sustainability

MDPI

Systematic Review

Transfer-Aware Anomaly Detection for Sustainable AIoT-Enabled Water Infrastructures: A Systematic Review of Graph-Based Models, Tabular Foundation Models, and Privacy-Preserving Adaptation

Bilal Boukoftane *, Antonio Pita-Lozano and Juan Miguel Navarro Ruiz *

Research Group in Advanced Telecommunications (GRITA), UCAM Universidad Católica de Murcia, Avenida de los Jerónimos 135, 30107 Guadalupe, Murcia, Spain

* Correspondence: bboukoftane@alu.ucam.edu (B.B.); jmnavarro@ucam.edu (J.M.N.R.)

Abstract

Early detection of leaks, bursts, process faults, and cyber-physical attacks is a direct lever for the sustainable management of water infrastructures: it reduces non-revenue water, the energy spent on treating and pumping lost water, and service interruptions (Sustainable Development Goal 6). Anomaly detectors for water distribution networks (WDNs) and wastewater treatment plants (WWTPs) are increasingly deployed through Artificial Intelligence of Things (AIoT) architectures, yet a detector that performs well on one system cannot be assumed to transfer to another with a different topology, sensor layout, demand pattern, and label availability. This systematic review examines the literature through this transfer problem, with graph neural networks (GNNs) as topology-aware models, tabular foundation models (TFMs; in the retrieved evidence, the TabPFN family) as candidate local screeners, federated learning as privacy-preserving adaptation, and digital twins and explainable artificial intelligence as validation and trust layers. It follows PRISMA 2020 and is registered retrospectively in the Open Science Framework (DOI 10.17605/OSF.IO/D7RXU). Peer-reviewed articles, theses, and preprints in English or Spanish were identified through modular block searches (Google Scholar, publisher sites, three Spanish repositories), author-supplied sources, and citation chasing, verified against primary records on 6–7 September 2026, screened in two stages, coded with a predefined framework, appraised with a six-item checklist, and synthesized descriptively because tasks, datasets, and metrics are not commensurable. The corpus comprises 150 unique works: 76 water-specific experimental studies, 21 benchmark or testbed studies, and 53 method, review, or contextual works. Graph-based approaches have the stronger water-specific evidence for topology-dependent tasks in WDNs, whereas TabPFN-family TFMs rest on general tabular-learning evidence and one water-adjacent prediction study, with no validated water anomaly-detection use. Only 6 of the 76 experimental studies evaluate an explicit transfer mechanism, federated evidence is limited to one water study, and no study reports an operator evaluation of explanations. The review contributes a transfer taxonomy with study counts, an evidence-class scheme, a gap map of demonstrated and proposed components on a transfer-aware AIoT layout, and an evaluation agenda based on source–target protocols, and identifies the evidence utilities need before transferable detectors can support sustainable operation at scale.

Keywords: anomaly detection; water distribution networks; sustainable water management; non-revenue water; transfer learning; AIoT; graph neural networks; tabular foundation models; TabPFN; federated learning

Check for updates

Academic Editor: Hossein Bonakdari, Bahram Gharabaghi, Silvio José Gumiere

Received: 5 October 2026
Revised: 22 October 2026
Accepted: 29 October 2026
Published: 9 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Water distribution networks (WDNs) and wastewater treatment plants (WWTPs) are exposed to anomalies that range from pipe bursts, leaks, and pump or valve faults to sensor failures, water-quality deviations, and cyber-physical attacks on supervisory control systems [1–3]. Late or missed detection translates into water loss, service interruption, public-health exposure, and damage to assets, while false alarms consume the limited attention of operators. Detecting these events reliably is therefore an operational requirement rather than an analytical convenience [4,5].

The stakes are also those of sustainability. Leaks and bursts waste treated water together with the energy used to pump and treat it, process faults in WWTPs release untreated loads to receiving waters, and attacks on control systems threaten service continuity; detecting these anomalies early is therefore one of the most direct contributions that AIoT can make to the sustainable management of water infrastructures, in line with Sustainable Development Goal 6 and with the resource-efficiency aims of smart-water programs. Transferability matters for the same reason: a detector that can be adapted to a new network with little labeled data lowers the cost of adoption for the many small and medium utilities that cannot train models from scratch, so that the sustainability benefit is not confined to the few well-instrumented systems in which most of the evidence reviewed here was produced.

The sensing basis for detection has changed. Pressure, flow, acoustic, water-quality, and smart-meter streams are now collected through Internet of Things (IoT) deployments that increasingly embed inference close to the sensor, a configuration commonly described as the Artificial Intelligence of Things (AIoT) [6–8]. AIoT settings impose constraints that offline benchmark studies rarely capture: intermittent connectivity, bandwidth and energy budgets on battery-powered nodes, latency requirements for local control, and restrictions on moving raw operational data between sites [8–10].

The data produced by these systems have two structures that call for different model families. Hydraulic and cyber-physical anomalies propagate along the pipe network, so the network topology is informative and graph neural networks (GNNs) can exploit it [11–13]. At the same time, much of the data at a single node or gateway is tabular: a small window of multivariate readings with no explicit topology. For this second structure, tabular foundation models (TFMs) such as TabPFN, which are pre-trained on synthetic tasks and perform in-context inference on small tables, have been proposed as fast local screeners [14,15]. Their evidence in water infrastructures is, however, thin, and this review treats them as emerging candidates rather than as validated water methods.

The key problem addressed in this review is not only whether anomalies can be detected inside one dataset. The harder problem is whether a detector trained or designed in one water system can be trusted in another system with a different topology, sensor placement, demand pattern, label availability, and operating conditions. Each utility, district metered area (DMA), or treatment train is a distinct domain, labeled anomalies are rare, and simulators can produce scenarios that are unsafe to reproduce in the field [13,16,17]. Transfer learning describes the family of strategies that reuse knowledge across such domains [18,19]; federated learning extends the idea to collaborative adaptation when raw data cannot leave the owner [10,20]. Anomaly detection and transfer are therefore linked: a detection method whose transfer behavior is unknown has unknown deployment value.

Existing reviews cover parts of this picture. Reviews of leak management [4,21], GNNs for WDN operation [12,22], cyber-security of water systems [2,23], WWTP anomaly management [5], and AIoT or IoT architectures for water management [6–8] each address one axis. None of the reviews retrieved by the search protocol (Table 1) treats source–target transfer as the organizing question, codes studies by transfer strategy and deployment

layer, or examines TFMs alongside graph models (Section 2.2); a review with this focus may exist outside the sources searched.

Methodologically, the review follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 statement [24]. Searches were run in source-specific modular blocks, records were screened in title/abstract and full-text stages, and every work cited in the synthesis was coded by water domain, application category, dataset, model family, transfer category, deployment layer, privacy mechanism, digital-twin and explainable artificial intelligence (XAI) links, validation protocol, and evidence class (Section 3).

The review makes four contributions. First, a reproducible protocol: modular search blocks, source-specific fields, explicit evidence classes, and a traceability register that links every cited work to its coding (Supplementary Tables S1–S3). Second, a transfer taxonomy derived from the coded literature, with study counts per category and an explicit distinction between established practice and proposed strategy (Section 5.1). Third, an evidence-based separation of water-specific results from general method claims for graph-based models and TFMs, with a deployment-layer evidence matrix and a data-source table that show where the water evidence is concentrated (Sections 4.3 and 6.1).Fourth, a research agenda in which each gap is tied to a count from the coded evidence base, together with a gap map that locates demonstrated and proposed components on the AIoT layers (Sections 7.1 and 7.3).

The remainder of the article is organized as follows. Section 2 defines the technical building blocks and positions the review against related reviews. Section 3 describes the systematic review methodology. Section 4 reports the descriptive results of the evidence base. Section 5 presents the transfer taxonomy and the critical synthesis. Section 6 compares the model families. Section 7 sets out the research gaps, the future research agenda, and the gap map derived from them, and Section 8 concludes.

2. Background and Related Reviews

This section defines the technical building blocks used in the review. It explains how AIoT, GNNs, TFMs, transfer and federated learning, digital twins, and XAI relate to transfer-aware anomaly detection, and it positions the review against existing reviews. Figure 1 shows the four-layer deployment architecture used in the coding framework to classify the deployment layer of each study; the diagrams of the other building blocks are provided as Supplementary Figures S1–S4 rather than in the main text.

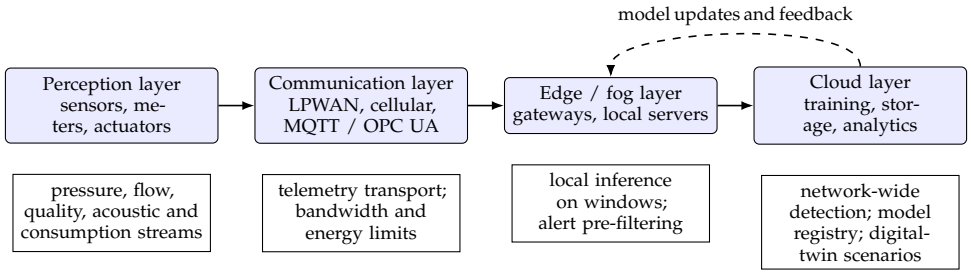


Figure 1. Four-layer AIoT deployment scheme used as the deployment-layer coding variable (perception, communication, edge/fog, cloud). The scheme names the functions of each layer only; which components have water-specific evidence and which remain proposals is shown in the gap map of Section 7.3.

2.1. Technical Building Blocks

AIoT deployment layers. An AIoT water system is usually described as four layers: a perception layer of sensors, meters, and actuators; a communication layer that relies on low-power wide-area protocols (e.g., LoRaWAN, NB-IoT) and industrial messaging (MQTT, OPC UA); an edge/fog layer of gateways or aggregation nodes that run lightweight

inference; and a cloud layer for model training, storage, and network-wide analytics [6–9]. For this review the layer matters because it fixes the constraints under which a detector must operate (latency, memory, energy, connectivity, and data locality) and therefore which transfer mechanisms are feasible.

Graph neural networks. GNNs learn node, edge, or graph representations by aggregating information over the graph neighborhood [11,25–27]. A WDN maps naturally to a graph with junctions, tanks, and sensors as nodes and pipes, pumps, and valves as edges, which is why GNNs have been applied to pressure estimation, leak detection and localization, consumption anomalies, and attack detection [13,17,28–30]. The relevant property for transfer is that a GNN trained on one graph may depend on that graph's topology and sensor placement; whether its representations generalize to unseen networks is an empirical question that only some studies test [13,17].

Tabular foundation models (TFMs). Among the emerging family of TFMs, TabPFN is a transformer-based model pre-trained on a large synthetic prior of tabular prediction tasks. At inference, it conditions on a small labeled table and predicts new samples without gradient-based training [14,15]. On general tabular benchmarks it has been reported to match or exceed tuned tree ensembles for small and medium datasets [14,15], the ensembles that were previously the reference for such data [31].

This general evidence motivates the consideration of TFMs for local tabular screening in water systems, where many sensor windows are small tables. The reviewed water literature, however, contains no water-specific TFM anomaly-detection study; one water-adjacent TabPFN study concerns irrigation water-quality index prediction rather than anomaly detection [32]; the remaining TabPFN evidence comes from industrial IoT intrusion detection [33,34] and from graph benchmarks [35]. TFMs are therefore a promising direction for water anomaly detection, but their capabilities have not yet been empirically demonstrated in this domain. Other tabular pre-training families exist (for example TabICL, CARTE, TabDPT, TP-BERTa, TabLLM, and XTab), but no water-related record using any of them was retrieved (Section 3.3); unless stated otherwise, TFM in this review therefore denotes the TabPFN family of prior-data fitted networks, and conclusions about TFMs are statements about that family.

Why TFMs and the AIoT layers remain axes of the review. TFMs and the AIoT deployment layers were fixed as axes before the search (blocks B4 and B5 of Table 2) because they are the components most often proposed for the local screening of sensor windows at gateways, and the review set out to test whether water evidence for them exists. It found none at a dedicated edge or fog layer and none for TFM anomaly detection in water (Sections 4.3 and 5.3.3). This absence is reported as a result, not assumed as a premise: every TFM and edge statement in the synthesis is labeled as general method evidence or as a proposal, and the title keeps both terms because the question the review answers is whether the evidence supports them. Time-series foundation models (for example Chronos [36] and TimesFM [37]) are a natural alternative candidate for local screening, since sensor streams are temporal; they were not a search axis, none of the 150 register works uses one, and the review does not assess them (Section 3.8). Their comparison with TabPFN-family screeners on water sensor windows is an agenda item (Section 7.2).

Transfer learning and federated learning. Following the standard definitions [18,19], a source domain provides data or a trained model, and a target domain is the system in which the detector must operate. Source and target may differ in feature space (sensor set), marginal distribution (demand regime, seasonality), task (pressure estimation versus leak detection), or label availability. Transfer strategies include parameter transfer and fine-tuning, domain adaptation, unsupervised or self-supervised pre-training, synthetic-to-real transfer, and pre-training across several graphs [13,16–19]. Federated learning is

treated here as a privacy-preserving form of collaborative transfer: clients (utilities, DMAs, or gateways) train locally and exchange model updates rather than raw data, subject to non-identically distributed data, communication limits, and security risks [10,20].

Digital twins and XAI. A digital twin is a calibrated simulation of the physical system that is updated with field data; for anomaly detection it can generate labeled scenarios, test detector hypotheses, and provide hydraulic context for alerts [38,39]. XAI methods such as SHAP and LIME, attention weights, and rule-based models make detector outputs interpretable for operators [40–42]; one WWTP prediction study (biogas-production forecasting, cited for method context only because it is not a detection task and lies outside the register) reports that transfer learning changed the stability of feature attributions [43]. In this review both are treated as supporting validation and trust layers rather than as independent topics.

2.2. Related Reviews and Positioning of the Present Review

Table 1 compares the reviews and surveys identified in the search with the present review along the dimensions that structure the synthesis. The comparison is limited to the reviews retrieved by the search protocol; it does not claim to be exhaustive. Two patterns are visible. First, water-specific reviews concentrate on one anomaly family (leaks, cyber-attacks, WWTP faults) or on one technology (GNNs, IoT architecture), and treat generalization across systems, where they mention it, as a future direction rather than as a coding dimension. Second, the reviews that cover transfer learning, federated learning, TFMs, or XAI in depth are general machine-learning surveys with no water content. The present review sits between these groups: it takes the water anomaly-detection literature as its corpus and codes it with the transfer, deployment, and trust variables that the method surveys define.

Table 1. Related reviews retrieved by the search protocol compared with the present review. Symbols: ✓ covered as a main theme; ○ mentioned; – not covered. TL: transfer learning; FL: federated learning; DT: digital twins; WDN/WDS: water distribution network/system; ICS/SCADA: industrial control systems/supervisory control and data acquisition; other abbreviations as in the text. The last column states what the present review adds relative to each review.

Review	Year	Water Domain	GNN	TL	TFM	AIoT	FL	DT	XAI	Main Limitation for the Present Question	Added Value of the Present Review
Rajan and Li [4]	2025	WDN leaks	○	–	–	○	–	–	–	Flow-data leak management only; no transfer or deployment coding	Cross-anomaly corpus; transfer and deployment-layer coding
Zuñiga-Uribe et al. [21]	2026	WDN leaks	○	–	–	–	–	–	–	Model/input/database mapping for leak detection; no generalization axis	Source–target validation counted per study
Vittori et al. [12]	2025	WDN operation	✓	○	–	–	–	○	○	GNN-centered; transfer to unseen networks noted but not coded	Transfer taxonomy with counts; comparison with tabular models
Liu and Liu [22]	2025	Infrastructure (incl. water)	✓	○	–	–	–	○	–	Reliability/resilience focus across infrastructures	Water anomaly detection focus; evidence classes
Tuptuk et al. [2]	2021	Water cyber-security	–	–	–	○	–	–	–	Security landscape; detection models not compared by transferability	Cyber-physical detection placed in the transfer framework
Habib et al. [23]	2026	WDN cyber-security	○	–	–	○	○	–	○	AI detection algorithms for attacks only	Attack detection integrated with hydraulic and quality anomalies
Yang et al. [5]	2025	WWTP	–	○	–	–	–	○	○	WWTP process anomalies only	WWTP evidence coded with WDN evidence under one scheme

Table 1. *Cont.*

Review	Year	Water Domain	GNN	TL	TFM	AIoT	FL	DT	XAI	Main Limitation for the Present Question	Added Value of the Present Review
Rahman et al. [6]	2026	Water quality/ resources	–	○	–	✓	○	○	○	AIoT applications survey; anomaly detection models not appraised	Detector evidence appraised by class and validation protocol
Amanatidis et al. [7]	2025	Water management	–	–	–	✓	○	○	–	Architecture-oriented; no systematic study coding	Systematic coding of models on the architecture layers
Pagano et al. [8]	2025	WDS (IoT)	–	–	–	✓	○	○	–	Communication and simulation tools for massive IoT	Model-side evidence for the same deployment layers
Liso et al. [1]	2024	Industry 4.0 (general)	○	○	–	○	–	–	–	Not water-specific	Water-specific corpus
Rana et al. [44] (preprint)	2025	ICS/SCADA (general)	–	–	–	○	–	–	–	Not water-specific; not peer reviewed	Peer-review status used as a coding variable
Biström [45] (thesis)	2025	Water cybersecurity	–	–	–	–	–	○	–	Bachelor's-level literature review comparing methods by F1 across studies; not peer reviewed	Evidence classes separate reviews and theses from peer-reviewed experiments
Pan and Yang; Weiss et al. [18,19]	2010/ 2016	None	–	✓	–	–	–	–	–	Method surveys without water content	Water instantiation of the transfer categories
Kairouz et al. [10]	2021	None	–	○	–	○	✓	–	–	Method survey without water content	FL treated as collaborative transfer for utilities
Wu et al. [11]	2021	None	✓	–	–	–	–	–	–	Method survey without water content	GNN transfer behavior on WDN graphs
Present review	2026	WDN, WWTP, water quality, testbeds	✓	✓	✓	✓	✓	✓	✓	Single-author coding; mixed retained corpus requiring evidence-class separation (Section 3.8)	Transfer-aware synthesis with counts, evidence classes, and a labeled framework

3. Systematic Review Methodology

The review was conducted and is reported in accordance with the PRISMA 2020 statement [24]; the completed PRISMA 2020 checklist and the PRISMA 2020 for Abstracts checklist are provided as Supplementary Material (Table S4). The review protocol was not registered before the searches were run. It was registered retrospectively, after data extraction and synthesis, in the Open Science Framework (OSF) Registries under the Generalized Systematic Review Registration template (registration DOI: 10.17605/OSF.IO/D7RXU; https://osf.io/d7rxu; registered on 6 October 2026). The registration documents the protocol as it was applied, and no amendments were made after registration. The protocol proceeds from research questions to search, selection, coding, synthesis, and results, and the coding framework is defined here, before any result is reported.

