource is a single point of failure.
	Single storage and compute infrastructure for querying and analytics.	Centralising high-throughput data may have prohibitive bandwidth and storage costs.
	Compute architecture is straightforward compared to federated ones.	Varying compliance, lack of trust, and regulatory regimes across institutions may limit data access and the analyses permitted.
Federated	Architectures are more robust to failures because of distribution.	Architectures are more complex and difficult to manage.
	No explicit data sharing, enhanced trust across the participating institutions, and reduced patient consent burden.	Data drift and uneven data distributions across institutions can impact training and the generalizability of models.
	More control across the participating institutions. Data are compliant with local governance.	Execution of queries and analytics can have relatively higher latencies in federated architectures.

generation (3). Such a data-driven approach, combined with domain knowledge, has the potential to validate established associations and discover novel ones for critical assessment, thereby generating new hypotheses and research questions. It also has the potential to provide novel system-level insights, a precursor to the development of targeted interventions, tailored treatment regimens, and disease management strategies (4). More importantly, it can complement the hierarchy of evidence from the *evidence pyramid*, including systematic reviews and *randomised controlled trials* (RCTs), which are widely used in medicine (3).

Networks have proven to be especially useful abstractions in this regard (5), with *nodes* representing the variables of interest, and *edges* their associations. These networks can be modelled across multiple scales (e.g., molecular, clinical, or demographic data) of varying granularity and resolution (6), as well as from cross-sectional and longitudinal profiles. While longitudinal profiles capture the explicit temporal evolution of multivariate processes, they are usually challenging to generate from economic and stationarity standpoints. The latter requires controlling several factors to preserve statistical properties over time (7). Not surprisingly, cross-sectional profiles that interrogate the multivariate process within a chosen time window, in conjunction with replicate measurements, have been prevalent. While earlier attempts focused on modelling these networks by estimating pairwise dependencies between variables [e.g., relevance networks (8)], there is increasing evidence that dependence between a pair of variables need not be direct, warranting the inclusion of conditional dependencies. Techniques such as *causal Bayesian networks* (CBNs) (9) have proven particularly helpful in this regard, with broad application in medicine and healthcare. These include deciphering associations from sequencing (10), molecular (11), epidemiological (12) and electronic health records (13) with a focus on oncology (14), neurology (15), rare (16) and infectious diseases (17), among others. CBNs have also been used successfully to model relationships from diverse data sets, including multi-centre (18), temporal (19), state-space (20), partially observed (21), and multimodal data (22) under the broad theme of *causal discovery* (CD) (23). We refer to [Supplementary Section S1](#), for further details.

This perspective examines applications of *federated causal discovery* (FCD) in medicine and healthcare. Specifically, it introduces and describes the “What”, “Why”, and “How” of causal networks, while offering a comprehensive and critical overview of FCD, including its opportunities and challenges. The manuscript does not cover core components of causal inference in the usual medical and methodological sense, such as explicit estimands, identification strategies, adjustment logic, estimation procedures, or sensitivity analysis (24, 25). A detailed description of causal representation, causal discovery, and federated learning is provided in the [Supplementary Sections S1 and S2](#) for completeness. Definitions and popular algorithms for FCD are discussed in Section 2. Section 3 describes the study that motivated us to write this paper, which is used to illustrate a practical application of FCD. Section 4 presents current trends, opportunities (4.1), and challenges (4.2) of FCD with a focus on medicine.

2 Federated causal discovery

Federated causal discovery (FCD) extends classical CD to multiple data sources without explicit data sharing. In this sense, it aims to find the structure of one (or more) CBNs that explain the data-generating mechanism behind the measured data. A single global CBN is most appropriate when the underlying data-generating process is assumed to be identical across all participating sites, such as in pooled homogeneous populations. In contrast, allowing the CBN to vary can capture site- or patient-specific causal mechanisms is more appropriate when there are differences in demographics, clinical protocols, or disease subtypes. An FL taxonomy is shown in [Table 2](#), and described in [Supplementary Section S2](#). FCD is a subset of FL. In particular:

- FCD is typically implemented on horizontally partitioned data, as it is challenging to compute sufficient statistics involving variables that are never observed jointly across clients. Typically, it cannot identify the true *causal graph* (CG) because multiple *directed acyclic graphs* (DAGs) may equally explain the data at hand (26).

TABLE 2 FL taxonomy based on different traits.

Trait	Type	Description
Partitioning	<u>Horizontal</u>	Full overlap in variables, no overlap in data points.
	Vertical	No overlap in variables, full overlap in data points.
	Hybrid	Partial overlap in variables, partial overlap in data points.
Learning	<u>Offline</u>	Model is learnt once from the available data.
	Online	Model is updated whenever new data are available.
Task	<u>Single-task</u>	Learning one global model.
	Multi-task	Learning different models for different clients.
Aggregation	<u>Synchronous</u>	Each client contributes to each aggregation.
	Asynchronous	Some clients may not contribute to each aggregation.
Sharing	<u>Model</u>	The entire local models are shared with the server.
	<u>Knowledge</u>	Model's intermediate information is shared.
	Synthetic data	Local synthetic data are shared instead of raw data.
Scale	<u>Cross-silo</u>	Few clients, high computational power, and availability.
	Cross-device	Many clients, limited computational power, and availability.
Topology	<u>Star schema</u>	One server is in charge of aggregation.
	Distributed	Multiple aggregators are allowed.

Yellow represents the data partitioning schema, blue represents the algorithmic structure and task, and red represents the client-server architecture. Underlined terms identify the framework of FCD within FL; exceptions are marked in [Supplementary Figure S2](#).

- FCD algorithms follow an offline learning procedure, as online CD is typically challenging even in single-source settings. This is not a significant limitation, as clinical data are more often collected retrospectively than prospectively.
- There are instances of multi-task FCD in the literature, such as FMT-NBN (27). They typically assume a shared CG (Section 2.1) to be realistic. How to account for distribution shifts among clients remains an open research question (Section 4.2).
- Synchronous aggregation is commonly carried out, although asynchronous learning may be practicable in principle.
- Most privacy-preserving techniques in FCD involve sharing scores or statistics, which are sometimes encrypted to protect the data. The extent to which original data can actually be traced back from scores or statistics remains an unexplored area (Section 4.2).
- CD requires significant computational resources and time for big data, which favours the cross-silo architecture. There are no examples of (fully) distributed FCD, but it could be a solution to develop personalised models for multi-task FL.

FCD often makes additional assumptions, which are discussed in the following section.

2.1 Assumptions

In the most general setting, a source-dependent *structural causal model* (SCM) is assumed to be responsible for generating each local data set $\mathcal{D}^1, \dots, \mathcal{D}^k$. Specifically, let $\mathcal{M}^1, \dots, \mathcal{M}^K$ be a collection of SCMs underlying the K data sources where $\mathcal{M}^k$ is given by:

$$\mathcal{M}^k = (\mathbf{U}^k, \mathbf{X}^k, \mathcal{F}^k, P^k(\mathbf{U}^k)), \\ X_i^k := f_i^k(\Pi_i^k, U_i^k).$$

Without loss of generality,¹ we set $\mathbf{U}^k = \mathbf{U}, \forall k$. Because FCD research and applications focus overwhelmingly on the horizontally-partitioned data, we set $\mathbf{X}^k = \mathbf{X}, \forall k$.

Figure 1 lists a set of common assumptions regarding both the model and the data. They are relevant beyond FL, but we discuss them in the context of FCD.

Assumption 1 Shared causal graph.

Cause-and-effect relationships are invariant across clients, that is $\Pi_i^k = \Pi_i, \forall i, k$. Hence, all CGs coincide: $\mathcal{G}^k = \mathcal{G}, \forall k$.

In real medical systems, measurement protocols, treatment policies, patient populations, and institutional workflows may differ substantially across sites. Therefore, [Assumption 1](#) may not always be realistic, as in the case of fMRI data (28). Even when the cause-and-effect relationships are population-invariant, the networks' parameters may vary (29). For instance, disease aetiology and drug mechanisms of action are similar regardless of where a physician collects measurements, but drug effectiveness may vary across subpopulations. For a more in-depth discussion, we refer the reader to Section 4.

Under [Assumption 1](#), FCD consists of obtaining a single $\mathcal{G}$ from $\mathcal{D}^1, \dots, \mathcal{D}^k$. Otherwise, we talk about *multi-task* FCD as the task of collaboratively finding $\mathcal{G}^k, \forall k$.

Assumption 2 Observational data.

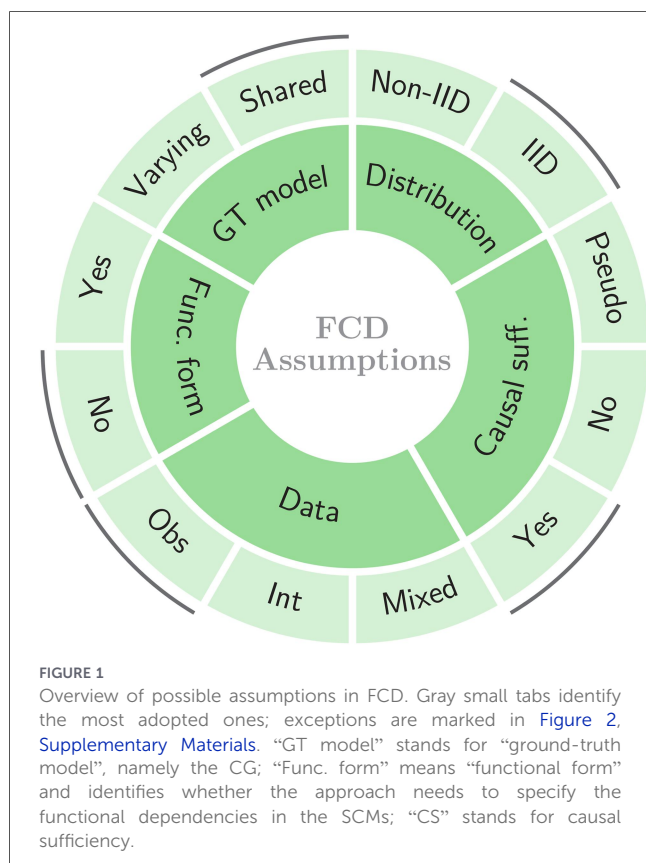
Data were collected during routine work without deliberate manipulations of the variables. In other words, the variables are not *intervened* on (9).