3.1. Research Questions

- **RQ1 (application evidence) .** Which anomaly-detection application areas in water infrastructures are addressed by graph-based and tabular models, with which datasets, validation protocols, and evidence classes?
- **RQ2 (transfer, the central question).** How can anomaly detectors transfer across heterogeneous water infrastructures while preserving accuracy, privacy, and operator trust? Which transfer strategies, including federated learning, have been evaluated, for which source–target mismatches, and with which validation protocols?
- **RQ3 (complementary roles).** Within an AIoT deployment, which roles do GNNs and TFMs play, and how do digital twins and XAI support the validation of transferred detectors and the trust of operators?

3.2. Search Strategy

Search dimensions. Two dimensions were primary search dimensions: water-infrastructure anomaly detection (the base block) and transfer-aware adaptation. The remaining dimensions (graph models, TFMs, federated learning, and the AIoT/edge, digital-twin, and XAI deployment context) were supporting technology blocks. A record could enter the corpus if it matched the base block and at least one technology block; it was not required to match all blocks at once. This modular design is what allows studies that address one or two dimensions (for example, a GNN leak detector without transfer learning, or a transfer study without a graph model) to be retrieved and coded. The synthesis categories used in Sections 4–7 (application category, evidence class, transfer category, deployment layer) were defined before coding (Section 3.6); the thematic labels used to describe the corpus emerged during coding.

Modular search blocks. Table 2 lists the modular search blocks. Searches were conducted in both English and Spanish, using the English terms shown in the table and their Spanish equivalents; the Spanish-language queries were run in Google Scholar, on the publisher sites, and in the repositories (Table 3). Each source of Table 3 was queried with the base block combined with one technology block at a time, that is, five block queries per source (B0 AND B1, B0 AND B2, . . . , B0 AND B5) adapted to the syntax of the source, and the union of the records retained after verification formed the identification set. General method references that define the technology blocks (surveys of transfer learning, federated learning, GNNs, XAI methods, and the TFM method papers) were added through backward citation chasing from included studies; they are labeled as general method evidence (class E3) in the register and are not counted as water-specific evidence.

Table 2. Modular search blocks. The base block (B0) was combined with each technology block in turn; a record was eligible if it matched B0 and at least one of B1–B5. Within each block, terms are ordered from the most general to the most specific. Spanish equivalents were used for the Spanish-language searches.

Block	Dimension	Terms (English; Spanish Equivalents in Supplementary Table S1)
B0	Water anomaly base block (primary)	(“anomaly detection” OR “fault detection” OR “leak detection” OR leakage OR “burst detection” OR “contamination detection” OR “attack detection” OR “intrusion detection”) AND (“water infrastructure” OR “water supply” OR “smart water” OR “water network” OR “water distribution” OR “wastewater treatment” OR “water quality” OR SWaT OR WADI OR BATADAL)
B1	Graph/GNN block	“topology-aware” OR “graph neural network” OR GNN OR “message passing” OR “graph convolutional” OR “graph attention”
B2	Transfer/domain adaptation block (primary)	“transfer learning” OR “domain adaptation” OR “domain generalization” OR “pre-training” OR “fine-tuning” OR “cross-network” OR “unseen network” OR “synthetic-to-real”
B3	Federated/privacy block	“privacy-preserving” OR “collaborative learning” OR “decentralized learning” OR “federated learning”
B4	Tabular foundation model (TFM) block	“tabular foundation model” OR (“in-context learning” AND tabular) OR “prior-data fitted network” OR TabPFN
B5	AIoT/edge/digital twin/XAI block	“Internet of Things” OR AIoT OR “edge computing” OR “fog computing” OR TinyML OR “digital twin” OR interpretable OR explainable OR XAI

Source-specific procedures. Table 3 reports, for each identification route, the field searched, the query form, the language, the stopping rule, and the date of the last search or verification. No publication-date limit was applied; the search boundary is defined by the last-search dates. No subscription bibliographic database (Scopus, Web of Science) was searched, because institutional access was not available during the search period; the retained corpus was therefore identified through the routes of Table 3 only, which PRISMA 2020 reports as identification via other methods (Figure 2).

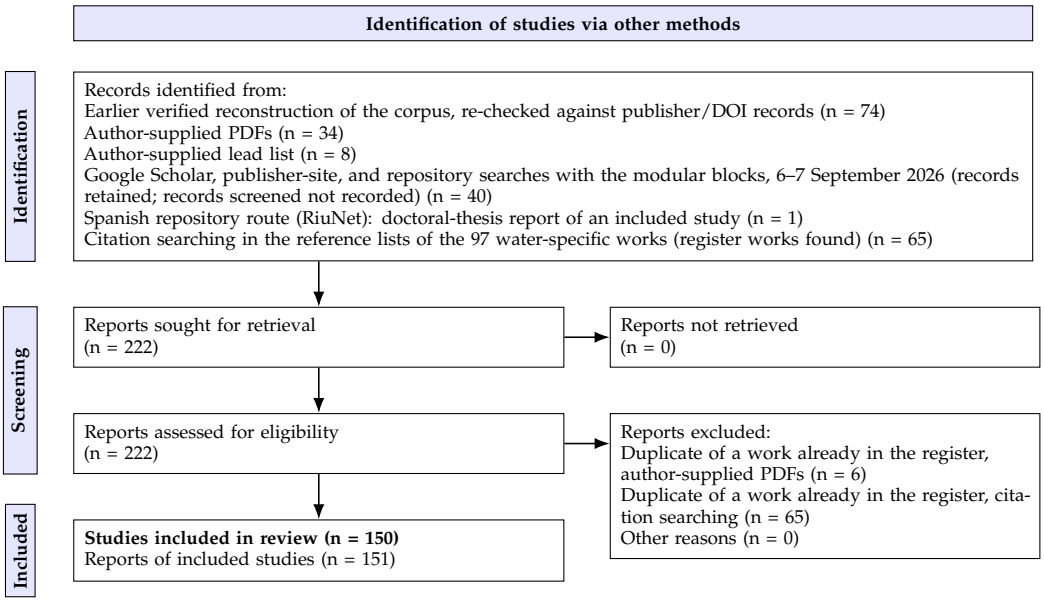


Figure 2. PRISMA 2020 flow diagram. No bibliographic-database search contributed records to the retained corpus, so all routes are reported as identification via other methods (Section 3.2); citation chasing was reconstructed from the Crossref reference lists of 93 of the 97 water-specific works. The only exclusions were duplicates of works already in the register. Evidence classes E1, E1p, E2, E3, and C1 are defined in Section 3.5.

Google Scholar was used with phrase-level block queries because it does not support long Boolean strings; for each block query the results were sorted by relevance and screened page by page until two consecutive pages yielded no new eligible record, with an upper bound of 200 results per query; this stopping rule is a pragmatic bound rather than a saturation criterion, and the Google Scholar, publisher-site, and repository route together contributed 40 of the 150 register works (Table 3; Figure 2). Publisher sites (for example, the MDPI platform) were searched with the block strings adapted to each site's syntax, and every candidate was followed to its DOI landing page. Spanish university repositories were searched directly with the Spanish base block in the title, abstract, and keyword fields. Works from an earlier verified reconstruction of the corpus (carried forward and re-checked against publisher or DOI records), author-supplied PDFs, and an author-supplied lead list were admitted only after the same verification against the primary publication record, and backward citation chasing was applied to the reference lists of the retained water-specific studies; forward citation chasing (works citing the retained studies) was not performed, because no citation index (Scopus, Web of Science) was available to the authors; this is recorded as a limitation (Section 3.8). Every record was retained only after title, author, year, and DOI (or stable URL) agreement with the primary record; the number of records screened before verification in the Google Scholar and publisher-site route was not recorded, which is stated as a limitation in Section 3.8.

Table 3. Identification routes and source-specific procedures (Supplementary Table S1 contains the search log and the verification outcome of each route).

Route	Field Searched	Query Form	Language; Stopping Rule	Last Search/Verification
Google Scholar	Full text (phrase blocks)	Phrase-level block pairs B0 AND B*i*, relevance order	EN, ES; two empty consecutive pages or 200 results	6–7 September 2026
Publisher sites and DOI landing pages (e.g., MDPI)	Title, abstract, keywords (site search)	B0 AND B*i* adapted to the site syntax; exact-title and DOI follow-up	EN, ES; all results screened	6–7 September 2026
RiuNet (Universitat Politècnica de València)	Title, abstract, keywords	Spanish B0 (and English B0)	ES, EN; all results screened	6–7 September 2026
UPCommons (Universitat Politècnica de Catalunya)	Title, abstract, keywords	Spanish B0 (and English B0)	ES, EN; all results screened	6–7 September 2026
RUO (Universidad de Oviedo)	Title, abstract, keywords	Spanish B0 (and English B0)	ES, EN; all results screened	6–7 September 2026
Earlier verified reconstruction of the corpus (carried forward)	Publisher or DOI record of each row	Re-check of title, authors, year, and DOI against the primary record	EN; 74 rows retained	6–7 September 2026
Author-supplied PDFs and lead list	Primary publication record	Exact-title and DOI verification of each file (34) and each lead-list entry (8)	EN, ES; 28 and 8 unique works retained	6–7 September 2026
Backward citation chasing	Reference lists of the retained water-specific studies	Manual screening; reconstructed from Crossref reference lists for Figure 2	EN, ES; 65 register works found, no new work	6–7 September 2026

3.3. Eligibility Criteria

Records were eligible when they met all of the following: (i) the study concerns anomaly, fault, leak, burst, contamination, or attack detection, or the hydraulic state estimation or leakage-risk assessment that supports such detection, in a WDN, WWTP, water-quality monitoring, smart-metering, river or environmental water sensing (hydrological telemetry), or water cyber-physical testbed setting; (ii) applies or evaluates at least one data-driven detector or estimator (classical machine learning, deep learning, graph-based, tabular foundation or federated model) in that setting, or is a general method or foundation reference for one of the five technology blocks of Table 2 (graph/GNN, transfer/domain adaptation, federated/privacy, TabPFN-family TFM, AIoT/edge/digital twin/XAI); the technology blocks are coding variables (Table 4), not inclusion conditions, so that water-specific studies with classical detectors are retained as the comparison base of the synthesis; (iii) the record is written in English or Spanish; and (iv) the record is a peer-reviewed journal article, a peer-reviewed conference paper, an examined thesis (bachelor's, master's, or doctoral) available in an institutional repository, or a public preprint. Reviews, surveys, assessment frameworks and conceptual architectures that meet criterion (i) but report no evaluated detector are admitted only as context works (class C1) for positioning and synthesis; they contribute no count to the results. Theses and preprints were admitted because they carry a substantial share of the Spanish-language and most recent evidence, but they were labeled by source type and were not counted as peer-reviewed water-specific evidence (Section 3.6). General method references were admitted only as class E3 evidence.

Table 4. Data extraction and coding variables (Supplementary Table S2 gives the value of each variable for every coded work; Supplementary Table S3 gives the full definitions).

Variable	Values/Definition
Review ID, citation, source title, source type, language	Identifier G001–G150; short citation; journal, proceedings, repository, or preprint server; journal/proceedings/thesis/preprint; EN/ES
Role in review	Primary water evidence; benchmark or testbed; conceptual architecture; review or survey; general method reference; reporting guideline
Water domain	WDN; WWTP; water quality; smart metering; water cyber-physical testbed; river or environmental water sensing; not water (general)
Application category	One of the eight predefined categories (primary category only)
Dataset	Simulated (hydraulic model); benchmark (BATADAL, SWaT, WADI, LeakDB); real operational or pilot data; not stated
Model family	GNN; classical ML/DL; TFM; digital-twin surrogate; federated method; none
Transfer category	T0 none; T1 conceptual mention only; T2 parameter transfer or fine-tuning; T3 domain adaptation or distribution alignment; T4 synthetic-to-real transfer with field validation on a system other than the one simulated (cross-network; same-network simulation-to-field validation is coded T0); T5 topology- or algorithm-informed pre-training across graphs; T6 federated or collaborative transfer; T7 pre-trained tabular prior with in-context inference; NA for reviews
Deployment layer	Perception/communication; edge; fog; cloud/offline; digital twin; testbed
Privacy mechanism	None; federated averaging; secure aggregation; differential privacy; not specified
Digital-twin link; XAI link	Yes/no/mentioned; type of explanation if any
Metrics reported; validation protocol	Metric names; in-distribution split; cross-network or unseen-network validation; cross-basin or cross-DMA transfer; field pilot; benchmark competition
Evidence class	E1, E1p, E2, E3, C1 (Table 5)
Quality appraisal	Dataset clarity; validation rigor; transfer protocol reported; deployment metrics reported; reproducibility (code or data); peer-review status (Supplementary Table S3)

Records were excluded when they concerned a different application domain that was retrieved through keyword overlap (for example, anomaly detection in energy systems, flood-prediction reviews, or educational-resource reviews), when they were news items, product pages, or other non-academic sources, when no method or evidence could be extracted, or when they duplicated another record. Block B4 included the generic terms "tabular foundation model" and "in-context learning" AND "tabular" as well as "prior-data fitted network" and "TabPFN", so other tabular pre-training families were not excluded by design; every retained water-related or water-adjacent B4 record nevertheless belongs to the TabPFN family (TabPFN, TabPFN-2.5, and variants), and the only non-TabPFN TFM in the register is TabICL, admitted as general method evidence (E3). Conclusions drawn about TFMs in this review therefore apply to the TabPFN family and are labeled as such (Section 3.8).

Table 5. Evidence classes used in the coding and in the synthesis, with the number of works of each class in the traceability register ($n = 150$).

Class	Definition	How It May Be Used in a Claim	*n*
E1	Peer-reviewed water-specific experimental study (journal or proceedings) with reported metrics	Water-specific evidence	75
E1p	Water-specific experimental study that is not journal peer reviewed (examined thesis or preprint)	Water-specific evidence of lower certainty; never the sole support of a claim	1
E2	Benchmark or testbed paper for water cyber-physical systems, or a general detector evaluated on such testbed data	Benchmark evidence; not equivalent to field deployment	21
E3	General method evidence: method papers, method surveys, and empirical studies outside the water domain	Method plausibility only; not water evidence	30
C1	Conceptual architecture, review, or survey without a new experiment	Positioning and proposal only	23

3.4. Study Selection and PRISMA Flow

Figure 2 shows the flow of records. Because no bibliographic-database search contributed records to the retained corpus (Section 3.2), all records are reported under identification via other methods. In total, 222 records were identified: 74 works from the earlier verified reconstruction of the corpus, 34 author-supplied PDFs, 8 works from the author-supplied lead list, 40 works retained from the Google Scholar, publisher-site, and repository searches of 6–7 September 2026 (the Spanish repositories contributed no additional unique work), one doctoral-thesis report retrieved from RiuNet, and 65 records found by backward citation chasing in the reference lists of the water-specific works.

All 222 reports were retrieved and assessed against the eligibility criteria of Section 3.3; 71 were excluded as duplicates of a work already in the register (6 author-supplied files that duplicated another file or an already represented work, and the 65 citation-chasing records, all of which were already in the register), no report was excluded for any other reason, and the thesis report was linked to the included study that it reports at thesis level (G112), so that 151 reports of 150 unique scholarly works were included. The number of records screened before verification in the Google Scholar and publisher-site route was not recorded (Section 3.8). Screening and full-text assessment were performed by the first author (B.B.); uncertain records were re-assessed after the coding framework had been fixed, and the decision recorded in Supplementary Table S1.

To quantify the risk of single-coder bias, a second coder (B. Al Khaled; see the Acknowledgments), who had not taken part in the screening or coding, re-assessed a random sample of 30 of the 150 register works (20%, stratified by evidence class: 14 E1, 1 E1p, 4 E2, 6 E3, 5 C1), blind to the register and using only the written definitions of Supplementary Table S3 and the works themselves: eligibility against the criteria of Section 3.3 and the four variables that drive every count of the synthesis (evidence class, application category, transfer category, deployment layer). Agreement was 25/30 for evidence class (Cohen's $\kappa = 0.76$), 24/30 for application category ($\kappa = 0.76$) and 23/30 for transfer category ($\kappa = 0.60$). For eligibility the second coder rated 11 works ineligible (19/30; κ is undefined because the register contains no ineligible work), in seven cases because criterion (ii) was read as requiring one of the five technology blocks and in four because of the scope of criteria (i) and (iv); the criteria are now worded as intended (Section 3.3) and no work was

excluded. For deployment layer, agreement was 14/30 ($\kappa = 0.04$); fourteen of the sixteen disagreements arose from two unstated conventions, the layer assigned to reviews and general method references and the layer assigned to detectors evaluated on testbed data, which are now written into Supplementary Table S3a. All 45 disagreements were resolved against the full text and are listed with their resolution in Supplementary Table S3d. Seven resolutions changed the register and two rules were then applied to the whole register for consistency: the transfer code of G020 is T0 rather than T5, because its multi-network model is tested on held-out samples of the training networks and not on an unseen network; the application category of G030 is water quality (drinking-water treatment works, not wastewater); the deployment layer of G130 is cloud/offline; the deployment layer of the 23 C1 rows is not applicable; and the transfer category is coded for the 76 water-specific experimental works only, so that eight E2 rows previously coded T0 are not applicable. Two bibliographic records (G055, G105) were corrected. The affected counts are updated in Sections 4.2, 4.3 and 5.1; no conclusion of the review changes. The check covers the eligibility and coding of retained works; it cannot cover the screening of excluded records, for which no log was retained (Section 3.8). Because the only exclusions at the verification stage were duplicates, a list of reports excluded with reasons (PRISMA 2020 item 16b) is not applicable; the citation-chasing matches are listed in Supplementary Table S1b.

3.5. Traceability of the Evidence Base

The retained corpus was reconstructed as an exact, row-level register of 150 unique scholarly works (Supplementary Table S1). Each work is assigned a stable identifier (G001–G150) and is linked to a DOI, stable publisher page, repository record, proceedings record, or other stable scholarly source; verification was completed on 6–7 September 2026. Deduplication was applied in order using the normalized DOI, the exact title, the normalized title, the author/year/venue tuple, and the preprint-to-version-of-record relationship, so that a DOI or title collision counts as one work.

The register is mixed by design: it combines empirical water studies, benchmark/testbed papers, general method/foundation works, and review or conceptual works. For this reason the manuscript does not treat the 150 records as a single empirical evidence set; instead, claims are qualified by evidence class. A small number of works are cited outside the register for positioning or method context only and are not counted in any figure of the review: the related reviews compared in Table 1; two cross-domain anomaly-detection studies and a WWTP biogas-prediction study discussed as method context in Section 5.3; the TabPFN water-quality-index prediction study of Section 5.3.3, which fails criterion (i) because it is a prediction rather than a detection task; the source papers of the LeakDB and BattLeDIM benchmarks; and the general method references of Section 2.1 that define the technology blocks. One included study has two reports: the journal article G112 and the doctoral thesis of its first author [46], retrieved through the RiuNet route; following PRISMA 2020, the two reports are linked to one register entry, so that the flow diagram counts 151 reports of 150 included studies. Of the 150 works, 76 are direct water-specific experimental studies (75 E1 and 1 E1p), 21 are water benchmark or testbed studies (E2), 30 are general method or foundation works (E3), and 23 are review, conceptual, or contextual works (C1). Quantitative statements about direct water experiments in Sections 4–7 refer to the 76 E1/E1p works, benchmark statements refer to the 21 E2 works, and E3 and C1 records support method definitions, architecture, and positioning without increasing the direct water-experiment count. Every work that supports a statement in the synthesis is listed in the register with its evidence class, search provenance, coding status, transfer category, deployment layer, and inclusion role, so that the evidence base is reconstructible record by record.

3.6. Data Extraction, Coding Variables, and Evidence Classes

Every register entry was coded with the variables in Table 4. The coding sheet is supplied in the source package (file `data/evidence_coding.csv`) and as Supplementary Table S2. Coding was based on the full text when the study was central to the synthesis and on the title, abstract, and methods summary otherwise; the coding basis is recorded per entry to support auditability. Application categories were fixed in advance as leak detection/localization, burst/hydraulic anomaly and state estimation, cyber-physical attack, water quality, sensor fault/data quality, consumption anomaly, WWTP process anomaly, and conceptual/framework; a study was assigned to its primary category. Transfer categories follow the taxonomy defined in Table 4 and elaborated in Section 5.1.

Evidence classes. Different kinds of evidence were separated rather than treated alike (Table 5). The classes are used throughout the synthesis to limit the strength of claims: a statement supported only by E3 evidence is a statement about the method in general, not about water systems; a statement supported only by C1 evidence is a proposal.

Quality appraisal. For each water-specific study (E1, E1p, E2) six items were appraised as reported / partly reported / not reported: dataset clarity (source, size, anomaly definition, prevalence), validation rigor (independent test data, repeated runs), transfer protocol (source and target defined, adaptation data size, target performance), deployment metrics (latency, memory, energy, or communication), reproducibility (code or data availability), and peer-review status. The appraisal is reported in Supplementary Table S3 and summarized in Section 4.4. It was used to qualify claims, not to exclude studies.

3.7. Synthesis Approach

The synthesis is descriptive. Counts are reported per application category, evidence class, transfer category, deployment layer, and gap indicator, and each count is traceable to the register entries listed in Supplementary Table S2. A pooled meta-analysis was not attempted because the reviewed studies report different tasks (classification, localization, state estimation), different datasets and anomaly definitions, different splits, and different metrics; pooling such values would produce a false sense of comparability [24]. Where reported performance values are mentioned, they are given only as reported-study examples with their protocol and are not treated as head-to-head comparisons. No statistical investigation of heterogeneity was performed; instead, every count is stratified by evidence class, application category, data source, deployment layer, and transfer category (Table 5; Sections 4 and 5.1), so that the composition behind each conclusion is visible.

No sensitivity analysis was performed; where benchmark evidence matters, counts are reported both for the 76 direct water-specific experimental works and for the 97 works that also include the 21 benchmark or testbed studies (Section 4.3), and claims that rest on a single study are identified as such. Risk of bias due to missing results was not formally assessed because no results were pooled; to keep non-journal evidence visible, examined theses and preprints were admitted and labeled by evidence class, and the absence of negative-transfer checks in the water literature is itself reported as a finding (Section 4.4).

Use of generative artificial intelligence. The tool (Claude, Anthropic; 2026 model versions) was used as a support tool during the revision for three purposes: proposing a data-source category for the 21 register entries previously recorded as “not stated”, each proposal being checked by the authors against the source text before the register was updated; checking that the counts reported in the text, tables and figures agree with the traceability register; and language editing, drafting of revised passages and LaTeX typesetting. It also states that the tool did not search for or screen records and made no eligibility, coding or quality-appraisal decision, that those decisions were made by the

authors and the second coder, and that all outputs were reviewed and edited by the authors, who take full responsibility for the content.

3.8. Limitations of the Review Process

Several limitations of the review process should be stated before the results. First, although the retained corpus is traceable at row level (Supplementary Table S1), it is a mixed scholarly corpus rather than a homogeneous empirical dataset; the main synthesis therefore separates peer-reviewed water experiments (E1), non-journal water experiments (E1p), benchmark or testbed evidence (E2), general method evidence (E3), and review or conceptual works (C1), so that general method and conceptual records are not used as direct water-domain evidence. Second, screening and coding were performed by a single author. The reliability check of Section 3.4 quantifies the resulting risk for the eligibility and coding of retained works (substantial agreement for evidence class and application category, moderate for transfer category) and led to the corrections listed in Supplementary Table S3d, but it cannot be extended to the records excluded at screening, for which no log was retained. Third, the register was assembled from publisher metadata and abstracts verified against the primary record, and the full text was consulted for coding where it was accessible to the authors; the text basis was not recorded for every row. For this revision the 21 water-specific experimental works whose data source had been recorded as not stated were re-examined against the full text (7 works) or the abstract and accessible excerpts (14 works), which established the data source of 18 of them (Section 4.3); no subscription access to the remaining paywalled full texts was available, so the coding of those works rests on the abstract and accessible excerpts, and the text basis of this re-examination and of the works coded in depth is recorded in Supplementary Tables S1 and S2. This is the main residual limitation of the coding.

Fourth, the reviewed studies use heterogeneous datasets, anomaly definitions, and metrics, so no pooled meta-analysis was attempted. Fifth, direct transfer validation, device-level deployment evidence, and operator evaluation of explanations are all limited in the water-specific literature, and the synthesis treats these as gaps rather than as settled results. Sixth, the Spanish-language searches were limited to Google Scholar, the publisher sites, and the three repositories (Section 3.2), so Spanish-language records indexed elsewhere may have been missed; and forward citation chasing was not performed (Section 3.2), so recent works that cite the retained studies but were not reached by the other routes may have been missed.

Seventh, reporting bias and the certainty of the evidence were not assessed with formal instruments; the evidence classes and the six-item quality appraisal serve that purpose and are coarser than GRADE-type ratings. Eighth, the corpus was identified without a subscription bibliographic database: Scopus and Web of Science were not accessible during the search period, the number of records screened before verification in the Google Scholar and publisher-site route was not recorded, and an exploratory search of OpenAlex and IEEE Xplore run on 3 September 2026 while the search blocks were being designed was not used to assemble the register. The review therefore does not claim exhaustive coverage of the indexed literature; its counts are those of a verified, traceable corpus (Supplementary Table S1), and the flow diagram reports the verification stage rather than an exhaustive retrieval. Ninth, the TFM category is in practice the TabPFN family: no water study using another tabular pre-training family (TabICL, CARTE, TabDPT, TP-BERTa, TabLLM, XTab) was retrieved, so the comparison of Section 6.1 and the TFM entries of the gap map are statements about prior-data fitted networks of the TabPFN type rather than about tabular foundation models as a class. Tenth, time-series foundation models were not a search

axis; none of the register works uses one, and the review therefore does not assess them (Section 2.1).

4. Descriptive Results of the Evidence Base

4.1. Composition of the Register and Coded Evidence Base

Table 6 summarizes the 150 works in the traceability register. By source type, 106 are journal articles, 38 are conference or workshop papers, 3 are book chapters, and 3 are a doctoral thesis, a preprint, and a technical review. According to the register, 147 records are classified as peer reviewed and 3 as not peer reviewed (an examined doctoral thesis, a preprint survey of class C1, and one water-specific experimental preprint, class E1p); examined theses and preprints are labeled by source type and evidence class and are not counted as peer-reviewed journal evidence. All 150 works are in English (for 4 works the language was inferred from the full text because the source metadata gives no language field); no Spanish-language work was retained, although the Spanish repositories were searched (Section 3.2), and the one Spanish-language doctoral thesis retrieved is linked to the register as a second report of an included study (Section 3.5). The register was assembled through modular English searches, targeted follow-up searches, and Spanish-language discovery, but after work-by-work bibliographic verification and deduplication the retained records are English-language scholarly works. 80 works were published in 2021–2024 and 23 in 2025–2026, while 47 are from 2020 or earlier, which reflects both the recency of the graph-based water literature and the older method and testbed foundations.

By evidence class, 75 works are peer-reviewed water-specific experiments (E1), 1 is a water-specific experiment without journal peer review (E1p), 21 are benchmark or testbed papers (E2), 30 are general method or foundation works (E3), and 23 are reviews, conceptual, or contextual works (C1). The direct water-specific experimental core (E1 and E1p) therefore comprises 76 works, and every quantitative statement about direct water evidence below refers to this set; benchmark statements refer to the 21 E2 works, and the 30 E3 and 23 C1 records support method definitions and positioning only.

Table 6. Composition of the traceability register ($n = 150$ unique retained scholarly works; full row-level classification in Supplementary Table S1).

Dimension	Categories	*n*
Source type	Journal article	106
	Conference or workshop paper	38
	Book chapter	3
	Thesis, preprint, or technical review	3
Peer-review status	Peer reviewed/not peer reviewed	147 / 3
Language (verified register)	English (4 inferred from the full text; no language field in the source metadata)	150
Publication year	≤2020/2021–2024/2025–2026	47/80/23
Evidence class	E1/E1p/E2/E3/C1	75/1/21/30/23
Role in review (by class)	Direct water-specific experiment (E1, E1p)	76
	Water benchmark or testbed (E2)	21
	General method or foundation (E3)	30
	Review, conceptual, or contextual (C1)	23
Water domain	WDN	70
	Water testbed (SWaT, WADI, BATADAL, LeakDB, analogous)	21

Table 6. *Cont.*

Dimension	Categories	*n*
	WWTP	5
	Water quality	7
	Other water (river/environmental, telemetry, security, treatment works, smart metering)	8
	Non-water / general method	39
Model family	Graph-based (GNN, topology-aware)	32
	Classical or temporal ML/DL	74
	Tabular foundation model (TFM)	4
	Digital-twin surrogate	4
	Federated method	3
	XAI method	6
	Review or survey	21
	Benchmark or testbed description	6

4.2. Application Categories (RQ1)

Table 7 reports the distribution of the 76 direct water-specific experimental works (classes E1 and E1p) across the predefined application categories, together with the model families and data sources recorded for each category in the traceability register (Supplementary Table S1). The 21 benchmark or testbed works (E2), the 30 general method or foundation works (E3), and the 23 review, conceptual, or contextual works (C1) are not assigned to these categories. Two patterns are visible. Hydraulic tasks dominate: leak detection and localization (43 works) and burst, hydraulic anomaly, or state estimation (9 works) together account for 52 of the 76 works. The remaining works are distributed across cyber-physical attack detection (8), water quality (9), WWTP process anomaly (4), and sensor fault or data quality (3), so the non-hydraulic categories remain comparatively sparse.

Table 7. Application categories of the 76 direct water-specific experimental works (E1, E1p). Model-family counts sum to *n* in each category; the data-source column lists the most frequent stated sources (up to three); example works are given by register identifier (Supplementary Table S1).

Category	*n*	Model Families (Register)	Data Sources (Register)	Example Works (Supplementary Table S1)
Leak detection/localization	43	classical ML/DL (19), GNN (11), digital-twin surrogate (3), LSTM/temporal DL (2), graph-based/topology-aware NN (2), XAI method (2), discrete-time dynamic graph network (1), probabilistic digital twin/generative deep learning (1), hybrid data-driven leak detector (1), IoT-enabled machine learning (1)	simulated WDN, EPANET simulation, real utility data	G002, G003, G004, G005, G006
Burst/hydraulic anomaly & state estimation	9	GNN (6), LSTM/temporal DL (2), classical ML/DL (1)	EPANET simulation, simulated WDN, real DMA data	G001, G014, G015, G019, G020
Cyber-physical attack	8	classical ML/DL (4), GNN (3), LSTM/temporal DL (1)	EPANET simulation, simulated WDN, BATADAL	G034, G035, G036, G037, G038

Table 7. *Cont.*

Category	*n*	Model Families (Register)	Data Sources (Register)	Example Works (Supplementary Table S1)
Water quality	9	GNN (5), classical ML/DL (2), XAI method (1), LSTM / temporal DL (1)	real water-quality sensor data, EPANET simulation, real utility data	G024, G025, G026, G027, G030, G111
WWTP process anomaly	4	classical ML/DL (3), LSTM/temporal DL (1)	real WWTP operational data	G031, G120, G122, G144
Sensor fault/data quality	3	classical ML/DL (1), LSTM/temporal DL (1), federated method (1)	real water-quality sensor data, hydrological telemetry	G028, G029, G032
Total	76	Direct water-specific experimental works (E1 + E1p)		

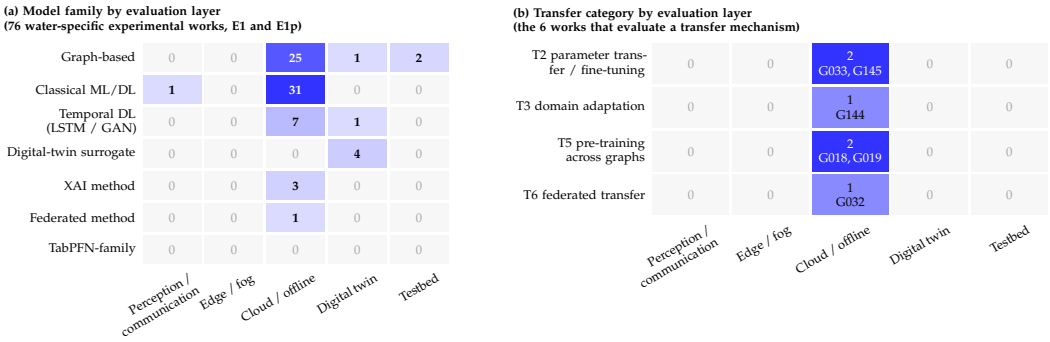
4.3. Datasets, Deployment Layers, and Validation Protocols

Across the 76 direct water-specific experimental works, 39 rely on hydraulic simulation (EPANET-type models or digital-twin scenarios), 30 use field, pilot, laboratory-rig, or real operational data, 4 use a public water benchmark or testbed (LeakDB, BATADAL, SWaT/WADI, or an analogous testbed), and 3 do not state the data source in the accessible text (Table 8); the 21 works whose data source had been recorded as not stated were re-examined for this revision against the full text or the accessible excerpts, which established the source of 18 of them. The 21 E2 benchmark or testbed works add the cyber-physical evidence, which is dominated by the SWaT and WADI testbeds and the BATADAL competition network; taken together, the 97 E1/E1p/E2 water empirical and benchmark works therefore rest heavily on simulation and on a small number of shared testbeds. Two public leak benchmarks deserve mention because they are the datasets best suited to source–target evaluation of leak detectors: LeakDB [47], a set of simulated leak scenarios on benchmark networks with labeled events, used by one register work (G124), and BattLeDIM [48], whose L-Town network provides a realistic topology with two years of SCADA-type data, one of them with labeled leaks, used by one register work (G113); neither has yet been used in the register for a cross-network protocol, although L-Town contains distinct areas that could serve as source and target. By deployment layer, 67 of the 76 direct water-specific works are cloud or offline evaluations, 6 combine cloud training with a digital twin, 2 are testbed evaluations, and 1 is a perception-and-communication pilot; no work in the register evaluates a direct water-specific detector at a dedicated edge or fog layer.

Table 8. Data sources recorded in the traceability register, for the 76 direct water-specific experimental works (E1, E1p) and for the 97 works that also include the 21 water benchmark or testbed works (E2). The last column states the class that contributes the benchmark and testbed evidence.