For a client k and variable X_i , a *perfect* intervention sets a structural assignment for X_i to $X_i := x_i$ for a certain x_i . On the other hand, an *imperfect* intervention modifies either $P^k(\mathbf{U})$ or $f_i^k(\cdot)$ in $\mathcal{M}^k$ without fixing them to a single, deterministic value.

Assumption 3 Pseudo-causal sufficiency.

Any variable causing two or more variables in $\mathbf{X}^k$ is included in $\mathbf{X}^k, \forall k$.

¹In fact, we may set $\mathbf{U} = \bigcup_k \mathbf{U}^k$, $P(\mathbf{U})$ its joint distribution, and $P^k(\mathbf{U}) = \sum_{i \neq k} P(\mathbf{U}^k, \mathbf{U}^i), \forall k$.



Assumption 3 coincides with the standard *causal sufficiency* assumption holding locally at each data source (30). However, multi-source causal sufficiency requires stronger guarantees, as one variable may cause two other variables measured at different sites.

Assumption 4 Causal sufficiency.

Any variable causing two or more variables in $\bigcup_k \mathbf{X}^k$ is included in $\mathbf{X}^k$, $\forall k$.

The hybrid and vertical data partitioning settings conflict with [Assumptions 3](#) and [4](#) because not all variables are necessarily measured across all sources. Note that causal sufficiency implies pseudo-causal sufficiency, but not vice versa. Unmeasured confounding, selection bias and measurement errors can all be seen as different instances of violations of causal sufficiency, and they can lead to incorrect causal edges even when CD runs in a federated setting. An illustrative example contrasting [Assumptions 3](#) and [4](#) will be presented in [Section 3](#).

Assumption 5 IID data.

Data are *independent and identically distributed* (IID) within and across the clients, hence $P^k(\mathbf{U}) = P(\mathbf{U})$, $\forall k$.

Non-IID data commonly arise from *distribution drifts*, or *shifts*, between different populations: $P^i(\mathbf{X}) \neq P^j(\mathbf{X})$, for some $i \neq j$. If [Assumptions 2](#) and [3](#) hold, we can account for them

using a set of *context variables* as illustrated in [Figure 2](#) (31). Consider the CBN $(\mathcal{G}^*, \theta^*)$, where (i) $\mathcal{G}^*$ is the enhanced version of $\mathcal{G}$ obtained by adding a set of variables $\mathbf{C}$ and an edge from $C \in \mathbf{C}$ to $X \in \mathbf{X}$ whenever $P^i(X) \neq P^j(X)$, $i \neq j$; (ii) θ^* parametrizes the joint distribution $P^*(\mathbf{X}, \mathbf{C})$. Each assignment $\mathbf{C} = \mathbf{c}$ defines specific environmental, individual, or measurement conditions (31). The local data set $\mathcal{D}^k$ is collected under the environment $\mathbf{C} = \mathbf{c}^k$, and the underlying CBN is $(\mathcal{G}, \theta^k)$. Here, θ^k parametrizes the joint distribution $P^k(\mathbf{X}) = P^*(\mathbf{X}, \mathbf{C} \mid \mathbf{C} = \mathbf{c}^k)$.

Hypothetically, $\mathcal{G}^*$ can be learned whenever the context $\mathbf{C}$ is recorded at each client. However, those variables are usually unknown, so at least one indicator variable C can serve as a proxy, with $C = k$ denoting the k -th client (2). Note that this approach requires the server to know which client each local update originates from, which is not always the case in highly private settings.

2.2 Algorithms

Causal discovery using multiple data sets is well-explored in the literature (31). Basic approaches to FCD involve performing CD locally and aggregating the learned graphs through edge voting, union, or intersection (32). These approaches prove effective when local sample sizes in individual clients are large enough and [Assumption 5](#) assumption holds. However, smaller sample sizes at the client level ([Figure 2](#)) may not be representative of the underlying population, leading to statistically significant differences in local distributions. Moreover, [Assumption 5](#) may not be satisfied in many scenarios.


[Table 3](#) presents the main FCD algorithms linked to the assumptions and settings discussed in [Section 2.1](#).² Each works under horizontal data partitioning, except for FMT-NBN, which can be applied to non-horizontal settings. The table columns and notable algorithmic instances are detailed in the caption.

Except for FedC²SL, in which clients share encoded statistics, no other method directly addresses data privacy. Almost all of them argue that data leakage is prevented by claiming that no client data are exposed. While Gao et al. (33) acknowledge that avoiding data sharing is only a partial solution, the extent to which the model's parameters or structure reveal sensitive information remains unclear in the literature. As acknowledged by Mian et al. (34), differential privacy (35) can be achieved by injecting a certain amount of noise into the broadcast from client to server. Moreover, techniques such as secure multi-party computation (36) and encryption could be exploited and revealed to be effective in other domains. However, the feasibility of these approaches in FCD, the computational burden they entail, and their implications for privacy protection are underexplored (see [Section 4.2](#)).

2.2.1 Constraint

The FedC²SL algorithm (36) extends PC and the *fast causal inference* (FCI) algorithms (23) to the federated setting. FCI

²The search and inclusion criteria for this review can be found in [Supplementary Section S3](#).



works in contexts where some variables may be unobserved, while PC assumes causal sufficiency (Assumption 3). The server iteratively constrains the DAG space using *federated conditional independence tests* that distribute the computation of sufficient statistics and securely aggregate them. The FedCDH algorithm (37) does the same, but clients release sufficient statistics to the server only once.

2.2.2 Score

The PERI algorithm (34) builds on the Greedy Equivalent Search (GES) algorithm (52). It minimises the worst-case *regret* of a function that measures how far a DAG is from the DAG that best fits the data across all clients. PERI ensures privacy and convergence guarantees. The FedCausal algorithm (39) extends the gradient NOTEARS-ADMM algorithm (41) by providing a global optimisation function.

2.2.3 Ensemble

Model ensemble techniques consist of finding a *consensus* DAG from a given collection (53). The FedGES algorithm (51) uses model ensemble techniques to combine locally learned DAGs at the server level. The procedure is iterative: the global DAG is fused with the client's DAG, becoming the starting point for the subsequent local optimisation. Note that naive aggregation of edges, such as union or averaging, does not ensure acyclicity, while edges lose their causal and statistical meanings. The FedLCS algorithm (50) is a local FL extension of the PC algorithm that finds only the edges incident on a single variable, rather than the entire CG.

3 A motivating example

Limited data availability frequently constrains cancer research, as do lengthy ethical approval processes and stringent requirements for protecting patient privacy. These challenges substantially hinder the development and validation of data-driven methodologies in real-world clinical settings, and motivate our use of FCD.

3.1 The clinical problem

Endometrial cancer (EC) is a cancer of the endometrium of the uterus. Approximately 90,000 patients die each year due to EC, calling for more research on personalised EC treatments (54). In this context, pelvic and para-aortic lymph node metastases (LNM) are among the most important prognostic factors for choosing adjuvant treatment and improving survival in node-positive EC. Approximately 10% of endometrial cancer patients present with LNMs at diagnosis according to clinical literature (54). The clinical problem is therefore to disentangle the interplay between administered treatments, LNM status, and other covariates to support clinicians in choosing the optimal therapy for each patient to maximise the survival and to reduce the risk of relapse (55).

3.2 The distributed data setting and why centralisation was not feasible

The European Network for Individualised Treatment of Endometrial Cancer (ENITEC) study and the PiPelle prospective

TABLE 3 Characterisation of the main FCD algorithms.