Data Source Category	*n* (E1, E1p; /76)	*n* (E1, E1p, E2; /97)	Note
Hydraulic simulation (EPANET-type models, digital-twin scenarios)	39	39	Simulated benchmark or utility networks; no field data by construction (two works add a real-network case)
Field, pilot, laboratory-rig, or real operational data	30	32	Real DMA, utility, treatment-works, water-quality, telemetry, acoustic, GIS/repair, or laboratory-rig and prototype data
Public water benchmark or testbed (LeakDB, BATADAL, SWaT, WADI, analogous)	4	22	The 21 E2 works supply most of the SWaT/WADI/BATADAL evidence
Not stated in the accessible text	3	4	Data source not identifiable in the abstract and accessible excerpts (full text paywalled); 21 works were re-examined for this revision (Section 3.8)
Total	76	97	

Figure 3 cross-tabulates, for the 76 direct water-specific experimental works, the layer at which each detector is evaluated with its model family and, for the 6 works that evaluate a transfer mechanism, the transfer category. Three properties of this matrix constrain the gap map of Section 7.3. First, the evidence is concentrated in one cell: 67 of the 76 works are cloud or offline evaluations, 6 combine cloud training with a digital twin, 2 are testbed evaluations, and 1 is a perception-and-communication pilot. Second, the layers in which a transfer-aware deployment would place its time-critical functions are empty: no direct water-specific work in the register is evaluated at a dedicated edge or fog layer, and no TFM has water evidence at any layer. Third, all six transfer-evaluating works are themselves cloud or offline evaluations, so transfer and deployment have not yet been demonstrated together in a water setting.



Row totals (a): graph-based 28, classical ML/DL 32, temporal DL 8, digital-twin surrogate 4, XAI 3, federated 1, TabPFN-family 0. Column totals (a): 1 / 0 / 67 / 6 / 2. Shading is proportional to the count within each panel.

Figure 3. Deployment-layer evidence matrix from the traceability register (Supplementary Table S1): (**a**) model family by evaluation layer for the 76 direct water-specific experimental works (E1, E1p); (**b**) transfer category by evaluation layer for the 6 works that evaluate an explicit transfer mechanism (T2, T3, T5, T6), with their register identifiers. No work is evaluated at a dedicated edge or fog layer, and all six transfer-evaluating works are cloud or offline evaluations.

4.4. Recurring Methodological Weaknesses

Table 9 reports how many of the water-specific experimental works (E1, E1p) and of the 97 works that also include the benchmark or testbed works (E2) meet each item of the six-item appraisal of Section 3.6. Two items, device-level deployment metrics and code or data availability, were not coded at row level and are reported through the deployment-layer coding and as a limitation (Section 3.8).

The appraisal yields four recurring weaknesses that constrain the synthesis. First, generalization is rarely tested: only 6 of the 76 direct water-specific experimental works evaluate an explicit transfer mechanism, 35 are validated on simulated data only, and none reports a negative-transfer check. Second, deployment feasibility is not measured: device-level latency, memory, energy, or communication metrics were not coded as a variable because 75 of the 76 works are evaluated offline, in a digital twin or on a testbed, and none reports device-level measurements for a TFM. Third, the cyber-physical evidence is concentrated on a few shared testbeds (SWaT, WADI, and the BATADAL competition network), so attack-detection results describe those testbeds rather than operational utilities. Fourth, the data source of 3 works could not be established from the accessible text, and the evidence class of every work is stated in Supplementary Table S1 so that the classification can be audited.

Table 9. Quality appraisal of the water-specific experimental works (E1, E1p; $n = 76$) and of all water empirical and benchmark works (E1, E1p, E2; $n = 97$): number of works meeting each item (Supplementary Table S3c).

Appraisal Item	*n* (/76)	*n* (/97)	Note
Data source stated	73	93	3 E1 works and 1 E2 work give no identifiable data source in the accessible text
Validation protocol: cross-system transfer evaluated	6	6	T2, T3, T5, T6 works (Table 10)
Validation protocol: same-network simulation-to-field validation	1	1	G017 (coded T0)
Validation protocol: benchmark or testbed split	4	25	LeakDB, BATADAL, SWaT/WADI, SWaDNet; all 21 E2 works
Validation protocol: in-distribution split on field or operational data	27	27	Single system, no transfer claim
Validation protocol: in-distribution split on simulated data only	35	35	Single simulated network or scenario set
Validation protocol: not stated	3	3	Same works as the first row
Transfer protocol reported (source, target, adaptation data)	6	6	All six works that claim transfer; two further works report same-network validation or maintenance (G017, G117), not transfer
Device-level deployment metrics	not coded	not coded	75 of 76 works are evaluated at the cloud/offline, digital-twin or testbed layer (Figure 3); the single perception-layer work is a smart-metering pilot; no work reports device-level metrics for a TFM
Code or data availability	not coded	not coded	Not recorded at row level (Section 3.8)
Peer reviewed	75	96	One preprint (E1p) among the experimental works; all E2 works peer reviewed

5. Taxonomy and Critical Synthesis

5.1. Transfer Taxonomy Derived from the Coded Literature (RQ2)

The taxonomy in Table 10 was built in two steps. The categories were first defined from the transfer-learning literature [18,19] and the federated-learning literature [10,20], then each direct water-specific experimental work was assigned the category that best describes what is reused between a source and a target. The counts are therefore not counts of papers that mention transfer learning; they count works in which a reusable object is moved from a source to a target and the target result is reported. Two properties of the distribution matter for the rest of the review. First, 69 of the 76 direct water-specific experimental works do not reuse anything across systems: they train and test within one network, testbed, or dataset. Second, transfer, when present, is thinly spread across categories, and no work reports cross-network synthetic-to-real transfer (T4; the one simulation-to-field validation [28] concerns the same networks and is coded T0, Table 11) or a pre-trained tabular prior applied to a water task (T7).

Table 10. Transfer taxonomy derived from the traceability register. Counts refer to the 76 direct water-specific experimental works (E1, E1p); transfer-evaluating works are identified by register ID (Supplementary Table S1). The supporting-evidence column lists non-water method evidence (E3) and proposals (C1) that inform a category without constituting water evidence.

Category	What is Reused (Transferable Object)	*n* Water	Water Studies	Status	Supporting Evidence and Typical Mismatch
T0 No transfer	Nothing; single-system training and testing	70	The remaining 70 E1/E1p works (Supplementary Table S1)	Default practice	Results describe one network, plant, or testbed
T1 Conceptual only	Transfer discussed, not evaluated	0	–	Proposal	Reviews and conceptual frameworks (C1)
T2 Parameter transfer/fine-tuning	Network weights re-trained on a small target set	2	G033, G145	Demonstrated in two water studies	Fine-tuning across a new site or regime; mismatch: same task, new site with few labels
T3 Domain adaptation/distribution alignment	Feature alignment, re-weighting, invariant embeddings	1	G144	Demonstrated in one WWTP study	Method surveys [18,19]; mismatch: demand regime, calibration, seasonality
T4 Synthetic-to-real transfer with field validation	Simulator-trained detector validated on field data of a system other than the one simulated; same-network simulation-to-field validation [28] is coded T0	0	–	Not demonstrated in water	Digital-twin scenario generation; mismatch: simulator bias, unmodeled demand and control
T5 Topology- or algorithm-informed pre-training across graphs	Graph encoder pre-trained across networks or on algorithmic reasoning, evaluated on unseen graphs	2	G018, G019	Demonstrated on simulated networks	GNN inductive learning [11,26]; mismatch: topology, sensor coverage, hydraulic regime
T6 Federated/collaborative transfer	Model updates exchanged between utilities, DMAs, or gateways	1	G032	Demonstrated in one water study (federated GAN for hydrological telemetry)	FedAvg and FL survey [10,20]; mismatch: non-IID clients, communication, poisoning
T7 Pre-trained tabular prior (in-context inference)	Synthetic-task prior conditioned on a small local table	0	–	Not demonstrated in water	TabPFN model papers [14,15]; general (E3) evidence only

5.2. Transfer Objects and Transfer Strategies for Graph-Based and Temporal Detectors

The taxonomy classifies studies by what is reused. The coded studies and the method literature show that the reused object can be one of several things: the weights of a trained network [16] (and, outside the register, in a WWTP biogas-prediction study [43]); a graph encoder or message-passing scheme that represents relations among junctions, pipes, tanks, pumps, and sensors [11,13,17]; a pre-trained tabular prior queried in context [14,32]; a simulator-generated scenario library or residual rule [28,38,46]; or, in the federated literature, aggregated model updates [10,20]. Which object is worth transferring depends on the target constraint rather than on the model family [18,19].

A target with a new sensor layout but ample normal-operation data favors transferring a graph encoder and adapting the output layer; the closest tested configuration is the multi-graph pre-training study, which reuses a pressure estimator on unseen simulated networks [13]. A target with tabular streams and no usable topology favors a tabular prior or a calibrated tree-based baseline [14,31], a configuration for which no water anomaly evidence exists. A target that cannot share raw data favors federated updates, whose non-IID and communication costs are documented only outside water [10]. Explanation patterns are a further transferable object: the one WWTP study that tracks them, a biogas-

production prediction study cited for method context outside the register, shows that feature attributions can change after transfer [43].

For graph-based detectors, transfer is the central unresolved question, because each utility has a different graph, sensor set, and hydraulic regime [12,13]. Table 11 lists the strategies that the coded studies evaluate or that the method literature suggests, with the source–target mismatch each addresses and the verification each requires. Two strategies have water-specific experimental support, both on simulated targets: multi-graph pre-training with randomized sensor placement [13] and algorithm-informed pre-training that embeds flow-algorithm reasoning before the leak task [17]. One strategy has field support for a temporal rather than a graph model, namely fine-tuning across DMAs [16], and one study validates a simulation-trained graph detector on real events without moving to another network [28]. The remaining strategies (domain adaptation, compression of a cloud graph model for fog placement, federated aggregation) are proposals supported by general method evidence. A graph-transfer study therefore has five elements to report: the source graph and its sensor placement, the pre-training task, the transferred object, the adaptation procedure and its data budget, and the target evaluation under the target topology and sensor coverage [13,17,18].

Table 11. Transfer strategies for graph-based and temporal water anomaly detectors, with the source–target mismatch each addresses, the verification each requires, and the evidence available in the coded set.

Strategy	Source	Target	Mismatch Addressed	Verification Required	Evidence in the Coded Set
Multi-graph pre-training with randomized sensor placement	Several simulated WDN graphs	Unseen simulated WDN with sparse pressure sensors	Topology; sensor coverage	Evaluation on unseen graphs and altered sensor layouts	Demonstrated in simulation [13] [E1]
Algorithm-informed pre-training	Synthetic graphs with known flow-algorithm behavior	Leak localization under out-of-distribution hydraulic conditions	Shortcut learning; hydraulic regime	Out-of-distribution test conditions	Demonstrated in simulation [17] [E1]
Simulation-to-field validation of a graph detector	EPANET-simulated leaks for two real networks	Real leak events in the same networks	Simulator bias	Reported simulation-to-field drop; transfer to other networks	Same-network simulation-to-field validation only [28] (E1), not counted as T4 because source and target are the same networks; drop not reported
Fine-tuning of a temporal detector across systems	DMA with labeled bursts	New DMA with few labels	Site change; label scarcity	Adaptation data size; negative-transfer check	Demonstrated on field data with an LSTM [16] [E1]; not tested for graph models
Retraining with explanation-stability tracking	Plant model under one operating regime	Changed regime or plant	Regime shift; explanation drift	Attribution stability before and after transfer	Demonstrated for a WWTP model [43] [E1]
Domain adaptation of graph or tabular features	Source network with normal-operation data	Target with different demand regime or calibration	Distribution shift without labels	Alignment quality; target false-alarm rate	Proposed; method evidence only [18,19] [E3]
Compression or distillation of a cloud graph model for fog placement	Cloud-trained graph model	Fog gateway with limited compute	Resource constraints	Latency and memory on the target hardware; detection loss after compression	Proposed; no water study; constraints documented for edge intrusion detection [34] and federated settings [10] [E3]

Table 11. *Cont.*

Strategy	Source	Target	Mismatch Addressed	Verification Required	Evidence in the Coded Set
Federated aggregation of detectors across utilities	Several utilities or DMAs with local data	Shared model without raw-data pooling	Data-sharing restrictions; non-IID clients	Non-IID robustness, communication rounds, poisoning resilience	One water instance (federated GAN for hydrological telemetry, G032; [E1]); otherwise general method evidence [10,20] [E3]

5.3. Critical Synthesis by Evidence Pattern

The following subsections synthesize the coded evidence by pattern rather than by study. For each pattern the text states what dominates, which assumptions the studies share, where the studies contradict one another or their own deployment claims, and which weaknesses recur.

5.3.1. Topology-Dependent Hydraulic Models

Graph-based models dominate the hydraulic evidence: graph neural networks and topology-aware models account for 28 of the 76 direct water-specific experimental works and are concentrated in the leak-detection and hydraulic-state categories (Table 7). The shared assumption is that the pipe network is known, static, and fully represented in the graph. This assumption is reasonable for the simulated benchmarks that are the data source of most of these works, but it is weaker in operational networks, where sensor coverage is sparse and changes over time. The register documents a small set of works that address part of this problem by pre-training across graphs or algorithmic structure and evaluating on networks not seen during training, including an algorithm-informed graph model [17] and a transferable-metamodel study (G019 in Supplementary Table S1); multi-graph pre-training with randomized sensor placement has also been reported for pressure estimation [13]. Even so, this evidence is obtained on simulated targets, and no graph study in the register reports validated transfer to a different field network, which is why cross-graph transfer is treated as demonstrated in simulation but not yet in the field.

5.3.2. Benchmark-Based Cyber-Physical Detection

Attack detection is evaluated almost exclusively on three sources: the SWaT treatment testbed, the WADI distribution testbed, and the BATADAL simulated network [3,49,50]. Detectors evaluated on these data in the coded set include a graph-structure-learning model [30], a spatiotemporal model [51], and a GNN-based sensor-placement method coupled with an LSTM detector on BATADAL that is updated monthly by fine-tuning with weight interpolation [52] (model maintenance on the same network, coded T0: the T2 category requires parameters transferred from a different source system or regime, Table 10); all report benchmark splits and none reports transfer to another plant or network. The security reviews in the register list many more benchmark studies and consistently identify the same limitation: high benchmark scores do not establish operational performance, because the testbeds are small, attacks are scripted, and normal operation differs from a utility with thousands of assets [2,23,44,45].

Two consequences follow for the transfer question. The source domain of almost all attack-detection evidence is a testbed, so any deployment is a synthetic-to-real or testbed-to-field transfer that the coded studies do not evaluate; and the graph in cyber-physical detection is usually a learned dependency graph between sensors [30] rather than the pipe topology, so results on these benchmarks do not transfer directly to the hydraulic GNN pattern. The same dependency-graph approach has been developed for anomaly detection

in computer-network telemetry [53,54]; these works are coded as general method evidence (E3) because, according to their abstracts, they were not evaluated on water data, and they should not be read as water-specific support for graph attention models.

5.3.3. Local Tabular Screening and TabPFN-Family TFMs

The general evidence for TFMs is stronger than the water evidence, and the two must be kept apart. On general benchmarks, TabPFN is reported to perform well on small tabular datasets without task-specific training [14,15], in a regime where tuned tree ensembles were previously the reference [31]. Outside water, TabPFN and its variants have been applied to intrusion detection in industrial IoT, with reported limits on sample and feature counts that motivated an ensemble variant for multi-access edge computing [33,34], and to graph benchmarks where tabularized node features let TabPFN compete with GNNs in some settings [35]. In the traceability register, however, all four TFM works (G066–G069) are classified as general method evidence (E3), and the register contains no water-specific TFM anomaly-detection study. The only water-adjacent TFM application identified, a TabPFN study of irrigation water-quality index prediction [32], addresses regression rather than anomaly detection, so it fails criterion (i) of Section 3.3 (it neither detects anomalies nor estimates the hydraulic state that supports detection); it is not part of the register and is cited as context, not as water-specific evidence. The role of TFMs as local anomaly screeners in water systems is therefore a proposal that follows from the tabular structure of node-level data, not from water experiments, and their feasibility on embedded or gateway devices in water is undemonstrated.

5.3.4. Transfer-Aware and Federated Adaptation

Among the 76 direct water-specific experimental works, only 6 evaluate an explicit transfer mechanism (Table 10). They show that fine-tuning can reduce the labeled data needed for detection at a new site or under a changed regime (G033, G145), that pre-training across graphs or algorithmic structure can produce detectors that work on unseen simulated networks (G018, G019), and that domain adaptation can be applied to a wastewater-treatment fault-diagnosis task (G144). Federated learning appears in the water-specific experimental evidence: one study implements a federated generative model for anomaly detection and repair of hydrological telemetry data (G032, class E1, transfer category T6). Because the conclusions on federated learning and on domain adaptation each rest on a single study, their eligibility is stated explicitly: G032 meets criterion (i) as anomaly detection in river and environmental water sensing (hydrological telemetry), a water setting named in Section 3.3 but not a distribution network, and G144 as fault diagnosis of a wastewater treatment process with transfer between operating periods of the same plant; neither concerns a water distribution network, and the synthesis does not generalize from them to WDNs. This evidence is nonetheless limited and emerging; it is a single water study set against a much larger body of general federated-learning method work [10,20], so privacy-preserving collaborative transfer in water systems is best described as sparsely evidenced rather than demonstrated at scale, and no water study reports domain adaptation with distribution alignment beyond the single fault-diagnosis case.

The premise that labeled anomalies are scarce is visible in the corpus itself: of the 76 water-specific experimental works, 39 train on simulated anomalies, 30 use field, pilot, laboratory-rig or real operational data, 4 use benchmark or testbed data, and 3 do not state their data source (Table 8), so detectors are far more often learned from simulated events or from unlabeled normal operation than from labeled field anomalies. Two cross-domain works, cited here for method context only and not counted in the register (Section 3.5), indicate the mechanisms that the water literature has not yet tested. A federated anomaly-

detection study on heart sounds shows that collaborative detection can be trained across clients with weak or no labels [55], which is the regime that a multi-utility extension of the single water study (G032) would face under T6. A few-shot anomaly-detection baseline built on an immunity-inspired prototype network [56] shows that a detector can be adapted from a handful of labeled anomalies, which is the data-efficiency claim that the two water fine-tuning studies (T2) make with far fewer comparisons. Neither work involves water data; they are cited only to locate the water evidence relative to current few-shot and weak-label methods.