#	Name	Method		Assumptions					Application		
		Class	Share	Acycl.	Suff.	Int.	SCM	Data	RWD	Scal.	Code
Wang et al. (36)	FedC ² SL	Constr.	Stat.	✓				D	✓	(100, 10), (50, 64)	✓
Li et al. (37)	FedCDH	Constr.	Stat.	✓	✓(p)			C, D	✓	(30, 10), (6, 32)	✓
Mian et al. (34)	PERI	Score (D)	Score	✓	✓			C (*)	✓	(15,9), (9, 100)	✓
Mian et al. (38)	RFCD	Score (D)	Score	✓	✓			C (*)		(10, 7)	
Yang et al. (39)	FedCausal	Score (C)	Param. (subset)	✓	✓			C	✓	(80, 10)	
Liu et al. (40)	NOTEARS-PFL	Score (C)	Param.		✓		✓	C	✓	(80, 10), (50, 64)	
Ng and Zhang (41)	NOTEARS-ADMM	Score (C)	Param.		✓			C	✓	(100, 10), (50, 64)	✓
Liu et al. (42)	FedCASL	Score (C)	Param.		✓			C	✓	(40, 10), (30, 32)	
Ye et al. (43)	DARLS	Average	Param., grad.	✓	✓		✓	C, D	✓	(70, 20)	✓
Huang et al. (44)	FedPC	Average	Struct.	✓	✓			C, D	✓	(441, 15)	✓
Gao et al. (33)	FedDAG	Average	Param.		✓		✓	C	✓	(40, 10)	✓
Qiu and Yang (45)	Bloom	Average	Param.		✓	✓	✓	C	✓	(100, -)	
Abyaneh et al. (46)	FedCDI	Average	Struct.		✓	✓		D	✓	(20, 8), (37, 4)	✓
Guo et al. (47)	FedECD	Average	Struct.		✓			D (*)		(27, 30)	
Guo et al. (48)	FedCSL	Average	Substruct.		✓			D (*)	✓	(5000, 12)	✓
Guo et al. (49)	FedACD	Average	Struct.		✓			C, D	✓	(37, 20)	✓
Yu et al. (50)	FedLCS	Average	Struct.		✓			C, D	✓	(400, 15)	
Yang et al. (27)	FMT-NBN	Average	Param.		✓		✓	C	✓	(20, 4)	
Torrijos et al. (51)	FedGES	Ensemble	Struct.	✓	✓			D (*)		(724, 100)	✓

Method. **Class:** “Constr.” (constraint-based method), “Score” (score-based, discrete (“D”) or continuous (“C”) score), “Average” (with a voting strategy server-side) or “Ensemble” (complex model fusion). **Share:** information sent to the central server, summary statistics (“stat.”), model score (“Score”), parameters (“Param.”), gradients (“grad.”) or the DAG structure (“Struct.”). The DAG could have edge weights (FedCSL) or an assigned belief (FedCDI). **Acycl.:** whether the algorithmic output is a DAG or requires additional processing. **Assumptions.** **Suff.:** whether the causal or pseudo-causal (“p”) sufficiency assumption is required. **Int.:** methods that can exploit interventional data. **SCM:** whether assumptions about the variables’ functional assignments are needed (such as linearity). **Application.** **Data:** handles discrete (“D”) or continuous (“C”) data; a (*) symbol is present whenever the extension to another data type is feasible, for instance, by changing local sufficient statistics. No method is available for mixed discrete/continuous data, but extensions may be possible. **RWD:** whether experiments on real-world data have been carried out. No method has been tested on survival or longitudinal data, but NOTEARS-PFL and FedDAG have been on imaging data sets. **Scalability:** the maximum number of (nodes, clients) on which the method has been tested.

ENDOMETrial carcinoma (PIPENDO) study (55) involved 19 gynaecological oncology clinics. Each clinic collected detailed data on patient demographics, histopathological features, administered treatments, and outcomes. Centralising these multi-institutional patient records in a single repository for analysis was not feasible due to stringent privacy regulations, lengthy ethical approval processes, and institutional data governance policies governing sensitive oncological data.

3.3 The choice of FCD method

Cancer research data often contain missing values, and the reasons for these missing values may not be random but rather informative about how data are collected. Zanga et al. (55) developed the first FCD algorithm that accounts for local missing-data distributions. In particular, the authors extended existing CD methods for missing data to the federated setting, relaxing the underlying assumption of a global missing mechanism because the missingness patterns in each clinic’s data differed, reflecting local clinical workflows, referral patterns, or documentation practices.

Failing to model federated missingness could bias the learned causal structure. While such an approach could be implemented using any of the methods mentioned in Section 2.2, score-based CD was the most convenient from a mathematical perspective.

3.4 How clinicians validated the learned causal graph

Involving clinicians was a vital part of this knowledge elicitation process: experts provided a set of causal and probabilistic statements that the CBN had to satisfy. Firstly, the model variables were selected based on which factors the clinical experts considered relevant for predicting survival and the presence of LNMs (56, 57). Secondly, the CBN’s causal pathways were validated against prior clinical knowledge through systematic literature reviews to corroborate the causal claims implied by the graph (55). Clinicians could also validate the CBN’s predicted probabilities for specific events against the statements they provided to assess its calibration.

3.5 Downstream clinical decision support

The learned CBN offers new insights into how treatments interact with the observed variables. This directly supports clinical decision-making by helping clinicians reason about the likely effect of alternative adjuvant treatments on LNM progression and patient survival for specific patient profiles, thereby moving toward a more personalised treatment strategy. FCD produced a CBN that generalised significantly better and achieved higher causal accuracy than those derived from single-source studies, consistent with other landmark studies using non-causal models (58–61). Furthermore, an overall sensitivity analysis showed that FCD also improved model robustness (55).

3.6 Contrasting pseudo-causal and causal sufficiency

Each of the 19 clinics collected the same set of patient variables, such as tumour grade, lymph node status, and treatment. Pseudo-causal sufficiency implies that within any single clinic's data set, there are no unmeasured common causes that jointly affect the recorded variables, even though clinics might differ in unobserved factors, such as local care pathways, that do not directly confound the measured relationships. In contrast, causal sufficiency would require that if a variable, such as a specific genetic mutation, caused, say, the choice of adjuvant therapy at one clinic and the survival outcome at another, that mutation would need to be measured and included at all 19 sites. This stronger condition is often impractical in federated healthcare settings where each institution may collect a slightly different set of covariates. Yet, pseudo-causal sufficiency still permits each clinic's local causal structure to be learned without requiring that every possible confounder is shared or recorded across the entire network.

4 Discussion

As noted earlier, increasing digitisation in medicine has led to an exponential growth in multivariate and multimodal observational data. Deciphering novel patterns and associations from these data sets has the potential to discover novel associations while validating what is known. However, non-causal machine learning approaches focus exclusively on predictive performance and are not designed for clinical reasoning (62).

FCD has the potential to decipher causal patterns from these observational data sets across, without explicit data sharing, overcoming sample size constraints. Beyond providing transparent and mechanistic views of observed systems, CBNs obtained through FCD provide the ground for *causal inference* (Supplementary Section S1.2). Under specific assumptions (63) such as correct measurement, model faithfulness, and the absence of unmeasured confounding (ensured by Assumption 4), CBNs may yield valid conclusions about treatment or exposure effects and risk-factor modification, with the potential to ultimately improve patient outcomes. Moreover, they may reduce the risk of mistaking correlation patterns for causation,

thereby providing a principled approach to investigating scenarios in which RCTs are infeasible and bridging the gap between observational and experimental evidence.

Still, they require large, well-balanced samples to produce transferable models that generalise across clinical and population cohorts and between different clinical cohorts. Crucially, the close relationship between several key concepts in causal inference and trial design (for instance, backdoor adjustment vs. randomisation, collider bias and stratification, path analysis and mediation) makes CBNs a powerful tool for boosting systematic reviews and RCTs.

In the remainder of this section, we will focus on opportunities and challenges of FCD in this context, as well as less-explored research areas.

4.1 Opportunities

Building on the effectiveness of FL (58, 59) and CD (62), FCD leverages decentralised *real-world evidence* (RWE) to ultimately improve drug discovery and patient outcomes, as observed above. By also pooling information from external control arms with real-world control patients from different institutions, the learned CBNs may be exploited to improve trial effectiveness, fairness, and data integration (64).

4.1.1 Clinical translation

The need for cross-institutional cooperation to foster medical research has been apparent for years. Experts' collaboration and data sharing enhance RWE, thereby strengthening the generalizability of analyses, reducing local biases, and promoting fairness (65). Popular architectures enabling multi-institutional storage and analytics of medical data include both centralised and federated architectures (Table 1).

Centralised architectures have been adopted to accelerate precision medicine research efforts, such as the *All of Us* initiative funded by the National Institutes of Health (NIH) (66). *All of Us* aims to provide equitable access to diverse data sets from participating institutions in a secure, centralised cloud environment, with centralised governance and a paid model to support analytics. Similarly, the European ELIXIR infrastructure supports the coordination and development of life sciences research, including a network of cross-domain experts, and promotes best practices for data analysis (67). The *Health Information Exchange* (HIE) also aims to improve care coordination while supporting surveillance (68). In addition, the *Global Alliance for Genomics and Health* (GA4GH) (69) aims to establish standards for sharing genomic and health data while ensuring privacy, security, and adherence to ethical values.

Unlike centralised architectures, federated architectures support querying of de-identified medical data, with the potential to accelerate clinical trials. An example is the *Shared Health Research Information Network* (SHRINE), funded by the *National Center for Advancing Translational Science* (NCATS) (70). SHRINE has been adopted by *Patient Centered Outcomes Research* (PCORI), enabling federated querying of healthcare data from PCORnet members. While these established projects

targeted cohort discovery, new developments are focusing on model learning through FL instead.

FL has gained traction in health research, with studies exploring different data types and applications (65). To this end, several initiatives promoted the creation of distributed infrastructures to host ready-to-use, private data. Among those, the *IDEA4RC* project focuses on creating a decentralised ecosystem of rare cancers data in Europe, compliant with the standardisation of common data models (71); *FeederNet* is a South Korean initiative supporting the development of a bio-health data ecosystem for federated analyses (72). In principle, these platforms could be leveraged for large-scale FCD.