5.3.5. Validation and Trust Layers: Digital Twins and XAI

Digital twins appear in two experimental studies, a peer-reviewed study in which a hydraulic model generates the scenarios on which a detector is trained and tested [38] and in one conceptual framework [39]; the doctoral thesis in which a graph-convolutional twin estimates pump speeds and supports state estimation on benchmark networks [46] reports the work of G112 at thesis level and is linked to that register entry rather than counted separately (Section 3.5); reviews mention them as a direction [6,12]. The evidence supports their use as scenario generators and hydraulic plausibility checks for alerts, but no coded study closes the loop from twin-generated training to field validation, which is why twin-to-field transfer is coded as proposed. XAI appears in three experimental studies: fuzzy rule explanations for leak detection [42], learned dependency graphs and attention scores for attack detection [30], and feature-attribution stability under transfer for a WWTP model [43]. The third result is the most relevant to the transfer question: it indicates that adapting a model to a new domain can change which features the explanation highlights, so explanations validated on the source cannot be assumed to hold on the target. No coded study evaluates explanations with operators, so operator trust remains an assumption in the reviewed literature rather than a measured outcome.

6. Comparative Analysis of the Model Families

6.1. Qualitative Comparison of Graph-Based Models and TabPFN-Family TFMs (RQ3)

Table 12 compares the two model families along the dimensions that the synthesis identified as decisive, and states for each cell the class of evidence that supports it. A numerical scoring of these dimensions was deliberately avoided: the dimensions are not measured on a common scale in the reviewed studies, and any score would be an author judgment presented with artificial precision. The comparison supports a division of roles rather than a ranking. Graph-based models have water-specific evidence for topology-dependent tasks and are the only family with evidence of generalization to unseen networks, although that evidence is simulated. TFMs have only general method evidence (E3) for fast inference on small tables; the one water-adjacent TFM study, a water-quality index regression [32], lies outside the register and does not address anomaly detection. Their role as local screeners is therefore a proposal that follows from the tabular structure of node-level data, not from water experiments.

Table 12. Qualitative comparison of graph-based models and tabular foundation models. The evidence class of each cell is given in brackets; "proposed" marks cells with no water-specific experimental support in the coded set.

Dimension	Graph-Based Models (GNNs, Topology-Aware Networks)	Tabular Foundation Models (TabPFN Family)
Data structure exploited	Pipe topology or learned sensor-dependency graph plus node time series [E1, E2]	Small multivariate tables without topology [E3]

Table 12. *Cont.*

Dimension	Graph-Based Models (GNNs, Topology-Aware Networks)	Tabular Foundation Models (TabPFN Family)
Water tasks with experimental evidence	Leak detection and localization, pressure estimation, consumption anomalies, contamination detection, and attack detection on testbeds [E1, E1p, E2]	One water-adjacent study, on water-quality index prediction [32]; no water anomaly-detection study in the register [proposed]
Generalization to unseen systems	Two studies on simulated networks [13,17] [E1]; one simulation-trained detector validated on real events in its own networks [28] [E1]; none across field networks	None in water; in-context inference is itself a form of transfer of a synthetic prior [E3]
Deployment-layer evidence	Cloud or offline for 67 of the 76 water works, with 6 combining cloud and a digital twin and 2 testbed evaluations [E1, E1p]; no dedicated edge or fog evaluation	Edge/MEC intrusion detection outside water [33,34] [E3]; no device-level water evidence [proposed]
Data requirements	Topology, sensor placement, and long normal-operation records; simulation often substitutes for field data	A small labeled context table; documented limits on rows and features [14,34] [E3]
Interpretability evidence	Attention or learned graphs [30] and fuzzy rules [42] [E2, E1p]	None reported in water; general post-hoc methods applicable [40,41] [E3]
Privacy and federation	One water federated study (federated GAN for hydrological telemetry, G032) [E1]; evidence limited and emerging; model updates are large	No federated water study [proposed]; the prior is shared while the context stays local by construction
Main transfer risk	Dependence on one topology and sensor layout; shortcut learning [13,17]	Distribution drift, calibration change, and sample or feature limits; unverified in water
Role suggested by the evidence	Network-level diagnosis in fog, cloud, or digital-twin layers	Candidate for local screening at gateway level, pending water-specific and device-level validation

6.2. Reported Performance Values Are Not Comparable Across Studies

The coded studies report F1-scores, detection rates, false-alarm rates, localization accuracies, and state-estimation errors on different datasets, anomaly definitions, splits, and prevalence levels. For this reason the review does not present a cross-study performance figure in the main text. The review does not tabulate reported performance values in the supplementary material either; two values may be compared only when the dataset version, preprocessing, anomaly definition, split, and metric definition are identical, which is the case for none of the pairs in the coded set except for models evaluated on the same public benchmark under its published split [30,51]. Even in that case the comparison describes benchmark performance, not transfer or deployment performance.

6.3. Failure Modes and Controls

A transfer-aware AIoT detector can fail at the data, model, transfer, privacy, and decision layers, and the coded evidence indicates which failures are currently unmeasured. Data-layer failures (missing or relocated sensors, drift, biased simulation, class imbalance) are addressed in simulation by two studies, one through randomized sensor placement [13] and one through the analysis of monitoring-system failures [46], and are otherwise unreported. Model-layer failures (shortcut learning, poor out-of-distribution behavior, false-alarm accumulation) are addressed by two graph studies [13,17] and are otherwise invisible in in-distribution evaluations. Transfer-layer failures (negative transfer when source and target are insufficiently related) are not reported by any water study, because none reports a negative-transfer check. Privacy-layer failures (raw-data leakage, poisoning, non-IID instability) are documented in the federated-learning literature [10] but have no water measurement. Decision-layer failures (uninterpretable alerts, unverified actions) are partly addressed by the three XAI studies and not at all by operator evaluations. The controls that follow from this list are the evaluation blocks of Section 7.2: report source–target relatedness, adaptation data size, target-domain validation, deployment cost, and explanation stability alongside detection metrics.

6.4. Answers to the Research Questions

RQ1 (application evidence). The traceability register contains 76 direct water-specific experimental works (E1, E1p) in six application categories (Table 7). Hydraulic tasks dominate: leak detection and localization (43 works) and burst, hydraulic anomaly, or state estimation (9) together account for 52 of the 76 works; cyber-physical attack detection (8), water quality (9), WWTP process anomaly (4), and sensor fault or data quality (3) are comparatively sparse. Graph neural networks and topology-aware models are used in 28 of the 76 works (the 32 graph-based works of Table 6 also count one E2 graph detector evaluated on the SWaT and WADI testbeds and three E3 graph-learning foundations) and are the main family for hydraulic tasks, while TFMs have no water-specific anomaly-detection instance. By data source, 39 works use simulation, 30 use field, pilot, laboratory-rig or real operational data, 4 use a public water benchmark or testbed, and 3 do not state the source in the accessible text; the 21 E2 benchmark or testbed works add the SWaT, WADI, and BATADAL evidence (Table 8).

RQ2 (transfer). Of the 76 direct water-specific experimental works, 6 evaluate an explicit transfer mechanism: two fine-tune across sites or regimes (T2), one applies domain adaptation to a wastewater-treatment task (T3), two pre-train across graphs or algorithmic structure (T5), and one uses federated collaborative transfer (T6). No work reports cross-network synthetic-to-real transfer (T4) or a pre-trained tabular prior applied to a water task (T7), and none reports a negative-transfer check. Federated learning is present in a single water study, and domain adaptation has one instance. The evidence therefore supports fine-tuning and cross-graph pre-training as feasible transfer strategies, largely in simulation for the graph models, and does not yet support broad claims about privacy-preserving collaborative transfer or twin-to-field transfer in water systems.

RQ3 (complementary roles). The evidence supports a division of roles rather than a ranking (Table 12). Graph-based models are evaluated almost entirely in cloud, offline, or digital-twin settings and provide network-level diagnosis; TFMs are candidates for local screening, a role that follows from general tabular evidence and from the tabular structure of node-level data rather than from water experiments (Section 5.3.3). Digital twins support scenario generation and XAI provides explanations in a few studies; neither has been evaluated with operators. The gap map of Section 7.3 (Figure 4) locates these roles on the AIoT layers, with each component labeled as demonstrated or proposed (Section 7.3).

7. Research Gaps and Future Research Agenda

7.1. Gaps Linked to the Coded Evidence

Table 13 lists the gaps that the synthesis identified, the count from the traceability register that documents each gap, and the implication for future studies. Each count refers to the 76 direct water-specific experimental works (E1, E1p) unless stated otherwise, and the works behind each count are identified by their register identifiers in Supplementary Table S1.

Table 13. Research gaps linked to counts from the traceability register ($n = 76$ direct water-specific experimental works; identifiers in Supplementary Table S1).

Gap	Coded Evidence	Implication
No standard source–target protocol	6 of 76 evaluate an explicit transfer mechanism; 0 report a negative-transfer check	Transfer claims should be accompanied by source and target definitions, adaptation data size, target performance, and a negative-transfer check
Synthetic-to-real transfer is rarely validated	39 of 76 train on simulation; 0 report cross-network synthetic-to-real transfer (T4), and one validates on field events of the same networks [28]	Twin-to-field studies with an explicit performance drop are needed before simulation-trained detectors can be recommended for deployment

Table 13. *Cont.*

Gap	Coded Evidence	Implication
Deployment metrics are absent	device-level latency, memory, energy, or communication metrics were not coded as a variable; 75 of 76 works are evaluated at the cloud/offline, digital-twin or testbed layer (Table 9); 0 report device-level metrics for a TFM	Edge and gateway claims, including for TFMs, should be supported by measurements on the intended hardware
Field validation is limited	30 of 76 use field, pilot, laboratory-rig or real operational data; 39 use simulation; 4 use a public water benchmark or testbed; 3 do not state the source; the 21 E2 works add testbed data	Benchmark and simulation results should be labeled as such; utility partnerships are needed for external validation
Federated learning is thinly evidenced	1 of 76 implements federated training (G032); otherwise the mechanism appears only in general method evidence [10,20] and in review proposals	Federated transfer across utilities or DMAs should be evaluated with non-IID, communication, and security metrics before it is presented as a solution to data-sharing constraints
XAI is rarely evaluated and never with operators	3 of 76 works report an explanation method (coded XAI method, Table 7); 0 report an operator evaluation	Explanation stability after transfer and operator usefulness should be reported as acceptance criteria
TFMs lack water anomaly evidence	0 of 76 use a TFM for water anomaly detection; the four TFM works in the register are general method evidence (E3), and the one water-adjacent TabPFN study concerns water-quality index prediction and lies outside the register	Water-specific anomaly studies with drift and calibration changes are needed before the class can be recommended as a local screener
Edge screening and network diagnosis are not integrated	0 of 76 are evaluated at a dedicated edge or fog layer; 0 combine local tabular screening with a network-level graph model	Two-stage designs should be evaluated for false-alarm reduction and localization gain
Reproducibility and benchmark alignment are weak	cyber-physical evidence rests on a few shared testbeds (SWaT, WADI, BATADAL); no public benchmark provides source–target splits for transfer	Benchmark packages with topology, sensor placement, anomaly labels, transfer splits, and twin scenario files would enable comparable transfer studies

7.2. Future Research Agenda

The agenda follows from the gaps and is framed as testable study designs rather than as generic recommendations. It merges the evaluation protocol and the reporting checklist that the synthesis implies.

Evaluation design. Future transfer-aware studies should report four evaluation blocks, because each answers a different question (Table 14). In-distribution detection shows whether the detector solves the benchmark task; out-of-distribution transfer shows whether it generalizes to a new topology, sensor layout, utility, or anomaly mechanism; deployment feasibility shows whether it can run at the intended AIoT layer; and decision usefulness shows whether operators receive stable and actionable evidence. Among the 76 direct water-specific experimental works, most report the first block, 7 report part of the second, only a small number report device-level feasibility for the third, and none report the fourth.

Prioritized study designs. Five designs would address the largest gaps. (1) A cross-network GNN transfer benchmark that trains on one or more source WDNs and evaluates on at least one unseen target network under altered sensor placement, extending the two simulated studies and the single simulation-to-field validation to field targets in other networks [13,17,28]. (2) A two-stage pipeline in which a TFM or compact tabular model screens local windows at a gateway and a graph model diagnoses spatial context at the fog or twin layer, evaluated for false-alarm reduction, localization gain, and device-level cost; this design would provide the water anomaly evidence that TFMs currently lack. (3) A federated transfer study across simulated utilities or real utility partitions, reporting rare-event detection, communication rounds, client imbalance, personalization strategy, and robustness to non-IID data [10,20]. (4) A synthetic-to-real digital-twin study that generates leak, burst, quality, and attack scenarios in a calibrated twin, adapts the detector to field data, and reports the performance drop [38,46]. (5) An explanation-stability study that

compares attributions, attention, fuzzy rules, and graph saliency before and after transfer and assesses them with operators [40–43]. (6) A local-screening study that compares a TabPFN-family model with a time-series foundation model [36,37] on the same water sensor windows, reporting detection quality, calibration under drift, and device-level cost at the gateway.

Table 14. Recommended evaluation design for transfer-aware AIoT water anomaly-detection studies.

Evaluation Block	Question Answered	Minimum Reporting	Why It Is Needed
In-distribution detection	Does the model solve the task under the reported split?	Precision, recall, F1, PR-AUC, false-alarm rate, anomaly prevalence, confusion matrix	Necessary but insufficient for rare events; the block reported by most coded studies
Out-of-distribution transfer	Does the model adapt to a new topology, sensor layout, utility, or anomaly mechanism?	Source and target definitions, source–target performance drop, adaptation data size, negative-transfer check	Tests the transfer claim directly; reported in part by only 6 of the 76 direct water-specific works
Deployment feasibility	Can the detector run at the intended layer?	Latency, memory, model size, communication rounds, energy proxy on the intended hardware	Edge, fog, cloud, and federated settings impose different constraints; 75 of the 76 water works are evaluated offline, in a twin or on a testbed, and device-level metrics are never reported for a TFM in water
Decision usefulness	Does the system support trusted operational action?	Explanation stability before and after transfer, localization clarity, operator feedback	Translates model output into operational decisions; reported by none of the works

Minimum reporting checklist. Studies that build on this review should report the data source and its type (simulation, testbed, field); the anomaly type, definition, and prevalence; topology availability and sensor placement; the train–validation–test strategy and, for transfer studies, the source–target split and adaptation data size; the computational footprint at the intended layer; the privacy mechanism, if any; and the explanation method and its stability. Reporting these items would make the literature usable by utilities, which need to know not whether a model worked in one paper but whether it can be transferred to their own infrastructure with limited labels, different sensors, strict data-sharing rules, and a need for explainable alarms.

7.3. A Gap Map for Transfer-Aware AIoT Anomaly Detection

Figure 4 is a gap map, not a proposed architecture. It places on the AIoT layers each component that the coded evidence has demonstrated in at least one water-specific experimental study (solid elements) and each component that remains a proposal resting on general method evidence or on reviews (dashed elements), and it is derived from the gaps of Section 7.1 and the agenda of Section 7.2: every dashed element corresponds to an agenda item. No reviewed study implements or validates the layout as a whole, and the map does not advance a new architecture; its purpose is to show where the water evidence stops. Table 15 makes the same distinction explicit for every component.

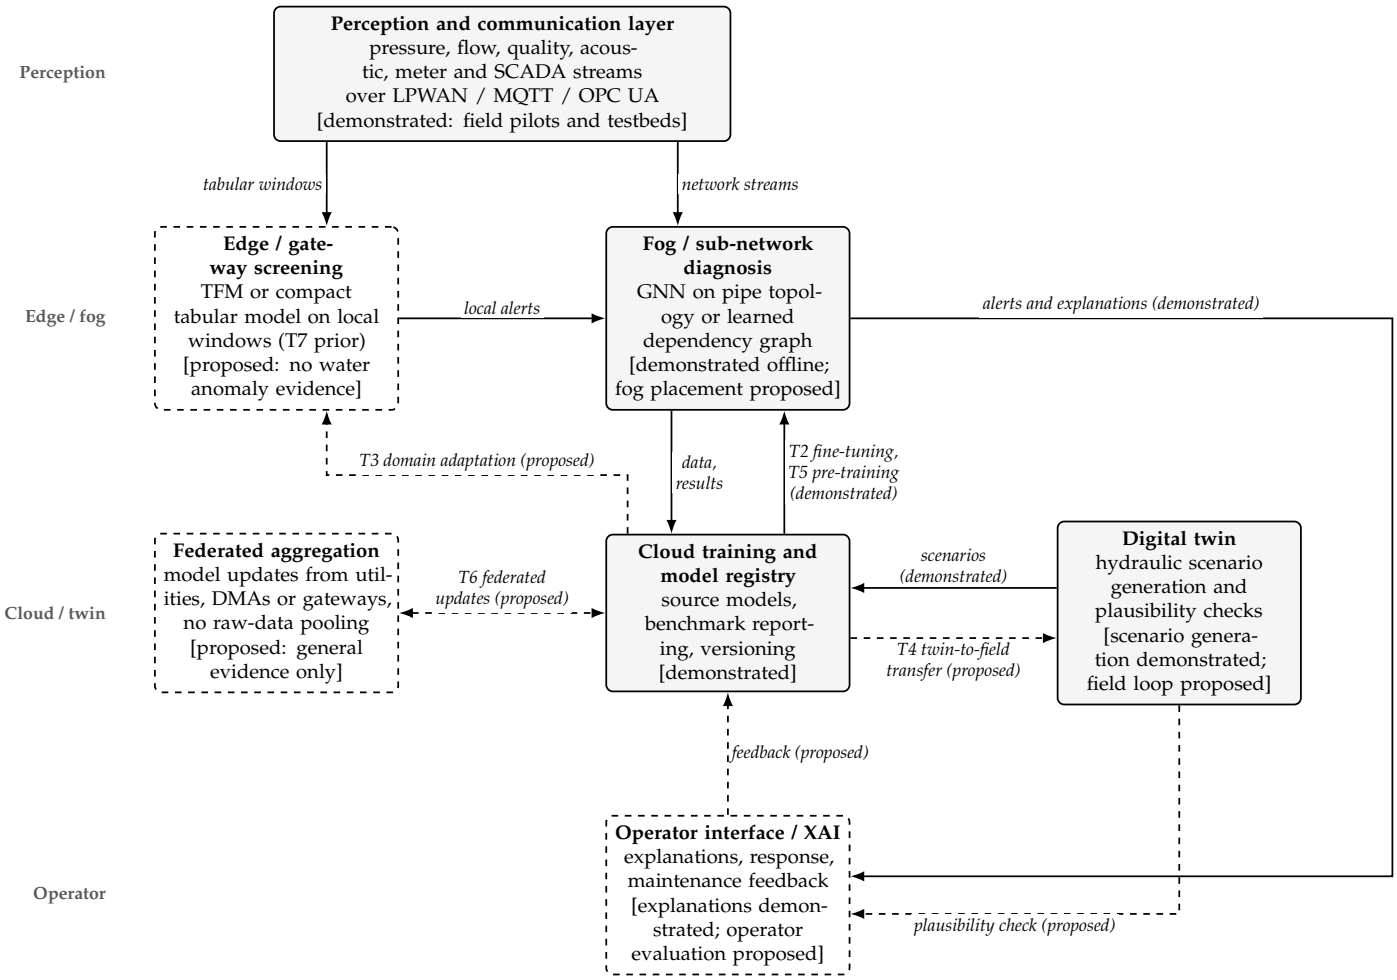


Figure 4. Gap map for transfer-aware AIoT anomaly detection, derived from the gaps and agenda of Section 7. Solid boxes and arrows have support in at least one water-specific experimental study; dashed ones are proposals resting on general method evidence or reviews (Table 15). T2–T7 are the transfer categories of Table 10; T4 has no instance (the single same-network simulation-to-field validation [28] is coded T0), so the twin-to-field arrow remains proposed. TFM: tabular foundation model.

Table 15. Components of the gap map, separated into those supported by reviewed empirical water literature and those that remain proposals with a stated validation need.

Component	Supported by Reviewed Empirical Water Literature	Proposed Synthesis and Future Validation Need
Perception and communication	LoRaWAN metering pilot [9]; testbed instrumentation [49,50]; IoT water-quality sensing [57]	Sensor-placement changes as an explicit transfer mismatch (tested in simulation only [13])
Edge screening with TFMs	No direct water-specific TFM anomaly-detection evidence. One water-adjacent TabPFN study concerns irrigation water-quality index prediction rather than anomaly detection; the remaining support is general tabular-learning or industrial-IoT evidence.	Water-specific anomaly validation; device-level latency, memory, and energy measurements; drift handling; comparison with calibrated tree-based and temporal baselines.
Fog or cloud graph diagnosis	Leak, pressure, consumption, contamination, and attack detection with graph models (28 graph-based works, Table 7; Figure 3)	Placement at the fog layer with compression; evaluation under sensor removal and topology change
Transfer between systems (T2, T5)	Fine-tuning across sites or regimes (G033, G145; T2) and cross-graph or algorithm-informed pre-training on unseen simulated networks (G018, G019; T5)	Standard source–target protocol with adaptation data size and negative-transfer checks; field targets

Table 15. *Cont.*

Component	Supported by Reviewed Empirical Water Literature	Proposed Synthesis and Future Validation Need
Domain adaptation (T3)	One wastewater-treatment fault-diagnosis study (G144)	Alignment across demand regimes and calibrations; comparison with fine-tuning
Synthetic-to-real transfer (T4)	None across networks; 39 of the 76 works train on simulation, one of which validates on field events of the same networks [28]	Twin-to-field evaluation across networks with the performance drop reported
Federated aggregation (T6)	One water study, a federated GAN for hydrological telemetry (G032) [E1]; otherwise general method evidence [10,20]	Non-IID robustness, communication cost, personalization, and poisoning resilience for water clients
Digital twin	Scenario generation for detector training and testing [38]; graph-convolutional twin for pump-speed and state estimation on benchmark networks [46]	Closed loop from twin scenarios to field alerts; calibration reporting
XAI and operator interface	Fuzzy rules [42]; attention and learned graphs [30]; attribution stability under transfer [43]	Explanation stability after transfer as an acceptance criterion; operator studies

8. Conclusions

This systematic review asked whether anomaly detectors for AIoT-enabled water infrastructures transfer across heterogeneous systems. On the basis of a traceable register of 150 works, of which 76 are direct water-specific experimental studies, the evidence supports three conclusions. First, graph-based models have the strongest water-specific evidence for topology-dependent tasks and are the only family with evidence of generalization to unseen networks, although only in simulation. Second, tabular foundation models, which in the retrieved evidence means the TabPFN family, rest on general tabular-learning evidence and one water-adjacent prediction study; their use as local anomaly screeners and their feasibility on embedded devices remain proposals. Third, transfer is the least developed axis: only 6 of the 76 experimental works evaluate an explicit transfer mechanism, no work reports cross-network synthetic-to-real transfer or a negative-transfer check, federated learning appears in a single water study, and no study evaluates explanations with operators.

The contribution of the review is the apparatus that makes these statements auditable rather than the statements themselves: an evidence-class scheme that keeps water-specific results apart from general method claims, a transfer taxonomy with counts, a gap map that locates demonstrated and proposed components on the AIoT layers, and an evaluation agenda built on source–target protocols. Its main limitations are the single-coder screening and coding, quantified for a sample by a second coder, the partial reliance on metadata rather than full text for some register entries, and the absence of a subscription-database search; every coding decision is recorded in the register so that it can be checked. The next step for the field is therefore not a new model family but a change in evaluation practice, in which water anomaly detectors are reported with their transfer behavior, deployment cost, and explanation stability. For sustainable water management the implication is concrete: the components already supported by evidence can reduce water losses and response times where topology and data exist, whereas the components that would carry these gains to data-poor utilities, cross-network transfer, federated adaptation, and embedded tabular screening, are those for which evidence is thinnest, and closing that gap is where research on AIoT anomaly detection can most directly serve the sustainability of water services.

Supplementary Materials: The following are available at https://www.mdpi.com/article/10.3390/su182111431/s1, Supplementary Table S1, the verified global traceability register of the exact 150 retained scholarly works (identifiers G001–G150, source information, DOI or stable URL, evidence class, search provenance, coding status, transfer category, deployment layer, and inclusion role), provided as a machine-readable file (`data/S1_verified_global_traceability_register_150.csv` and `supplementary/S1_verified_global_traceability_register_150.xlsx`) together with the

modular search blocks, the search-route and verification log, and the PRISMA counts; Supplementary Table S2, the coding variables applied to the register; Supplementary Table S3, evidence-class definitions and quality-appraisal criteria; Supplementary Figures S1–S4, layer diagrams of GNN message passing, transfer-learning pipeline, federated learning, and digital-twin/XAI integration; Supplementary Table S4, the completed PRISMA 2020 checklist and PRISMA 2020 for Abstracts checklist.

Author Contributions: Conceptualization, B.B., A.P.-L. and J.M.N.R.; methodology, B.B., A.P.-L. and J.M.N.R.; investigation, B.B.; data curation, B.B.; writing—original draft preparation, B.B.; writing—review and editing, B.B., A.P.-L. and J.M.N.R.; visualization, B.B.; supervision, A.P.-L. and J.M.N.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The verified traceability register (Supplementary Table S1, machine-readable), the per-study coding sheet (Supplementary Table S2; `data/evidence_coding.csv`), the second-coder sample and disagreement list (Supplementary Table S3d) and the PRISMA checklists (Supplementary Table S4) are deposited in the Open Science Framework project associated with the registration of this review (DOI: 10.17605/OSF.IO/D7RXU), and are also provided with the submission. No raw utility sensor data were used in this review.

Acknowledgments: The authors thank B. Al Khaled for the independent second coding of the 30-work reliability sample described in Section 3.4. During the preparation of this manuscript, the authors used Claude (Anthropic; 2026 model versions) for language editing and the drafting of revised text, for the LaTeX typesetting of tables and figures, for consistency checks of the reported counts against the traceability register, and for support in re-examining the data-source field of 21 register entries during the revision. Screening, eligibility and coding decisions were made by the authors and by the second coder. The authors reviewed and edited all outputs and take full responsibility for the content of this publication.

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

Abbreviations

The following abbreviations are used in this manuscript:

AIoT	Artificial Intelligence of Things
BATADAL	Battle of the Attack Detection Algorithms
DL	Deep learning
DMA	District metered area
DT	Digital twin
FL	Federated learning
GAN	Generative adversarial network
GNN	Graph neural network
ICS	Industrial control system
IID	Independent and identically distributed
IoT	Internet of Things
LSTM	Long short-term memory
ML	Machine learning
OSF	Open Science Framework
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
RQ	Research question
SCADA	Supervisory control and data acquisition
SWaT	Secure Water Treatment testbed
TabPFN	Tabular Prior-Data Fitted Network

TFM	Tabular foundation model
TL	Transfer learning
WADI	Water Distribution testbed
WDN	Water distribution network
WWTP	Wastewater treatment plant
XAI	Explainable artificial intelligence

Appendix A. Included Works of the Traceability Register

Table A1 lists the 150 unique scholarly works of the traceability register (Supplementary Table S1) with their evidence class, water domain, application category or role, and reference, so that every included work is cited. Evidence classes are defined in Table 5: E1, peer-reviewed water-specific experiment; E1p, water-specific experiment not journal peer reviewed; E2, benchmark or testbed study; E3, general method evidence; C1, review, conceptual, or contextual work. The full coding of every work (dataset, model family, transfer category, deployment layer, quality appraisal, and coding basis) is given in Supplementary Tables S1 and S2.

Table A1. Included works ($n = 150$) by register identifier, evidence class, water domain, and application category or role, with the corresponding reference.

ID	Class	Water Domain	Application Category or Role	Ref.
G001	E1	WDN	burst/hydraulic anomaly and state estimation	[58]
G002	E1	WDN	leak detection/localization	[59]
G003	E1	WDN	leak detection/localization	[60]
G004	E1	WDN	leak detection/localization	[61]
G005	E1	WDN	leak detection/localization	[62]
G006	E1	WDN	leak detection/localization	[63]
G007	E1	WDN	leak detection/localization	[64]
G008	E1	WDN	leak detection/localization	[65]
G009	E1	WDN	leak detection/localization	[66]
G010	E1	WDN	leak detection/localization	[67]
G011	E1	WDN	leak detection/localization	[68]
G012	E1	WDN	leak detection/localization	[69]
G013	E1	WDN	leak detection/localization	[70]
G014	E1	WDN	burst/hydraulic anomaly and state estimation	[71]
G015	E1	WDN	burst/hydraulic anomaly and state estimation	[72]
G016	E1	WDN	leak detection/localization	[73]
G017	E1	WDN	leak detection/localization	[28]
G018	E1p	WDN	leak detection/localization	[17]
G019	E1	WDN	burst/hydraulic anomaly and state estimation	[74]
G020	E1	WDN	burst/hydraulic anomaly and state estimation	[75]
G021	E1	WDN	burst/hydraulic anomaly and state estimation	[13]
G022	E1	WDN	burst/hydraulic anomaly and state estimation	[76]
G023	E1	WDN	burst/hydraulic anomaly and state estimation	[77]
G024	E1	water quality	water quality	[78]
G025	E1	WDN	water quality	[79]

Table A1. *Cont.*

ID	Class	Water Domain	Application Category or Role	Ref.
G026	E1	river or environmental water sensing	water quality	[80]
G027	E1	water quality	water quality	[81]
G028	E1	water quality	sensor fault/data quality	[82]
G029	E1	water quality	sensor fault/data quality	[83]
G030	E1	water treatment works	water quality	[84]
G031	E1	WWTP	WWTP process anomaly	[85]
G032	E1	hydrological telemetry	sensor fault/data quality	[86]
G033	E1	WDN	burst/hydraulic anomaly and state estimation	[16]
G034	E1	WDN	cyber-physical attack	[87]
G035	E1	WDN	cyber-physical attack	[88]
G036	E1	WDN	cyber-physical attack	[89]
G037	E1	WDN	cyber-physical attack	[90]
G038	E1	WDN	cyber-physical attack	[91]
G039	E1	WDN	cyber-physical attack	[92]
G040	E2	water cyber-physical testbed	cyber-physical attack	[49]
G041	E2	water cyber-physical testbed	cyber-physical attack	[50]
G042	E2	water cyber-physical testbed	cyber-physical attack	[93]
G043	E2	water cyber-physical testbed	cyber-physical attack	[3]
G044	E2	water cyber-physical testbed	cyber-physical attack	[94]
G045	C1	WDN	review/survey	[95]
G046	C1	WDN	review/survey	[96]
G047	C1	WDN	review/survey	[97]
G048	C1	water cyber-physical testbed	review/survey	[2]
G049	C1	water cyber-physical testbed	review/survey	[98]
G050	C1	WDN	review/survey	[99]
G051	C1	WDN	review/survey	[100]
G052	C1	WDN	review/survey	[101]
G053	C1	water quality	review/survey	[102]
G054	C1	WDN	review/survey	[103]
G055	C1	smart metering	review/survey	[104]
G056	E3	general (not water)	general method/reference	[24]
G057	E3	general (not water)	general method/reference	[18]
G058	E3	general (not water)	general method/reference	[105]
G059	E3	general (not water)	general method/reference	[106]
G060	E3	general (not water)	general method/reference	[20]
G061	E3	general (not water)	general method/reference	[10]
G062	E3	general (not water)	general method/reference	[25]
G063	E3	general (not water)	general method/reference	[26]
G064	E3	general (not water)	general method/reference	[107]
G065	E3	general (not water)	general method/reference	[11]
G066	E3	general (not water)	general method/reference	[108]
G067	E3	general (not water)	general method/reference	[14]

Table A1. *Cont.*

ID	Class	Water Domain	Application Category or Role	Ref.
G068	E3	general (not water)	general method/reference	[109]
G069	E3	general (not water)	general method/reference	[110]
G070	E3	general (not water)	general method/reference	[41]
G071	E3	general (not water)	general method/reference	[40]
G072	E3	general (not water)	general method/reference	[111]
G073	E3	general (not water)	general method/reference	[112]
G074	E3	general (not water)	general method/reference	[113]
G075	E3	general (not water)	general method/reference	[114]
G076	E3	general (not water)	general method/reference	[115]
G077	E3	general (not water)	general method/reference	[116]
G078	E3	general (not water)	general method/reference	[117]
G079	E2	water treatment testbed	cyber-physical anomaly detection	[118]
G080	E2	water cyber-physical testbed	cyber-physical anomaly detection	[119]
G081	E2	water cyber-physical testbed	cyber-physical attack detection	[120]
G082	C1	general (not water)	general method/reference	[121]
G083	E3	general (not water)	general method/reference	[122]
G084	C1	urban water security	water security assessment	[123]
G085	C1	general (not water)	general method/reference	[124]
G086	C1	general (not water)	general method/reference	[125]
G087	C1	general (not water)	general method/reference	[126]
G088	E2	water cyber-physical testbed	multivariate anomaly detection	[30]
G089	C1	general (not water)	general method/reference	[127]
G090	C1	general (not water)	general method/reference	[128]
G091	E2	water cyber-physical testbed	cyber-physical attack detection	[129]
G092	E3	general (not water)	general method/reference	[130]
G093	C1	watershed water security	water security assessment	[131]
G094	E3	general (not water)	general method/reference	[132]
G095	C1	general (not water)	general method/reference	[133]
G096	C1	water security	water security assessment review	[134]
G097	E3	general (not water)	general method/reference	[135]
G098	C1	general (not water)	general method/reference	[136]
G099	E2	water cyber-physical testbed	cyber-physical anomaly detection	[137]
G100	E2	clean-water supply testbed	anomaly detection	[138]
G101	E3	general (not water)	general method/reference	[139]
G102	E2	water treatment testbed	predictive safety / anomaly impact	[140]
G103	E2	water cyber-physical testbed	multivariate anomaly detection	[141]
G104	C1	general (not water)	general method/reference	[142]
G105	E2	water treatment testbed	continuous cyber-physical security	[143]
G106	E3	general (not water)	general method/reference	[144]
G107	E2	water treatment testbed	cyber-physical attack detection	[145]
G108	E2	water treatment testbed	cyber-physical attack detection	[146]

Table A1. *Cont.*

ID	Class	Water Domain	Application Category or Role	Ref.
G109	E2	water treatment testbed	cyber-physical attack detection	[147]
G110	E3	general (not water)	general method/reference	[148]
G111	E1	WDN	water quality	[149]
G112	E1	WDN	cyber-physical attack	[150]
G113	E1	WDN	leak detection/localization	[151]
G114	E1	WDN	leak detection/localization	[152]
G115	E1	WDN	leak detection/localization	[153]
G116	E1	water quality	water quality	[154]
G117	E1	WDN	cyber-physical attack	[52]
G118	E1	WDN	leak detection/localization	[155]
G119	E1	WDN	leak detection/localization	[156]
G120	E1	WWTP	WWTP process anomaly	[157]
G121	E1	WDN	leak detection/localization	[158]
G122	E1	WWTP	WWTP process anomaly	[159]
G123	E1	water quality	water quality	[160]
G124	E1	WDN	leak detection/localization	[161]
G125	E1	WDN	leak detection/localization	[162]
G126	E1	WDN	leak detection/localization	[163]
G127	E1	WDN	leak detection/localization	[164]
G128	E2	water cyber-physical testbed	cyber-physical attack	[165]
G129	E1	WDN	leak detection/localization	[166]
G130	E1	WDN	leak detection/localization	[167]
G131	E1	WDN	leak detection/localization	[168]
G132	E1	WDN	leak detection/localization	[169]
G133	E1	WDN	leak detection/localization	[170]
G134	E1	WDN	leak detection/localization	[38]
G135	E1	WDN	leak detection/localization	[171]
G136	E1	WDN	leak detection/localization	[172]
G137	E1	WDN	leak detection/localization	[173]
G138	E1	WDN	leak detection/localization	[174]
G139	E1	river or environmental water sensing	water quality	[175]
G140	E1	WDN	leak detection/localization	[176]
G141	E1	WDN	leak detection/localization	[177]
G142	E1	WDN	leak detection/localization	[178]
G143	E2	WDN	cyber-physical attack	[179]
G144	E1	WWTP	WWTP process anomaly	[180]
G145	E1	WDN	leak detection/localization	[181]
G146	E1	WDN	leak detection/localization	[182]
G147	E1	WDN	leak detection/localization	[183]
G148	E2	WWTP	WWTP process anomaly	[184]
G149	E1	WDN	leak detection/localization	[185]
G150	E1	WDN	leak detection/localization	[186]

References

1. Liso, A.; Cardellicchio, A.; Patruno, C.; Nitti, M.; Ardino, P.; Stella, E.; Renò, V. A Review of Deep Learning-Based Anomaly Detection Strategies in Industry 4.0 Focused on Application Fields, Sensing Equipment, and Algorithms. *IEEE Access* **2024**, *12*, 93911–93923. https://doi.org/10.1109/ACCESS.2024.3424488.
2. Tuptuk, N.; Hazell, P.; Watson, J.; Hailes, S. A Systematic Review of the State of Cyber-Security in Water Systems. *Water* **2021**, *13*, 81. https://doi.org/10.3390/w13010081.