However, limited applied research on FCD has been published to date. Zanga *et al.* (55) evaluated their FCD algorithm on decentralised data sets related to endometrial cancer. Their method is suited to incomplete data, especially when the missing mechanism differs across centres. Chen *et al.* (73) modelled gene expression by leveraging a time-dependent version of the CG. Zhang *et al.* (74) proposed a method to aggregate outputs from different CD algorithms and applied it to the treatment of acute kidney injury. Other applied works (75, 76) focused on *federated causal inference* by assuming the CGs are known instead of learning them. Furthermore, no existing distributed medical data set currently provides a ground-truth DAG for FCD benchmarking and experimentation.

Cross-institutional collaborations may aim to learn more robust CBNs through FCD than individual institutions can; the reference population may be clearly specified and taken into account (see Section 4.2.2) (77). Involving their respective experts in the CD process increases the trust in the learned CBN and fosters its adoption in clinical practice. However, it may present challenges in harmonising heterogeneous and potentially conflicting domain knowledge (78).

4.1.2 Efficiencies in RCTs

RCTs are limited in their ability to translate to real-world scenarios, such as cancer care, due to their inclusion and exclusion criteria (79). Non-causal models have shown mixed performance in providing information comparable to that from RCTs from RWE available in electronic health records (80). CBNs, on the other hand, can effectively harness RWE to better understand the real-world patient experience and outcomes thanks to their ability to learn digital twins of such settings (13). CBNs can identify relevant hypotheses to test and subpopulations with different treatment responses; provide prior estimates of effect sizes to identify relevant biomarkers for drug development (81); perform power analysis for sample size determination; and reduce the reuse of trial control arms, which limits transferability (82). In addition, they could be used more often to answer secondary analyses of existing RCT data (83).

Furthermore, CBNs can inform RCT design and assess its feasibility using observational RWE by emulating it within the “target trial” framework. Making sound design decisions about the length of follow-up, sample size, potential confounders, relevant subpopulations, minimisation of lost-to-follow-up, and expected cost of treatment delivery is essential for trials to provide useful information at their conclusion (84). At the same time, RCTs are increasingly expensive, reaching a median

budget of \$650,000, of which 27.4% is spent in the planning phase and 12.7% in the finalisation phase (85). Even so, about 70% of RCTs exceeded the budget by over 50%. Furthermore, 57% of RCTs had one or more substantial amendments, each costing \$141,000 to \$535,000 (86). Of these, 45% originated from protocol design flaws, inconsistencies in the protocol narrative, and infeasible eligibility criteria, which were “avoidable” and would have been flagged by a “target trial” CBN built from readily available RWE. Any additional planning efficiencies and reduced overruns resulting from more targeted designs are essential to advancing evidence-based healthcare in a cost-effective manner.

4.1.3 Fairness

Fairness is a well-documented issue in clinical trials. In addition to biases in historical control data (87), they may also be biased due to limitations in the trial design (88), self-selection among minorities (89), physicians’ implicit biases (90), and other factors. FCD offers a comprehensive solution to this class of issues because of its unique combination of FL and causal inference.

FL can potentially provide better coverage of a target population by allowing multiple institutions to pool information without sharing patient-level data, which is one of the major concerns that pushes minorities to self-select themselves out of trial enrollment. Indirectly, it would also reduce the impact of physicians’ implicit biases by giving them access to scrutinised, larger control arms (91).

In addition to mitigating known sources of biases, better population coverage provides CD with the data it needs to construct a CBN that captures the characteristics of the trial and its patients. Such a CBN serves as the foundation for achieving *counterfactual fairness* (92), which builds on counterfactuals to examine differential outcomes as a function of legally protected attributes. This is the most rigorous framework for fairness assessment and remediation in the literature; notably, it goes beyond descriptive statistics and allows for disambiguation between biases mediated by distinct causal pathways (93).

4.1.4 Synthetic data generation

Artificial intelligence models require large amounts of patient health data to build healthcare systems that improve diagnosis and understanding of disease. Issues with patient privacy and limited data availability make synthetic data generation an attractive solution to train those models (94).

CBNs are increasingly explored as generative models for synthetic data generation (95), with successful applications using the UK Clinical Practice Research Datalink (96). FCD can, at the same time, provide larger data sets to train more realistic CBNs for this purpose and serve as a reproducible data source for studying the performance, biases and systematic distortions of FCD algorithms. There is a strong research interest in investigating how a CBN, learned from multiple data sources and thus susceptible to various biases and limitations, can or must be used to generate reliable data.

4.2 Challenges

4.2.1 Data heterogeneity

Most clinical data are stored in heterogeneous data silos with different logical and physical structures, each one tailored to meet specific technical needs. Combining different databases may lead to inappropriate and biased results, even within the same institution (97).

A software solution to data harmonisation is given by *Beacon v2* (98), which provides a secure and flexible protocol for querying heterogeneous databases. However, Beacon does not support FL. Collaborative networks such as *Observational Health Data Sciences and Informatics* (OHDSI) (72) provide *common data models* (CDMs) to address data heterogeneity in medicine, thereby improving the generalizability of findings. *Common* refers to the data logical structure and shared vocabularies, fostering homogeneous semantics. A CDM enhances data owners' management, improves data users' interoperability while complying with privacy and security standards, and facilitates the development of standardised analytical tools. The *Observational Medical Outcomes Partnerships* (OMOP) CDM (99), developed within OHDSI, provides a longitudinal view of each patient. Its *oncology extension* further increases the information granularity to best support cancer research.

Translation into CDMs via *extract-transform-load* (ETL) processes poses several challenges. Only a subset of a centre's raw data is typically mapped, and ineffective cooperation among centres will impact the final data homogeneity. A lack of ETL expertise can also undermine the final data quality, leading to spurious cause-and-effect relationships in FCD. Both issues must be considered in FCD, and the ETL process must be well-documented for this purpose. Methodological research should also focus on developing FCD techniques for longitudinal and censored data (100), as well as modalities such as medical imaging (101) and natural language processing (64).

4.2.2 Distribution and semantic drifts

Variable distributions may differ between sources due to genetic population structure, heterogeneous environmental conditions, including unmeasured economic and social factors, and other covariates, such as age and sex (102). The observational or interventional nature of various parts of the data is also a difficult assumption to test. Heterogeneous machine calibration, both across centres and over time, differing institutional policies, and various inductive and deductive processes employed by physicians, may effectively amount to unobservable interventions. Note that interventions differ from distribution drift across populations, as the latter holds irrespective of whether an intervention has been performed (103). Also, there may be *semantic drift* when variables have different semantics across data sources.

Modelling incomplete data poses similar challenges. The reason behind the presence of missing values is inherently causal; each data source in FCD may have not only varying degrees of missingness but also different missingness mechanisms. Expert prior knowledge and specialised algorithms are required to handle them (55). Discarding incomplete observations or assuming an identical missingness mechanism will bias both ETL and FCD itself.

4.2.3 Aggregation bias

Learning CBNs from a single data source is a well-explored problem; the same is not true in FCD. Under [Assumption 1](#), one may average all the local distributions, weighting them by the local sample size. However, this would not account for distribution drifts across sources. Even a single low-quality or highly biased data set can disproportionately bias the global CBN when few clients are present or when the data are class-imbalanced. Druzdzal and Diez (77) demonstrated how not combining knowledge from different sources or using only data from the setting in which the CBN will be used is neither necessary nor sufficient to ensure model correctness. Hence, whether to adopt FL depends on the population(s) of interest, the research question, and the validity of assumptions. A sensible approach could be to perform a posteriori global model personalisation for individual clients, especially when the research question pertains to the local population. However, determining the value of each client's contribution to FCD is an open research question.

4.2.4 Privacy and security

FCD algorithms must implement privacy and security by design to fulfil legal and ethical regulations, considering that institutions may be located in different jurisdictions (104). Privacy involves safeguarding and controlling personal information, while security means protecting the system's integrity, availability, and confidentiality. The nature of attacks against privacy (e.g., *membership inference*) and security (e.g., *model poisoning*) depends on the threat model and the attacker's goal (105, 106). Adversaries can be both internal and external, but a defence-in-depth strategy that couples FL's architecture with cybersecurity best practices and encryption can render many such attacks impractical with current technologies (107).

Privacy and security should be addressed on multiple fronts (64). Membership inference attacks aim to determine whether a specific individual's data was included in the training set, model poisoning attacks maliciously manipulate training data or model updates to corrupt the learned model, and model inversion attacks seek to reconstruct sensitive training data from model outputs or parameters. Both can originate from participating parties as well as external adversaries.

Cybersecurity best practices will prevent the most common data leaks. FL is designed to avoid data sharing, thereby preserving privacy by design; most privacy attacks in the literature are only feasible under unrealistic assumptions (108). Indeed, FL's privacy-by-design approach aligns with the organisational measures required by GDPR (Articles 24, 25, and 32) and thereby facilitates compliance with both GDPR and the EU AI Act (107). Secure multi-party computation (36) (MPC, for encryption) and differential privacy (35) (DP, for privacy) can reach an acceptable trade-off between privacy, performance, and fairness (105). However, MPC incurs substantial computational overhead, and DP may degrade diagnostic or prognostic accuracy. Research to better understand the optimal cost-benefit application of these techniques is ongoing (109, 110). Notably, this privacy-accuracy trade-off is orthogonal to the federated learning framework; it is instead specific to the

type of causal model being federated, as demonstrated by differentially private causal discovery algorithms that exhibit similar accuracy losses even in centralised settings (105, 111).