3. Taormina, R.; Galelli, S.; Tippenhauer, N.O.; Salomons, E.; Ostfeld, A.; Eliades, D.G.; Aghashahi, M.; Sundararajan, R.; Pourahmadi, M.; Banks, M.K.; et al. Battle of the Attack Detection Algorithms: Disclosing Cyber Attacks on Water Distribution Networks. *J. Water Resour. Plan. Manag.* **2018**, *144*, 04018048. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000969.
4. Rajan, G.; Li, S. A Systematic Literature Review on Flow Data-Based Techniques for Automated Leak Management in Water Distribution Systems. *Smart Cities* **2025**, *8*, 78. https://doi.org/10.3390/smartcities8030078.
5. Yang, S.; Behzadian, K.; Coleman, C.; Holloway, T.G.; Campos, L.C. Application of AI-based techniques for anomaly management in wastewater treatment plants: A review. *J. Environ. Manag.* **2025**, *392*, 126886. https://doi.org/10.1016/j.jenvman.2025.126886.
6. Rahman, A.; Chung, G.C.; Ng, Y.H. Applications, Challenges, and Future Trends of Artificial Intelligence of Things (AIoT)-Enabled Water Quality and Resource Management. *Water* **2026**, *18*, 919. https://doi.org/10.3390/w18080919.
7. Amanatidis, P.; Lyratzis, E.; Angelopoulos, V.; Kouloumpris, E.; Skaperdas, E.; Bassiliades, N.; Vlahavas, I.; Maris, F.; Emmanouloudis, D.; Karampatzakis, D. Intelligent Water Management Through Edge-Enabled IoT, AI, and Big Data Technologies. *IoT* **2025**, *7*, 5. https://doi.org/10.3390/iot7010005.
8. Pagano, A.; Garlisi, D.; Tinnirello, I.; Giuliano, F.; Garbo, G.; Falco, M.; Cuomo, F. A survey on massive IoT for water distribution systems: Challenges, simulation tools, and guidelines for large scale deployment. *Ad Hoc Netw.* **2025**, *168*, 103714. https://doi.org/10.1016/j.adhoc.2024.103714.
9. Slaný, V.; Lučanský, A.; Koudelka, P.; Mareček, J.; Krčálová, E.; Martínek, R. An Integrated IoT Architecture for Smart Metering Using Next Generation Sensor for Water Management Based on LoRaWAN Technology: A Pilot Study. *Sensors* **2020**, *20*, 4712. https://doi.org/10.3390/s20174712.
10. Kairouz, P.; McMahan, H.B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A.N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. Advances and Open Problems in Federated Learning. *Found. Trends Mach. Learn.* **2021**, *14*, 1–210. https://doi.org/10.1561/2200000083.
11. Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; Yu, P.S. A Comprehensive Survey on Graph Neural Networks. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *32*, 4–24. https://doi.org/10.1109/TNNLS.2020.2978386.
12. Vittori, G.; Falkouskaya, Y.; Jimenez-Gutierrez, D.M.; Cattai, T.; Chatzigiannakis, I. Graph neural networks to model and optimize the operation of Water Distribution Networks: A review. *J. Ind. Inf. Integr.* **2025**, *47*, 100880. https://doi.org/10.1016/j.jii.2025.100880.
13. Truong, H.; Tello, A.; Lazovik, A.; Degeler, V. Graph Neural Networks for Pressure Estimation in Water Distribution Systems. *Water Resour. Res.* **2024**, *60*, e2023WR036741. https://doi.org/10.1029/2023WR036741.
14. Hollmann, N.; Müller, S.; Purucker, L.; Krishnakumar, A.; Körfer, M.; Hoo, S.B.; Schirrmeister, R.T.; Hutter, F. Accurate predictions on small data with a tabular foundation model. *Nature* **2025**, *637*, 319–326. https://doi.org/10.1038/s41586-024-08328-6.
15. Grinsztajn, L.; Flöge, K.; Key, O.; Birkel, F.; Jund, P.; Roof, B.; Jäger, B.; Safaric, D.; Alessi, S.; Hayler, A.; et al. TabPFN-2.5: Advancing the State of the Art in Tabular Foundation Models. *arXiv* **2025**, arXiv:2511.08667.
16. Glynis, K.; Kapelan, Z.; Bakker, M.; Taormina, R. Leveraging Transfer Learning in LSTM Neural Networks for Data-Efficient Burst Detection in Water Distribution Systems. *Water Resour. Manag.* **2023**, *37*, 5953–5972. https://doi.org/10.1007/s11269-023-03637-3.
17. Zhang, Z.; Fink, O. Algorithm-informed graph neural networks for leakage detection and localization in water distribution networks. *Reliab. Eng. Syst. Saf.* **2026**, *265*, 111494. https://doi.org/10.1016/j.ress.2025.111494.
18. Pan, S.J.; Yang, Q. A Survey on Transfer Learning. *IEEE Trans. Knowl. Data Eng.* **2010**, *22*, 1345–1359. https://doi.org/10.1109/TKDE.2009.191.
19. Weiss, K.; Khoshgoftaar, T.M.; Wang, D. A survey of transfer learning. *J. Big Data* **2016**, *3*, 9. https://doi.org/10.1186/s40537-016-0043-6.
20. McMahan, H.B.; Moore, E.; Ramage, D.; Hampson, S.; Agüera Y Arcas, B. Communication-Efficient Learning of Deep Networks from Decentralized Data. *Proc. 20th Int. Conf. Artif. Intell. Stat. (AISTATS)* **2017**, *PMLR 54*, 1273–1282.
21. Zuñiga-Uribe, M.; Rojas-Galván, R.; Álvarez-Alvarado, J.M.; Aviles, M.; Pérez-Soto, G.I.; Pérez-Moreno, V. Artificial Intelligence in Water Distribution Networks: A Systematic Review of Models, Input Variables, Databases, and Output Strategies for Leak Detection. *Smart Cities* **2026**, *9*, 45. https://doi.org/10.3390/smartcities9030045.
22. Liu, T.; Liu, F. Graph Neural Networks for Evaluating the Reliability and Resilience of Infrastructure Systems: A Systematic Review of Models, Applications and Future Directions. *IEEE Access* **2025**, *13*, 164883–164904. https://doi.org/10.1109/ACCESS.2025.3611333.
23. Habib, M.A.; Jurcut, A.D.; Ahmed, H.; Wei, W.; Salauddin, M. Cybersecurity in Water Distribution Networks: A Systematic Review of AI-Based Detection Algorithms. *Water* **2026**, *18*, 519. https://doi.org/10.3390/w18040519.
24. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ* **2021**, *372*, n71. https://doi.org/10.1136/bmj.n71.
25. Kipf, T.N.; Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. In Proceedings of the 5th International Conference on Learning Representations (ICLR), Toulon, France, 24–26 April 2017.

26. Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; Bengio, Y. Graph Attention Networks. In Proceedings of the 6th International Conference on Learning Representations (ICLR), Vancouver, BC, Canada, 30 April–3 May 2018.
27. Battaglia, P.W.; Hamrick, J.B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; et al. Relational inductive biases, deep learning, and graph networks. *arXiv* **2018**, arXiv:1806.01261.
28. Li, X.; Wu, Y. A Convolutional Graph Neural Network Model for Water Distribution Network Leakage Detection Based on Segment Feature Fusion Strategy. *Water* **2024**, *16*, 3555. https://doi.org/10.3390/w16243555.
29. Caetano, J.; Carriço, N.; Brentan, B.; Menapace, A.; Covas, D. Topology-Aware Neural Networks for Abnormal Consumption Detection and Location in Water Distribution Networks. *Water Resour. Res.* **2026**, *62*, e2025WR041195. https://doi.org/10.1029/2025WR041195.
30. Deng, A.; Hooi, B. Graph Neural Network-Based Anomaly Detection in Multivariate Time Series. *Proc. AAAI Conf. Artif. Intell.* **2021**, *35*, 4027–4035. https://doi.org/10.1609/aaai.v35i5.16523.
31. Grinsztajn, L.; Oyallon, E.; Varoquaux, G. Why do tree-based models still outperform deep learning on typical tabular data? *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 507–520.
32. Huang, X.; Yao, R.; Zhang, Y.; Li, X.; Yu, Z.; Guo, H. A novel irrigation water quality index (PCIWQI) assessment and prediction: Insights into model performance and computational efficiency from TabPFN and tree-based models. *J. Environ. Chem. Eng.* **2026**, *14*, 120934. https://doi.org/10.1016/j.jece.2025.120934.
33. Ruiz-Villafranca, S.; Roldán-Gómez, J.; Castelo Gómez, J.M.; Carrillo-Mondéjar, J.; Martinez, J.L. A TabPFN-based intrusion detection system for the industrial internet of things. *J. Supercomput.* **2024**, *80*, 20080–20117. https://doi.org/10.1007/s11227-024-06166-x.
34. Ruiz-Villafranca, S.; Roldán-Gómez, J.; Carrillo-Mondéjar, J.; Martinez, J.L.; Gañán, C.H. WFE-Tab: Overcoming limitations of TabPFN in IIoT-MEC environments with a weighted fusion ensemble-TabPFN model for improved IDS performance. *Future Gener. Comput. Syst.* **2025**, *166*, 107707. https://doi.org/10.1016/j.future.2025.107707.
35. Choi, J.; Kang, W.; Kim, M.; Kim, J.; Park, N. Can TabPFN Compete with GNNs for Node Classification via Graph Tabularization? *arXiv* **2025**, arXiv:2512.08798. .
36. Ansari, A.F.; Stella, L.; Turkmen, C.; Zhang, X.; Mercado, P.; Shen, H.; Shchur, O.; Rangapuram, S.S.; Pineda Arango, S.; Kapoor, S.; et al. Chronos: Learning the Language of Time Series. *Trans. Mach. Learn. Res.* **2024**, arXiv:2403.07815.
37. Das, A.; Kong, W.; Sen, R.; Zhou, Y. A decoder-only foundation model for time-series forecasting. *Proc. 41st Int. Conf. Mach. Learn. (ICML 2024)* **2024**, *235*, 10148–10167.
38. Pandey, P.; Mücke, N.T.; Jain, S.; Ramachandran, P.; Bohté, S.M.; Oosterlee, C.W. Machine Learning-Based Digital Twin for Water Distribution Network Anomaly Detection and Localization. *Eng. Proc.* **2024**, *69*, 201. https://doi.org/10.3390/engproc2024069201.
39. Alourani, A.; Alam, M.; Ali, A.; Khan, I.R.; Samal, C.K. Hybrid AI-IoT Framework with Digital Twin Integration for Predictive Urban Infrastructure Management in Smart Cities. *Comput. Mater. Contin.* **2026**, *86*, 1–32. https://doi.org/10.32604/cmc.2025.070161.
40. Lundberg, S.M.; Lee, S.I. A Unified Approach to Interpreting Model Predictions. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 4765–4774.
41. Ribeiro, M.T.; Singh, S.; Guestrin, C. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; pp. 1135–1144. https://doi.org/10.1145/2939672.2939778.
42. Khaled, Q.; De Marinis, P.; Louati, M.; Ferras, D.; Genga, L.; Kaymak, U. Explainable fuzzy GNNs for leak detection in water distribution networks. *arXiv* **2026**, arXiv:2601.03062. .
43. Oliveira, P.; Bessa, A.; Pereira, J.; Silva, S.; Duarte, M.S.; Durães, D.; Novais, P. Exploring Transfer Learning's Impact on the Explainability of Deep Learning Models for Wastewater Treatment Plants' Biogas Production. *Expert Syst.* **2026**, *43*, e70235. https://doi.org/10.1111/exsy.70235.
44. Rana, S.; Bajwa, A.; Tonoy, A.A.R.; Ahmed, I. Cybersecurity in Industrial Control Systems: A Systematic Literature Review on AI-Based Threat Detection for SCADA and IoT Networks. *SSRN* **2025**, preprint 5267824. https://doi.org/10.2139/ssrn.5267824.
45. Biström, C. Anomaly Detection Methods in Water Infrastructure. Bachelor's Thesis, Aalto University, Espoo, Finland, 2025.
46. Bonilla Granados, C.A. Monitorización de Sistemas de Distribución de Agua y Detección de Anomalías. Ph.D. Thesis, Universitat Politècnica de València, Valencia, Spain, 2025.
47. Vrachimis, S.G.; Kyriakou, M.S.; Eliades, D.G.; Polycarpou, M.M. LeakDB: A benchmark dataset for leakage diagnosis in water distribution networks. In Proceedings of the WDSA/CCWI Joint Conference 2018, Kingston, ON, Canada, 23–25 July 2018; Volume 1.
48. Vrachimis, S.G.; Eliades, D.G.; Taormina, R.; Kapelan, Z.; Ostfeld, A.; Liu, S.; Kyriakou, M.; Pavlou, P.; Qiu, M.; Polycarpou, M.M. Battle of the Leakage Detection and Isolation Methods. *J. Water Resour. Plan. Manag.* **2022**, *148*, 04022068. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001601.

49. Mathur, A.P.; Tippenhauer, N.O. SWaT: A water treatment testbed for research and training on ICS security. In Proceedings of the 2016 International Workshop on Cyber-physical Systems for Smart Water Networks (CySWater), Vienna, Austria, 11 April 2016; pp. 31–36. https://doi.org/10.1109/CySWater.2016.7469060.
50. Ahmed, C.M.; Palleti, V.R.; Mathur, A.P. WADI: A water distribution testbed for research in the design of secure cyber physical systems. In Proceedings of the 3rd International Workshop on Cyber-Physical Systems for Smart Water Networks (CySWater), Pittsburgh, PA, USA, 18-21 April 2017; pp. 25–28. https://doi.org/10.1145/3055366.3055375.
51. Ge, D.; Cheng, Y.; Cao, S.; Ma, Y.; Wu, Y. An enhanced abnormal information expression spatiotemporal model for anomaly detection in multivariate time-series. *Complex Intell. Syst.* **2024**, *10*, 2937–2950. https://doi.org/10.1007/s40747-023-01306-x.
52. Rustam, F.; Jurcut, A.D.; Salauddin, M. GSPO-LAD: A Graph Neural Network-Based Sensor Placement and Anomaly Detection in Water Distribution Systems. *IEEE Access* **2025**, *13*, 149462–149477. https://doi.org/10.1109/ACCESS.2025.3598996.
53. Latif-Martínez, H.; Suárez-Varela, J.; Cabellos-Aparicio, A.; Barlet-Ros, P. GAT-AD: Graph Attention Networks for contextual anomaly detection in network monitoring. *Comput. Ind. Eng.* **2025**, *200*, 110830. https://doi.org/10.1016/j.cie.2024.110830.
54. Latif Martínez, H. Network Anomaly Detection with Graph Neural Networks. Ph.D. Thesis, Universitat Politècnica de Catalunya, Barcelona, Spain, 2025.
55. Qiu, W.; Quan, C.; Yu, Y.; Kara, E.; Qian, K.; Hu, B.; Schuller, B.W.; Yamamoto, Y. Federated Abnormal Heart Sound Detection with Weak to No Labels. *Cyborg Bionic Syst.* **2024**, *5*, 0152. https://doi.org/10.34133/cbsystems.0152.
56. Ma, Z.; Chen, Z.; Zheng, X.; Wang, T.; You, Y.; Zou, S.; Wang, Y. A Biological Immunity-Based Neuro Prototype for Few-Shot Anomaly Detection with Character Embedding. *Cyborg Bionic Syst.* **2024**, *5*, 0086. https://doi.org/10.34133/cbsystems.0086.
57. Anggraini, L.; Senoo, E.E.K.; Soares, I.R.; Shwe, T.; Mendonça, I.; Aritsugi, M. Edge-Driven Water Quality Monitoring and Prediction: A Spatio-Temporal GNN-based IoT Approach for Environmental Sensing. In Proceedings of the IEEE/ACM 12th International Conference on Big Data Computing, Applications and Technologies (BDCAT 2025), Nantes, France, 1–4 December 2025; pp. 1–10. https://doi.org/10.1145/3773276.3774298.
58. Wang, X.; Guo, G.; Liu, S.; Wu, Y.; Xu, X.; Smith, K. Burst Detection in District Metering Areas Using Deep Learning Method. *J. Water Resour. Plan. Manag.* **2020**, *146*, 04020031. https://doi.org/10.1061/(asce)wr.1943-5452.0001223.
59. Zhou, B.; Lau, V.; Wang, X. Machine-Learning-Based Leakage-Event Identification for Smart Water Supply Systems. *IEEE Internet Things J.* **2020**, *7*, 2277–2292. https://doi.org/10.1109/jiot.2019.2958920.
60. Liu, Y.; Ma, X.; Li, Y.; Tie, Y.; Zhang, Y.; Gao, J. Water Pipeline Leakage Detection Based on Machine Learning and Wireless Sensor Networks. *Sensors* **2019**, *19*, 5086. https://doi.org/10.3390/s19235086.
61. Wu, J.; Ma, D.; Wang, W. Leakage Identification in Water Distribution Networks Based on XGBoost Algorithm. *J. Water Resour. Plan. Manag.* **2022**, *148*, 04021107. https://doi.org/10.1061/(asce)wr.1943-5452.0001523.
62. Zhou, M.; Yang, Y.; Xu, Y.; Hu, Y.; Cai, Y.; Lin, J.; Pan, H. A Pipeline Leak Detection and Localization Approach Based on Ensemble TL1DCNN. *IEEE Access* **2021**, *9*, 47565–47578. https://doi.org/10.1109/access.2021.3068292.
63. Guo, G.; Yu, X.; Liu, S.; Ma, Z.; Wu, Y.; Xu, X.; Wang, X.; Smith, K.; Wu, X. Leakage Detection in Water Distribution Systems Based on Time–Frequency Convolutional Neural Network. *J. Water Resour. Plan. Manag.* **2021**, *147*, 04020101. https://doi.org/10.1061/(asce)wr.1943-5452.0001317.
64. Li, J.; Zheng, W.; Lu, C. An Accurate Leakage Localization Method for Water Supply Network Based on Deep Learning Network. *Water Resour. Manag.* **2022**, *36*, 2309–2325. https://doi.org/10.1007/s11269-022-03144-x.
65. Soldevila, A.; Boracchi, G.; Roveri, M.; Tornil-Sin, S.; Puig, V. Leak detection and localization in water distribution networks by combining expert knowledge and data-driven models. *Neural Comput. Appl.* **2021**, *34*, 4759–4779. https://doi.org/10.1007/s00521-021-06666-4.
66. Kammoun, M.; Kammoun, A.; Abid, M. LSTM-AE-WLDL: Unsupervised LSTM Auto-Encoders for Leak Detection and Location in Water Distribution Networks. *Water Resour. Manag.* **2022**, *37*, 731–746. https://doi.org/10.1007/s11269-022-03397-6.
67. Fan, X.; Yu, X.B. An innovative machine learning based framework for water distribution network leakage detection and localization. *Struct. Health Monit.* **2021**, *21*, 1626–1644. https://doi.org/10.1177/14759217211040269.
68. Kang, J.; Park, Y.J.; Lee, J.; Wang, S.H.; Eom, D.S. Novel Leakage Detection by Ensemble CNN-SVM and Graph-Based Localization in Water Distribution Systems. *IEEE Trans. Ind. Electron.* **2018**, *65*, 4279–4289. https://doi.org/10.1109/tie.2017.2764861.
69. Chuang, W.Y.; Tsai, Y.L.; Wang, L.H. Leak Detection in Water Distribution Pipes Based on CNN with Mel Frequency Cepstral Coefficients. In Proceedings of the 2019 3rd International Conference on Innovation in Artificial Intelligence, Suzhou, China, 15–18 March 2019; pp. 83–86. https://doi.org/10.1145/3319921.3319926.
70. Cody, R.A.; Tolson, B.A.; Orchard, J. Detecting Leaks in Water Distribution Pipes Using a Deep Autoencoder and Hydroacoustic Spectrograms. *J. Comput. Civ. Eng.* **2020**, *34*, 04020001. https://doi.org/10.1061/(asce)cp.1943-5487.0000881.
71. Zanfei, A.; Menapace, A.; Brentan, B.M.; Righetti, M.; Herrera, M. Novel approach for burst detection in water distribution systems based on graph neural networks. *Sustain. Cities Soc.* **2022**, *86*, 104090. https://doi.org/10.1016/j.scs.2022.104090.
72. Menapace, A.; Zanfei, A.; Felicetti, M.; Avesani, D.; Righetti, M.; Gargano, R. Burst Detection in Water Distribution Systems: The Issue of Dataset Collection. *Appl. Sci.* **2020**, *10*, 8219. https://doi.org/10.3390/app10228219.