The scarcity of existing literature limits a comprehensive analysis of privacy and security in FCD. Membership inference attacks are not possible by sharing an unparameterised CG, but the presence or absence of specific edges might reveal sensitive information about certain subpopulations (112). For this reason, Xu *et al.* (111) developed the *PrivPC* algorithm to perform CD under DP constraints. Murakonda *et al.* (113) provided a theoretical bound on the error and power of membership inference attacks on an exposed network. Zhang *et al.* (114) showed how to generate synthetic data from a CBN under DP. Rocchi *et al.* (115) used credal networks to mask the released model without degrading its utility, in contrast to DP-based protection techniques. From a causal perspective, privacy protection must be carefully balanced against the need to retain all necessary confounders; failing to collect or to expose key variables due to overly restrictive privacy measures can bias the causal discovery process (62). Further research may shed light on practical privacy leaks in existing methods, clarifying what can and cannot be shared, to better enforce security standards while maintaining the scalability and effectiveness of FCD approaches.

4.2.5 Software availability

The lack of comprehensive software solutions for FCD hinders its development and evaluation on real-world data. Some FCD algorithms have been implemented independently by the respective authors using different programming languages and file formats. Moreover, most implementations require data to be in a tabular format and lack interfaces to learn from CDMs, forcing researchers to use ETL to preprocess CDMs into a tabular form. For instance, Schulz *et al.* (116) conducted a CD study using a single OMOP data set, first ensuring data compatibility with the employed algorithm. The OHDSI community provides the ATLAS and HARES tools for extracting a cohort from an OMOP CDM and populating a data table. At this point, the OHDSI's ARACHNE system may orchestrate federated epidemiological studies, while the *Vantage6* platform (117) may facilitate more complex FL analyses.

Finally, FCD may incur greater communication and computational overhead than centralised efforts (76). CD is known to be resource-intensive for big data, so most experimentation in the FCD literature is limited to simple CBNs. Future research should thoroughly investigate the computational bottlenecks of FCD and engineer reliable, efficient infrastructures.

5 Conclusions

The digitisation of healthcare systems has accelerated the adoption of data-driven, evidence-based medicine. In this context, FCD is emerging as a pivotal paradigm for uncovering causal structures across decentralised data sets without sharing raw data, potentially offering informed RCT designs, improved fairness, and reliable synthetic data generation as presented in Section 4.1. FCD serves as an essential enabler of causal inference, which is equally critical for estimating personalised

treatment effects, but depends on first securing a reliable, privacy-preserving causal graph. Individualised care plans require estimating the causal effect of each decision on the patient, often from observational RWE data because randomising those decisions is often unethical or unscalable (62).

FCD is a relatively young field that continues to evolve, and its maturation is best guided by close interaction with clinicians and their domain expertise. Key challenges we described in Section 4.2 include data heterogeneity, population-level biases, the questionable validity of causal assumptions and the lack of quality software implementations. A further critical obstacle is privacy, as discussed in Section 4.2.4: regulations such as GDPR (118) and HIPAA (119) severely restrict the sharing of patient-level data, creating a risk of non-compliance or breach. Federated learning architectures with industry-standard IT security keep raw data local, essentially removing the privacy risks associated with data transfer while enabling collaborative causal discovery (108). Indeed, FL's privacy-by-design approach aligns with the organisational measures required by GDPR and facilitates compliance with the EU AI Act (120), as discussed in the context of bioinformatics in Malpetti *et al.* (107). Beyond privacy, federated systems must also address robustness to adversarial attacks, such as membership inference and model inversion, which can originate from both participating parties and external adversaries seeking to extract information about the training data. For settings with the highest risk of patient reidentification, such as rare disease studies, multi-party computation and differential privacy can further strengthen protections. However, multi-party computation incurs substantial computational overhead, and differential privacy may degrade diagnostic or prognostic accuracy. Institutions should adopt these techniques only after a careful risk assessment and a cost-benefit analysis specific to their clinical context.

To translate theoretical methods into safe clinical practice, a structured roadmap is required.

1. **Readiness and common ground.** All institutions should adopt a common data model, such as OMOP, to achieve logical, structural, and vocabulary standardisation (see Section 4.2.1). Local quality checks, semantic harmonisation, and explicit modelling of missing-data mechanisms are mandatory. In parallel, institutions should establish data use agreements, ethics approvals, and federated infrastructure (e.g., secure aggregation nodes). A few successful examples are presented in Section 4.1.1.
2. **Federated causal discovery.** Discovery algorithms execute under appropriate privacy guarantees, with the choice of technique guided by the institutional risk assessment and sensitivity of the data. Because causality formalises knowledge about cause-and-effect relationships and the assumptions we make regarding the data-generating process, prior clinical knowledge must be explicitly encoded as forbidden edges, temporal constraints, or Bayesian priors to guide the search, suppress spurious associations, and make the underlying assumptions transparent (121). The output is a CBN whose validity must be demonstrable. As CBNs are inherently transparent graphical models, they are directly interpretable using classical inference techniques (9, 122), obviating the need

for post hoc explainable AI (xAI) methods, which are typically designed for opaque black-box models and can inadvertently capture spurious correlations.

3. **Expert evaluation and iterative refinement.** A multidisciplinary panel should assess the CBN's plausibility and consistency with a focus on established associations and published evidence. Discrepancies trigger algorithmic refinement and re-evaluation before progressing, a best practice known as a human-in-the-loop model validation (123).
4. **Independent validation and fairness auditing.** The refined CBN undergoes external validation on independent cohorts not used during discovery. A thorough fairness audit should check for structural disparities (for example, missing or inverted edges) across race, gender, socioeconomic status, and other sensitive attributes, ensuring the causal model itself does not encode bias against specific subgroups (see also Section 4.1.3).
5. **Clinical integration and oversight.** The validated CBN is embedded into institutional clinical decision support systems under clear governance (see Section 4.1.1). This integration should enable future causal inference studies to estimate treatment effects without compromising privacy. Continuous monitoring for drift, fairness, and new evidence triggers regular updates.

A fundamental obstacle to realising this roadmap is the absence of standardised evaluation benchmarks and software, as reported in Section 4.2.5. Existing multi-centre clinical trials, with their well-characterised cohorts and known treatment effect estimates, offer a natural starting point for constructing semi-synthetic benchmarks with ground-truth causal structures (55). Federated cancer registries, which already operate across institutional and national boundaries, provide another rich source of real-world distributed data (124). The community must leverage such resources to collaboratively build a benchmark suite that includes a known (real or semi-synthetic) causal ground truth, multiple heterogeneous data sets with controlled distribution shifts and non-random missingness, and realistic privacy constraints. Such a resource would accelerate methodological progress and provide the evidence needed for regulatory acceptance.

By methodically addressing these practical and methodological challenges, FCD can evolve from a research prototype into a mature, safe, and clinically indispensable tool—and, in so doing, unlock the full potential of causal inference across distributed real-world data. The legal and technical landscape will continue to evolve, with legislation such as the EU AI Act progressively shaping how data are shared and used. FL-based approaches are inherently well-positioned to adapt to these developments because they already embed privacy-by-design principles.

Author contributions

NR: Conceptualization, Writing – original draft, Writing – review & editing. MS: Writing – review & editing, Conceptualization. AZ: Conceptualization, Writing – review &

editing. RN: Writing – review & editing, Conceptualization. FS: Project administration, Writing – review & editing, Funding acquisition.

Funding

The author(s) declared that financial support was received for this work and/or its publication. This work was supported by the MUR under the grant “Dipartimenti di Eccellenza 2023-2027” of the Department of Informatics, Systems and Communication of the University of Milano-Bicocca, Milan, Italy, and by the National Plan for NRRP Complementary Investments (Project n. PNC0000003 - AdvANced Technologies for Human-centrEd Medicine (ANTHEM)). Niccolò Rocchi is funded by NextGeneration EU and Fondazione IRCCS Istituto Nazionale dei Tumori, Milan, Italy, under the Horizon Europe project IDEA4RC-Project (101057048). This study is partially supported by the grant K25-2170 from the International Joint Usage/Research Center, the Institute of Medical Science, the University of Tokyo.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fdgth.2026.1846020/full#supplementary-material>

References

- Acosta JN, Falcone GJ, Rajpurkar P, Topol EJ. Multimodal biomedical AI. *Nat Med.* (2022) 28:1773–84. doi: 10.1038/s41591-022-01981-2
- Li R, Romano JD, Chen Y, Moore JH. Centralized and federated models for the analysis of clinical data. *Annu Rev Biomed Data Sci.* (2024) 7:179–99. doi: 10.1146/annurev-biodatasci-122220-115746
- Berlin JA, Golub RM. Meta-analysis as evidence: building a better pyramid. *JAMA.* (2014) 312:603–6. doi: 10.1001/jama.2014.8167
- Abernethy A, Adams L, Barrett M, Bechtel C, Brennan P, Butte A, et al. The promise of digital health: then, now, and the future. *NAM Perspect.* (2022) 6:1–24. doi: 10.31478/202206e
- Kivela M, Arenas A, Barthelemy M, Gleeson JP, Moreno Y, Porter MA. Multilayer networks. *J Complex Netw.* (2014) 2:203–71. doi: 10.1093/comnet/cnu016
- Zitnik M, Leskovec J. Predicting multicellular function through multi-layer tissue networks. *Bioinformatics.* (2017) 33:i190–8. doi: 10.1093/bioinformatics/btx252
- Chatfield C. *The Analysis of Time Series: Theory and Practice.* New York, New York: Springer (1975). doi: 10.1007/978-1-4899-2925-9
- Butte AJ, Tamayo P, Slonim D, Golub TR, Kohane IS. Discovering functional relationships between RNA expression and chemotherapeutic susceptibility using relevance networks. *Proc Natl Acad Sci.* (2000) 97:12182–6. doi: 10.1073/pnas.220392197
- Pearl J. *Causal Inference in Statistics.* Chichester: Wiley (2021).
- Carapito R, Li R, Helms J, Carapito C, Gujja S, Rolli V, et al. Identification of driver genes for critical forms of COVID-19 in a deeply phenotyped young patient cohort. *Sci Transl Med.* (2022) 14(628):eabj7521. doi: 10.1126/scitranslmed.abj7521
- Brown AA, Fernandez-Tajes JJ, Hong M, Brorsson CA, Koivula RW, Davtian D, et al. Genetic analysis of blood molecular phenotypes reveals common properties in the regulatory networks affecting complex traits. *Nat Commun.* (2023) 14:5062. doi: 10.1038/s41467-023-40569-3
- Petersen AH, Ekström CT, Spirtes P, Osler M. Constructing causal life-course models: comparative study of data-driven and theory-driven approaches. *Am J Epidemiol.* (2023) 192:1917–27. doi: 10.1093/aje/kwad144
- Papanastasiou G, Scutari M, Tachdjian R, Hernandez-Trujillo V, Raasch J, Billmeyer K, et al. Causal modeling in large-scale data to improve identification of adults at risk for combined and common variable immunodeficiencies. *npj Digit Med.* (2025) 8:361. doi: 10.1038/s41746-025-01761-5
- Balordi A, Bernasconi A, Andreotti A, Guzzinati S, Cabañas De Paz R, Zanga A. On counterfactual explanations of cardiovascular risk in adolescent and young adult breast cancer survivors. *J Med Syst.* (2025) 49:140. doi: 10.1007/s10916-025-02273-1
- Delucchi M, Spinner GR, Scutari M, Bijlenga P, Morel S, Friedrich CM, et al. Bayesian network analysis reveals the interplay of intracranial aneurysm rupture risk factors. *Comput Biol Med.* (2022) 147:105740. doi: 10.1016/j.combiomed.2022.105740
- Lim RG, Al-Dalahmah O, Wu J, Gold MP, Reidling JC, Tang G, et al. Huntington disease oligodendrocyte maturation deficits revealed by single-nucleus RNAseq are rescued by thiamine-biotin supplementation. *Nat Commun.* (2022) 13:7791. doi: 10.1038/s41467-022-35388-x
- Belyaeva A, Cammarata L, Radhakrishnan A, Squires C, Yang KD, Shivashankar GV, et al. Causal network models of SARS-CoV-2 expression and aging to identify candidates for drug repurposing. *Nat Commun.* (2021) 12:1024. doi: 10.1038/s41467-021-21056-z
- Zanga A, Bernasconi A, Lucas PJF, Pijnenborg H, Reijnen C, Scutari M, et al. Causal discovery with missing data in a multicentric clinical study. In: *Proceedings of the 21st International Conference on Artificial Intelligence in Medicine (AIME23)* (2023). p. 40–4. doi: 10.1007/978-3-031-34344-5_5
- Deshpande A, Chu L, Stewart R, Gitter A. Network inference with granger causality ensembles on single-cell transcriptomics. *Cell Rep.* (2022) 38:110333. doi: 10.1016/j.celrep.2022.110333
- Scutari M, Kerob D, Krutmann J, Salah S. Causal networks of infodemiological data: modelling dermatitis. In: *Proceedings of the 23rd International Conference on Artificial Intelligence in Medicine (AIME25)* (2025). p. 397–407. doi: 10.1007/978-3-031-95838-0_39
- Mohammad-Taheri S, Zucker J, Hoyt CT, Sachs K, Tewari V, Ness R, et al. Do-calculus enables estimation of causal effects in partially observed biomolecular pathways. *Bioinformatics.* (2022) 38:i350–8. doi: 10.1093/bioinformatics/btac251
- Sturma N, Squires C, Drton M, Uhler C. Unpaired multi-domain causal representation learning. *Adv Neural Inf Process Syst.* (2023) 36:34465–92. doi: 10.52202/075280-1496
- Spirtes P, Glymour C, Scheines R. *Causation, Prediction, and Search.* Cambridge, Massachusetts: MIT Press (2000). doi: 10.7551/mitpress/1754.001.0001
- Fang Y. *Causal Inference in Pharmaceutical Statistics.* Boca Raton, Florida: Chapman and Hall (2024).
- Prosperi M, Guo Y, Sperrin M, Koopman JS, Min JS, He X, et al. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nat Mach Intell.* (2020) 2:369–75. doi: 10.1038/s42256-020-0197-y
- Tillman RE, Eberhardt F. Learning causal structure from multiple datasets with similar variable sets. *Behaviormetrika.* (2014) 41:41–64. doi: 10.2333/bhmk.41.41
- Yang X, Niu B, Lan T, Zhang C. Federated multi-task bayesian network learning in the presence of overlapping and distinct variables. *IJSE Trans.* (2025) 57:773–87. doi: 10.1080/24725854.2024.2443177
- Ramsey J, Hanson S, Hanson C, Halchenko Y, Poldrack R, Glymour C. Six problems for causal inference from fMRI. *Neuroimage.* (2010) 49:1545–58. doi: 10.1016/j.neuroimage.2009.08.065
- Schölkopf B. *Causality for Machine Learning.* 1st ed. New York, NY, USA: Association for Computing Machinery (2022). p. 765–804. chap. 39. doi: 10.1145/3501714.3501755
- Zanga A, Ozkirimli E, Stella F. A survey on causal discovery: theory and practice. *Int J Approx Reason.* (2022) 151:101–29. doi: 10.1016/j.ijar.2022.09.004
- Mooij JM, Magliacane S, Claassen T. Joint causal inference from multiple contexts. *J Mach Learn Res.* (2020) 21:1–108.
- Ajayi OT, Cheng Y. Decentralized learning of Bayesian networks from private data with applications to global pandemic. In: *2023 IEEE 43rd International Conference on Distributed Computing Systems (ICDCS)* (2023). p. 961–2. doi: 10.1109/icdcs57875.2023.00105
- Gao E, Chen J, Shen L, Liu T, Gong M, Bondell H. FedDAG: federated DAG structure learning. *Trans Mach Learn Res.* (2023).
- Mian O, Kaltenpoth D, Kamp M, Vreeken J. Nothing but regrets—privacy-preserving federated causal discovery. *Proc Mach Learn Res.* (2023) 206 (AISTATS):8263–78.
- Dwork C, McSherry F, Nissim K, Smith A. Calibrating noise to sensitivity in private data analysis. *Theory Cryptogr.* (2006) 7(3):265–84. doi: 10.29012/jpc.v7i3.405
- Wang Z, Ma P, Wang S. Towards practical federated causal structure learning. *Mach Learn Knowl Discov Databases Res Track.* (2023) 14170:351–67. doi: 10.1007/978-3-031-43415-0_21
- Li L, Ng I, Luo G, Huang B, Chen G, Liu T, et al. Federated causal discovery from heterogeneous data. In: Kim B, Yue Y, Chaudhuri S, Fragiadakis K, Khan M, Sun Y, editors. *International Conference on Learning Representations.* Vol. 2024 (2024). p. 10901–34.
- Mian O, Kaltenpoth D, Kamp M. Regret-based federated causal discovery. *Proc Mach Learn Res.* (2022) 185 (KDD Workshop on Causal Discovery):61–9.
- Yang D, He X, Wang J, Yu G, Domeniconi C, Zhang J. Federated causality learning with explainable adaptive optimization. *Proc AAAI Conf Artif Intell.* (2024) 38:16308–15. doi: 10.1609/aaai.v38i15.29566
- Liu S, Yan X, Guo X, Qi S, Wang H, Chang X. Federated Bayesian network learning from multi-site data. *J Biomed Inform.* (2025) 163:104784. doi: 10.1016/j.jbi.2025.104784
- Ng I, Zhang K. Towards federated Bayesian network structure learning with continuous optimization. *Proc Mach Learn Res.* (2022) 151 (AISTATS):8095–111.
- Liu C, Wang S, Wang Z, Zhuang J, Li W, Hao Z, et al. Federated causal structure learning with a bi-level optimization model. In: *2024 IEEE 14th International Conference on Cyber Technology in Automation, Control, and Intelligent Systems (Cyber)* (2024). p. 708–13. doi: 10.1109/cyber63482.2024.10749349
- Ye Q, Amini AA, Zhou Q. Federated learning of generalized linear causal networks. *IEEE Trans Pattern Anal Mach Intell.* (2024) 46:6623–36. doi: 10.1109/TPAMI.2024.3381860
- Huang J, Guo X, Yu K, Cao F, Liang J. Towards privacy-aware causal structure learning in federated setting. *IEEE Trans Big Data.* (2023) 9:1525–35. doi: 10.1109/TBDATA.2023.3285477
- Qiu C, Yang K. Interventional causal structure discovery over graphical models with convergence and optimality guarantees. *IEEE Trans Netw Sci Eng.* (2025) 12:156–72. doi: 10.1109/TNSE.2024.3487301
- Abyaneh A, Scherrer N, Schwab P, Bauer S, Schölkopf B, Mehrjou A. Federated causal discovery from interventions. *arXiv* (2022).
- Guo X, Yi L, Wu X, Yu K, Wang G. Enhancing causal discovery in federated settings with limited local samples. In: Yu H, Li X, Xu Z, Goebel R, King I, editors. *Federated Learning in the Age of Foundation Models - FL 2024 International Workshops.* Cham: Springer Nature Switzerland (2025). p. 164–79. doi: 10.1007/978-3-031-82240-7_12
- Guo X, Yu K, Liu L, Li J. FedCSL: a scalable and accurate approach to federated causal structure learning. *Proc AAAI Conf Artif Intell.* (2024) 38:12235–43. doi: 10.1609/aaai.v38i11.29113
- Guo X, Yu K, Wang H, Cui L, Yu H, Li X. Sample quality heterogeneity-aware federated causal discovery through adaptive variable space selection. In: *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence* (2024). p. 4071–9.
- Yu K, Rong C, Wang H, Cao F, Liang J. Federated local causal structure learning. *Sci China Inf Sci.* (2025) 68(3):132105:1–15. doi: 10.1007/s11432-023-4203-6

51. Torrijos P, Gámez JA, Puerta JM. Fedges: a federated learning approach for bayesian network structure learning. *Mach Learn.* (2025) 115:9. doi: 10.1007/s10994-025-06939-2
52. Chickering DM. Optimal structure identification with greedy search. *J Mach Learn Res.* (2003) 3:507–54.
53. Puerta JM, Aledo JA, Gámez JA, Laborda JD. Efficient and accurate structural fusion of bayesian networks. *Inf Fus.* (2021) 66:155–69. doi: 10.1016/j.inffus.2020.09.003
54. Bray F, Ferlay J, Soerjomataram I, Siegel RL, Torre LA, Jemal A. Global cancer statistics 2018: globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA Cancer J Clin.* (2018) 68:394–424. doi: 10.3322/caac.21492
55. Zanga A, Bernasconi A, Lucas PJF, Pijnenborg JM, Reijnen C, Scutari M, et al. Federated causal discovery with missing data in a multicentric study on endometrial cancer. *J Biomed Inform.* (2025) 169:104887. doi: 10.1016/j.jbi.2025.104877
56. Reijnen C, Gogou E, Visser NCM, Engerud H, Ramjith J, van der Putten LJM, et al. Preoperative risk stratification in endometrial cancer (endorisk) by a bayesian network model: a development and validation study. *PLoS Med.* (2020) 17: e1003111. doi: 10.1371/journal.pmed.1003111
57. Zanga A, Bernasconi A, Lucas PJF, Pijnenborg H, Reijnen C, Scutari M, et al. *Causal Discovery with Missing Data in a Multicentric Clinical Study*. Cham: Springer Nature Switzerland (2023). p. 40–4.
58. Roth HR, Chang K, Singh P, Neumark N, Li W, Gupta V, et al. Federated learning for breast density classification: a real-world implementation. In: *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning: 2nd MICCAI Workshop, DART 2020, and 1st MICCAI Workshop, DCL 2020* (2020). p. 181–91. doi: 10.1007/978-3-030-60548-3_18
59. Pati S, Baid U, Edwards B, Sheller M, Wang S, Reina GA, et al. Federated learning enables big data for rare cancer boundary detection. *Nat Commun.* (2022) 13:7346. doi: 10.1038/s41467-022-33407-5
60. Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, et al. Federated learning for predicting clinical outcomes in patients with COVID-19. *Nat Med.* (2021) 27:1735–43. doi: 10.1038/s41591-021-01506-3
61. Heyndrickx W, Mervin L, Morawietz T, Sturm N, Friedrich L, Zalewski A, et al. MELLODDY: cross-pharma federated learning at unprecedented scale unlocks benefits in qsar without compromising proprietary information. *J Chem Inf Model.* (2023) 64:2331–44. doi: 10.1021/acs.jcim.3c00799
62. Glocker B, Musolesi M, Richens J, Uhler C. Causality in digital medicine. *Nat Commun.* (2021) 12:5471. doi: 10.1038/s41467-021-25743-9
63. Pearl J, Glymour M, Jewell NP. *Causal Inference in Statistics: A Primer*. Chichester: Wiley (2016).
64. Xu J, Glucksberg BS, Su C, Walker P, Bian J, Wang F. Federated learning for healthcare informatics. *J Healthc Inf Res.* (2021) 5:1–19. doi: 10.1007/s41666-020-00082-4
65. Teo ZL, Jin L, Liu N, Li S, Miao D, Zhang X, et al. Federated machine learning in healthcare: a systematic review on clinical applications and technical architecture. *Cell Rep Med.* (2024) 5:101419. doi: 10.1016/j.xcrm.2024.101419
66. The All of Us Research Program Investigators. The “all of us” research program. *N Engl J Med.* (2019) 381:668–76. doi: 10.1056/NEJMr1809937
67. Elixir. Elixir Scientific Programme, 2024–2028 (2023). Available online at: <https://elixir-europe.org/how-we-work/scientific-programme>.
68. Hersh WR, Totten AM, Eden KB, Devine B, Gorman P, Kassakian SZ, et al. Outcomes from health information exchange: systematic review and future research needs. *JMIR Med Inform.* (2015) 3:e39. doi: 10.2196/medinform.5215
69. Rehm HL, Page AJ, Smith L, Adams JB, Alterovitz G, Babb LJ, et al. GA4GH: international policies and standards for data sharing across genomic research and healthcare. *Cell Genom.* (2021) 1:100029. doi: 10.1016/j.xgen.2021.100029
70. Weber GM, Murphy SN, McMurphy AJ, MacFadden D, Nigrin DJ, Churchill S, et al. The shared health research information network (shrine): a prototype federated query tool for clinical data repositories. *J Am Med Inform Assoc.* (2009) 16:624–30. doi: 10.1197/jamia.M3191
71. Elixir. *Intelligent ecosystem to improve the governance, the sharing and the re-use of health data for rare cancers* (Tech. rep.). European Commission (2022). doi: 10.3030/101057048
72. You SC, Lee S, Choi B, Park RW. Establishment of an international evidence sharing network through common data model for cardiovascular research. *Korean Circ J.* (2022) 52:853. doi: 10.4070/kcj.2022.0294
73. Chen J, Ma Y, Yue X. Federated learning of dynamic bayesian network via continuous optimization from time series data. *IEEE Trans Artif Intell.* (2026) 7:2693–709. doi: 10.1109/TAI.2025.3622545
74. Zhang M, Zhang X, Dai M, Wu L, Liu K, Wang H, et al. Development and validation of a multi-causal investigation and discovery framework for knowledge harmonization (mindmerge): a case study with acute kidney injury risk factor discovery using electronic medical records. *Int J Med Inform.* (2024) 191:105588. doi: 10.1016/j.ijmedinf.2024.105588
75. Xiong R, Koenecke A, Powell M, Shen Z, Vogelstein JT, Athey S. Federated causal inference in heterogeneous observational data. *Stat Med.* (2023) 42:4418–39. doi: 10.1002/sim.9868
76. Meurisse M, Estupiñán-Romero F, González-Galindo J, Martínez-Lizaga N, Royo-Sierra S, Saldner S, et al. Federated causal inference based on real-world observational data sources: application to a SARS-CoV-2 vaccine effectiveness assessment. *BMC Med Res Methodol.* (2023) 23:248. doi: 10.1186/s12874-023-02068-3
77. Druzdzel MJ, Díez FJ. Combining knowledge from different sources in causal probabilistic models. *J Mach Learn Res.* (2003) 4:295–316.
78. Borboudakis G, Tsamardinos I. Incorporating causal prior knowledge as path-constraints in Bayesian networks and maximal ancestral graphs. In: *Proceedings of the 29th International Conference on International Conference on Machine Learning* (2012). p. 427–34.
79. Miksad RA, Abernethy AP. Harnessing the power of real-world evidence (RWE): a checklist to ensure regulatory-grade data quality. *Clin Pharmacol Ther.* (2017) 103:202–5. doi: 10.1002/cpt.946
80. Franklin JM, Patorno E, Desai RJ, Glynn RJ, Martin D, Quinto K, et al. Emulating randomized clinical trials with nonrandomized real-world evidence studies: first results from the RCT duplicate initiative. *Circulation.* (2021) 143:1002–13. doi: 10.1161/circulationaha.120.051718
81. Wu JQ, Horeweg N, de Bruyn M, Nout RA, Jürgenliemk-Schulz IM, Lutgens LCHW, et al. Automated causal inference in application to randomized controlled clinical trials. *Nat Mach Intell.* (2022) 4:436–44. doi: 10.1038/s42256-022-00470-y
82. Marzano L, Darwich AS, Dan A, Tendler S, Lewensohn R, Petris LD, et al. Exploring the discrepancies between clinical trials and real-world data: a small-cell lung cancer study. *Clin Transl Sci.* (2024) 17:e13909. doi: 10.1111/cts.13909
83. Farmer RE, Kounali D, Walker AS, Savović J, Richards A, May MT, et al. Application of causal inference methods in the analyses of randomised controlled trials: a systematic review. *Trials.* (2018) 19:23. doi: 10.1186/s13063-017-2381-x
84. Glick HA, Doshi JA, Sonnad SS, Polsky D. *Economic Evaluation in Clinical Trials*. 2nd ed. Oxford: Oxford University Press (2015).
85. Griessbach A, Speich B, Amstutz A, Hausheer L, Covino M, Wnfried Ramirez H, et al. Resource use and costs of investigator-sponsored randomized clinical trials in Switzerland, Germany, and the United Kingdom: a meta-research study. *J Clin Epidemiol.* (2024) 176:111536. doi: 10.1016/j.jclinepi.2024.111536
86. Getz KA, Stergiopoulos S, Short M, Surgeon L, Krauss R, Pretorius S, et al. The impact of protocol amendments on clinical trial performance and cost. *Ther Innov Regul Sci.* (2016) 50:436–41. doi: 10.1177/2168479016632271
87. Viele K, Berry S, Neuenschwander B, Amzal B, Chen F, Enas N, et al. Use of historical control data for assessing treatment effects in clinical trials. *Pharm Stat.* (2013) 13:41–54. doi: 10.1002/pst.1589
88. Kahan BC, Rehal S, Cro S. Risk of selection bias in randomised trials. *Trials.* (2015) 16:405. doi: 10.1186/s13063-015-0920-x
89. Catz DS, Green NS, Tobin JN, Lloyd-Puryear MA, Kyler P, Umemoto A, et al. Attitudes about genetics in underserved, culturally diverse populations. *Public Health Genomics.* (2005) 8:161–72. doi: 10.1159/000086759
90. Chapman EN, Kaatz A, Carnes M. Physicians and implicit bias: how doctors may unwittingly perpetuate health care disparities. *J Gen Intern Med.* (2013) 28:1504–10. doi: 10.1007/s11606-013-2441-1
91. Ogier du Terrail J, Klopstein Q, Li H, Mayer I, Loiseau N, Hallal M, et al. FedECA: federated external control arms for causal inference with time-to-event data in distributed settings. *Nat Commun.* (2025) 16:7496. doi: 10.1038/s41467-025-62525-z
92. Kusner M, Loftus J, Russell C, Silva R. Counterfactual fairness. *Adv Neural Inf Process Syst.* (2017) 31:4066–76.
93. Gultchin L, Guo S, Malek A. Pragmatic fairness: developing policies with outcome disparity control. *Proc Mach Learn Res.* (2024) 236 (CLeaR):243–64.
94. Pezoulas VC, Zaridis DI, Mylona E, Androutsos C, Apostolidis K, Tachos NS, et al. Synthetic data generation methods in healthcare: a review on open-source tools and methods. *Comput Struct Biotechnol J.* (2024) 23:2892–910. doi: 10.1016/j.csbj.2024.07.005
95. Tracy DH, Sherman J, Baron M. Abstract 7373: generative bayesian networks for augmentation of molecular data from commercial genetics panels. *Cancer Res.* (2024) 84:7373–. doi: 10.1158/1538-7445.AM2024-7373
96. Tucker A, Wang Z, Rotalinti Y, Myles P. Generating high-fidelity synthetic patient data for assessing machine learning healthcare software. *npj Digit Med.* (2020) 3:147. doi: 10.1038/s41746-020-00353-9
97. Madigan D, Ryan PB, Schuemie M, Stang PE, Overhage JM, Hartzema AG, et al. Evaluating the impact of database heterogeneity on observational study results. *Am J Epidemiol.* (2013) 178:645–51. doi: 10.1093/aje/kwt010
98. Rambla J, Baudis M, Ariosa R, Beck T, Fromont LA, Navarro A, et al. Beacon v2 and beacon networks: a “lingua franca” for federated data discovery in biomedical genomics, and beyond. *Hum Mutat.* (2022) 43(6):791–9. doi: 10.1002/humu.24369

99. Voss EA, Makadia R, Matcho A, Ma Q, Knoll C, Schuemie M, et al. Feasibility and utility of applications of the common data model to multiple, disparate observational health databases. *J Am Med Inform Assoc.* (2015) 22:553–64. doi: 10.1093/jamia/ocu023
100. Rocchi N, Zanga A, Bernasconi A, Gronchi A, Callegaro D, Borghi A, et al. A causal discovery workflow for rare diseases: experts-in-the-loop analysis of sparse longitudinal data. *J Med Syst.* (2026) 50:9. doi: 10.1007/s10916-025-02327-4
101. Castro DC, Walker I, Glocker B. Causality matters in medical imaging. *Nat Commun.* (2020) 11:3673. doi: 10.1038/s41467-020-17478-w
102. Zhao Y, Li M, Lai L, Suda N, Civin D, Chandra V. Federated learning with non-IID data. *Corr abs/1806.00582* (2018).
103. Li A, Jaber A, Bareinboim E. Causal discovery from observational and interventional data across multiple environments. *Adv Neural Inf Process Syst.* (2023) 36:16942–56. doi: 10.52202/075280-0741
104. Truong N, Sun K, Wang S, Guitton F, Guo Y. Privacy preservation in federated learning: an insightful survey from the GDPR perspective. *Comput Secur.* (2021) 110:102402. doi: 10.1016/j.cose.2021.102402
105. Dwork C, Smith A, Steinke T, Ullman J. Exposed! a survey of attacks on private data. *Annu Rev Stat Appl.* (2017) 4:61–84. doi: 10.1146/annurev-statistics-060116-054123
106. Blanco-Justicia A, Domingo-Ferrer J, Martínez S, Sánchez D, Flanagan A, Tan KE. Achieving security and privacy in federated learning systems: survey, research challenges and future directions. *Eng Appl Artif Intell.* (2021) 106:104468. doi: 10.1016/j.engappai.2021.104468
107. Malpetti D, Scutari M, Gualdi F, van Setten J, van der Laan S, Haitjema S, et al. Technical and legal aspects of federated learning in bioinformatics: applications, challenges and opportunities. *Front Digit Health.* (2025) 7:1644291. doi: 10.3389/fdgth.2025.1644291
108. Wainakh A, Zimmer E, Subedi S, Keim J, Grube T, Karuppayah S, et al. Federated learning attacks revisited: a critical discussion of gaps, assumptions, and evaluation setups. *Sensors.* (2022) 23:31. doi: 10.3390/s23010031
109. Ficek J, Wang W, Chen H, Dagne G, Daley E. Differential privacy in health research: a scoping review. *J Am Med Inform Assoc.* (2021) 28:2269–76. doi: 10.1093/jamia/ocab135
110. Ballhausen H, Corradini S, Belca C, Bogdanov D, Boldrini L, Bono F, et al. Privacy-friendly evaluation of patient data with secure multiparty computation in a European pilot study. *npj Digit Med.* (2024) 7:280. doi: 10.1038/s41746-024-01293-4
111. Xu D, Yuan S, Wu X. Differential privacy preserving causal graph discovery. In: *2017 IEEE Symposium on Privacy-Aware Computing (PAC)* (2017). p. 60–71. doi: 10.1109/pac.2017.24
112. Ma P, Ji Z, Pang Q, Wang S. NoLeaks: differentially private causal discovery under functional causal model. *IEEE Trans Inf Forens Secur.* (2022) 17:2324–38. doi: 10.1109/TIFS.2022.3184263
113. Kumar Murakonda S, Shokri R, Theodorakopoulos G. Quantifying the privacy risks of learning high-dimensional graphical models. *Proc Mach Learn Res.* (2021) 130 (AISTATS):2287–95.
114. Zhang J, Cormode G, Procopiuc CM, Srivastava D, Xiao X. PrivBayes: private data release via Bayesian networks. *ACM Trans Database Syst.* (2017) 42:1–41. doi: 10.1145/3134428
115. Rocchi N, Stella F, de Campos C. Towards privacy-aware Bayesian networks: a credal approach. In: *ECAI 2025*, Vol. 413. IOS Press (2025). p. 5043–50. doi: 10.3233/faia251419
116. Schulz NA, Carus J, Wiederhold AJ, Johannis O, Peters F, Rath N, et al. Learning debiased graph representations from the omop common data model for synthetic data generation. *BMC Med Res Methodol.* (2024) 24:136. doi: 10.1186/s12874-024-02257-8
117. Smits D, Van Beusekom B, Martin F, Veen L, Geleijnse G, Moncada-Torres A. An improved infrastructure for privacy-preserving analysis of patient data. In: *Advances in Informatics, Management and Technology in Healthcare*. IOS Press (2022). p. 144–7. doi: 10.3233/shiti220682
118. European Union. General Data Protection Regulation (GDPR). *Official Journal of the European Union*. Document number 32016R0679 (2016).
119. US Congress. Health Insurance Portability and Accountability Act of 1996 (1996). Pub. L. No. 104-191, 110 Stat. 1936.
120. European Union. AI Act. *Official Journal of the European Union*. Document number 32024R1689 (2024).
121. Constantinou AC, Guo Z, Kitson NK. The impact of prior knowledge on causal structure learning. *Knowl Inf Syst.* (2023) 65:3385–434. doi: 10.1007/s10115-023-01858-x
122. Koller D, Friedman N. *Probabilistic Graphical Models*. Cambridge, Massachusetts: MIT Press (2010).
123. Scutari M, Malvestio M. *The Pragmatic Programmer for Machine Learning: Engineering Analytics and Data Science Solutions*. Boca Raton, Florida: Chapman & Hall (2023).
124. De Angelis R, Demuru E, Baili P, Troussard X, Katalinic A, Chirlaque Lopez MD, et al. Complete cancer prevalence in Europe in 2020 by disease duration and country (Eurocare-6): a population-based study. *Lancet Oncol.* (2024) 25:293–307. doi: 10.1016/S1470-2045(23)00646-0