73. Adedeji, K.; Hamam, Y.; Abe, B.; Abu-Mahfouz, A. Leakage Detection and Estimation Algorithm for Loss Reduction in Water Piping Networks. *Water* **2017**, *9*, 773. https://doi.org/10.3390/w9100773.
74. Kerimov, B.; Taormina, R.; Tscheikner-Gratl, F. Towards transferable metamodels for water distribution systems with edge-based graph neural networks. *Water Res.* **2024**, *261*, 121933. https://doi.org/10.1016/j.watres.2024.121933.
75. Kerimov, B.; Bentivoglio, R.; Garzón, A.; Isufi, E.; Tscheikner-Gratl, F.; Steffelbauer, D.B.; Taormina, R. Assessing the performances and transferability of graph neural network metamodels for water distribution systems. *J. Hydroinform.* **2023**, *25*, 2223–2234. https://doi.org/10.2166/hydro.2023.031.
76. Li, Z.; Liu, H.; Wang, Y.; Zhang, C.; Guo, X.; Fu, G. Graph neural networks for hydraulic predictions in water distribution networks: Incorporating sensor data and physical constraints. *Water Res. X* **2025**, *29*, 100448. https://doi.org/10.1016/j.wroa.2025.100448.
77. Xing, L.; Sela, L. Graph Neural Networks for State Estimation in Water Distribution Systems: Application of Supervised and Semisupervised Learning. *J. Water Resour. Plan. Manag.* **2022**, *148*, 04022018. https://doi.org/10.1061/(asce)wr.1943-5452.0001550.
78. Li, Z.; Liu, H.; Zhang, C.; Fu, G. Gated graph neural networks for identifying contamination sources in water distribution systems. *J. Environ. Manag.* **2024**, *351*, 119806. https://doi.org/10.1016/j.jenvman.2023.119806.
79. Li, Z.; Liu, H.; Zhang, C.; Fu, G. Real-time water quality prediction in water distribution networks using graph neural networks with sparse monitoring data. *Water Res.* **2024**, *250*, 121018. https://doi.org/10.1016/j.watres.2023.121018.
80. Sun, A.Y.; Jiang, P.; Yang, Z.L.; Xie, Y.; Chen, X. A graph neural network (GNN) approach to basin-scale river network learning: The role of physics-based connectivity and data fusion. *Hydrol. Earth Syst. Sci.* **2022**, *26*, 5163–5184. https://doi.org/10.5194/hess-26-5163-2022.
81. Muharemi, F.; Logofătu, D.; Leon, F. Machine learning approaches for anomaly detection of water quality on a real-world data set. *J. Inf. Telecommun.* **2019**, *3*, 294–307. https://doi.org/10.1080/24751839.2019.1565653.
82. Leigh, C.; Alsibai, O.; Hyndman, R.J.; Kandanaarachchi, S.; King, O.C.; McGree, J.M.; Neelamraju, C.; Strauss, J.; Talagala, P.D.; Turner, R.D.; et al. A framework for automated anomaly detection in high frequency water-quality data from in situ sensors. *Sci. Total Environ.* **2019**, *664*, 885–898. https://doi.org/10.1016/j.scitotenv.2019.02.085.
83. El-Shafeiy, E.; Alsabaan, M.; Ibrahem, M.I.; Elwahsh, H. Real-Time Anomaly Detection for Water Quality Sensor Monitoring Based on Multivariate Deep Learning Technique. *Sensors* **2023**, *23*, 8613. https://doi.org/10.3390/s23208613.
84. Riss, G.; Romano, M.; Memon, F.A.; Kapelan, Z. Detection of water quality failure events at treatment works using a hybrid two-stage method with CUSUM and random forest algorithms. *Water Supply* **2021**, *21*, 3011–3026. https://doi.org/10.2166/ws.2021.062.
85. Prabu, P.; Alluhaidan, A.S.; Aziz, R.; Basheer, S. Comparative analysis of machine learning models for detecting water quality anomalies in treatment plants. *Sci. Rep.* **2025**, *15*, 30453. https://doi.org/10.1038/s41598-025-15517-4.
86. Chen, Z.; Ni, X.; Li, H.; Kong, X. FedLGAN: A method for anomaly detection and repair of hydrological telemetry data based on federated learning. *PeerJ Comput. Sci.* **2023**, *9*, e1664. https://doi.org/10.7717/peerj-cs.1664.
87. Housh, M.; Ohar, Z. Model-based approach for cyber-physical attack detection in water distribution systems. *Water Res.* **2018**, *139*, 132–143. https://doi.org/10.1016/j.watres.2018.03.039.
88. Taormina, R.; Galelli, S. Deep-Learning Approach to the Detection and Localization of Cyber-Physical Attacks on Water Distribution Systems. *J. Water Resour. Plan. Manag.* **2018**, *144*, 04018065. https://doi.org/10.1061/(asce)wr.1943-5452.0000983.
89. Tsiami, L.; Makropoulos, C. Cyber—Physical Attack Detection in Water Distribution Systems with Temporal Graph Convolutional Neural Networks. *Water* **2021**, *13*, 1247. https://doi.org/10.3390/w13091247.
90. Kadosh, N.; Frid, A.; Housh, M. Detecting Cyber-Physical Attacks in Water Distribution Systems: One-Class Classifier Approach. *J. Water Resour. Plan. Manag.* **2020**, *146*, 04020060. https://doi.org/10.1061/(asce)wr.1943-5452.0001259.
91. Brentan, B.; Rezende, P.; Barros, D.; Meirelles, G.; Luvizotto, E.; Izquierdo, J. Cyber-Attack Detection in Water Distribution Systems Based on Blind Sources Separation Technique. *Water* **2021**, *13*, 795. https://doi.org/10.3390/w13060795.
92. Housh, M.; Kadosh, N.; Haddad, J. Detecting and Localizing Cyber-Physical Attacks in Water Distribution Systems without Records of Labeled Attacks. *Sensors* **2022**, *22*, 6035. https://doi.org/10.3390/s22166035.
93. Goh, J.; Adepu, S.; Junejo, K.N.; Mathur, A. A Dataset to Support Research in the Design of Secure Water Treatment Systems. In *Lecture Notes in Computer Science*; Springer International Publishing: Cham, Switzerland, 2017; pp. 88–99.https://doi.org/10.1007/978-3-319-71368-7_8.
94. Audibert, J.; Michiardi, P.; Guyard, F.; Marti, S.; Zuluaga, M.A. USAD: UnSupervised Anomaly Detection on Multivariate Time Series. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Virtual, 6–10 July 2020; pp. 3395–3404. https://doi.org/10.1145/3394486.3403392.
95. Wu, Y.; Liu, S. A review of data-driven approaches for burst detection in water distribution systems. *Urban Water J.* **2017**, *14*, 972–983. https://doi.org/10.1080/1573062x.2017.1279191.
96. Hu, Z.; Chen, B.; Chen, W.; Tan, D.; Shen, D. Review of model-based and data-driven approaches for leak detection and location in water distribution systems. *Water Supply* **2021**, *21*, 3282–3306. https://doi.org/10.2166/ws.2021.101.

97. Garzón, A.; Kapelan, Z.; Langeveld, J.; Taormina, R. Machine Learning-Based Surrogate Modeling for Urban Water Networks: Review and Future Research Directions. *Water Resour. Res.* **2022**, *58*, e2021WR031808. https://doi.org/10.1029/2021wr031808.
98. Hassanzadeh, A.; Rasekh, A.; Galelli, S.; Aghashahi, M.; Taormina, R.; Ostfeld, A.; Banks, M.K. A Review of Cybersecurity Incidents in the Water Sector. *J. Environ. Eng.* **2020**, *146*, 03120003. https://doi.org/10.1061/(asce)ee.1943-7870.0001686.
99. Vaquet, V.; Hinder, F.; Artelt, A.; Ashraf, I.; Strotherm, J.; Vaquet, J.; Brinkrolf, J.; Hammer, B. Challenges, Methods, Data – a Survey of Machine Learning in Water Distribution Networks. *arXiv* **2024**, arXiv:2410.12461. https://doi.org/10.48550/arxiv.2410.12461.
100. Taloma, R.J.L.; Cuomo, F.; Comminiello, D.; Pisani, P. Machine learning for smart water distribution systems: Exploring applications, challenges and future perspectives. *Artif. Intell. Rev.* **2025**, *58*, 120. https://doi.org/10.1007/s10462-024-11093-7.
101. Wen, H.; Gu, X.; Han, J.; Huo, X.; Ye, Z.; Lu, W. A Review of Anomaly Detection Methods for Smart Water Distribution Networks. In Proceedings of the International Conference on Artificial Intelligence and Advanced Autonomous Systems, Linzhi, China, 3–5 April 2026; pp. 69–75. https://doi.org/10.1145/3813723.3813732.
102. Balta, D.; Balta Kaç, S.; Balta, M.; Eken, S. Water Quality Estimation and Anomaly Detection: A Review. *EAI Endorsed Trans. Internet Things* **2023**, *9*, e2. https://doi.org/10.4108/eetiot.v9i4.3660.
103. Canivete, C.R.C.; Chitauro, M.; Flörke, M.; Okorie, M.E. Deep Learning Applications for Leak Detection and Localisation in Water Distribution Systems: A Systematic Literature Review. *Intell. Infrastruct. Constr.* **2026**, *2*, 10. https://doi.org/10.3390/iic2030010.
104. Kanyama, M.; Bhunu Shava, F.; Gamundani, A.; Hartmann, A. Machine learning applications for anomaly detection in Smart Water Metering Networks: A systematic review. *Phys. Chem. Earth Parts A/B/C* **2024**, *134*, 103558. https://doi.org/10.1016/j.pce.2024.103558.
105. Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; Lempitsky, V. Domain-Adversarial Training of Neural Networks. *J. Mach. Learn. Res.* **2016**, *17*, 1–35.
106. Zhuang, F.; Qi, Z.; Duan, K.; Xi, D.; Zhu, Y.; Zhu, H.; Xiong, H.; He, Q. A Comprehensive Survey on Transfer Learning. *Proc. IEEE* **2021**, *109*, 43–76. https://doi.org/10.1109/jproc.2020.3004555.
107. Hamilton, W.L.; Ying, R.; Leskovec, J. Inductive Representation Learning on Large Graphs. *arXiv* **2017**, arXiv:1706.02216. https://doi.org/10.48550/arxiv.1706.02216.
108. Hollmann, N.; Müller, S.; Eggensperger, K.; Hutter, F. TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. *arXiv* **2022**, arXiv:2207.01848. https://doi.org/10.48550/arxiv.2207.01848.
109. Nagler, T. Statistical Foundations of Prior-Data Fitted Networks. Proc. 40th Int. Conf. Mach. Learn. (ICML 2023) PMLR *Proc. Mach. Learn. Res.* **2023**, *202*, 25660–25676.
110. Qu, J.; Holzmüller, D.; Varoquaux, G.; Morvan, M.L. TabICL: A Tabular Foundation Model for In-Context Learning on Large Data. *arXiv* **2025**, arXiv:2502.05564. https://doi.org/10.48550/arxiv.2502.05564.
111. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215. https://doi.org/10.1038/s42256-019-0048-x.
112. Chandola, V.; Banerjee, A.; Kumar, V. Anomaly detection: A survey. *ACM Comput. Surv.* **2009**, *41*, 1–58. https://doi.org/10.1145/1541880.1541882.
113. Pang, G.; Shen, C.; Cao, L.; Hengel, A.V.D. Deep Learning for Anomaly Detection: A Review. *ACM Comput. Surv.* **2021**, *54*, 1–38. https://doi.org/10.1145/3439950.
114. Macas Carrasco, M.A.; Wu, C. Enhanced Cyber-Physical Security through Deep Learning Techniques. In Proceedings of the CPS Summer School PhD Workshop 2019. In *CEUR Workshop Proceedings*; CEUR-WS: Aachen, Germany, 2019; Volume 2457.
115. Perales Gómez, A.L. Cyberattacks Detection in Industrial Scenarios Using Machine Learning and Deep Learning Techniques. Ph.D. Thesis, Universidad de Murcia, Murcia, Spain, 2021.
116. Perales Gómez, Á.L.; Fernández Maimó, L.; Huertas Celdrán, A.; García Clemente, F.J.; Gil Pérez, M.; Martínez Pérez, G. SafeMan: A unified framework to manage cybersecurity and safety in manufacturing industry. *Softw. Pract. Exp.* **2020**, *51*, 607–627. https://doi.org/10.1002/spe.2879.
117. Gallego-Madrid, J.; Bru-Santa, I.; Ruiz-Rodenas, A.; Sanchez-Iborra, R.; Skarmeta, A. Machine learning-powered traffic processing in commodity hardware with eBPF. *Comput. Netw.* **2024**, *243*, 110295. https://doi.org/10.1016/j.comnet.2024.110295.
118. Inoue, J.; Yamagata, Y.; Chen, Y.; Poskitt, C.M.; Sun, J. Anomaly Detection for a Water Treatment System Using Unsupervised Machine Learning. In *Proceedings of the 2017 IEEE International Conference on Data Mining Workshops (ICDMW)*; IEEE: New York, NY, USA, 2017; pp. 1058–1065. https://doi.org/10.1109/icdmw.2017.149.
119. Li, D.; Chen, D.; Jin, B.; Shi, L.; Goh, J.; Ng, S.K. MAD-GAN: Multivariate Anomaly Detection for Time Series Data with Generative Adversarial Networks. In *Lecture Notes in Computer Science*; Springer International Publishing: Cham, Switzerland, 2019; pp. 703–716. https://doi.org/10.1007/978-3-030-30490-4_56.
120. Elnour, M.; Meskin, N.; Khan, K.; Jain, R. A Dual-Isolation-Forests-Based Attack Detection Framework for Industrial Control Systems. *IEEE Access* **2020**, *8*, 36639–36651. https://doi.org/10.1109/access.2020.2975066.

121. Sittón-Candanedo, I.; Alonso, R.S.; Corchado, J.M.; Rodríguez-González, S.; Casado-Vara, R. A review of edge computing reference architectures and a new global edge proposal. *Future Gener. Comput. Syst.* **2019**, *99*, 278–294. https://doi.org/10.1016/j.future.2019.04.016.
122. Mubarak, S.; Hadi Habaebi, M.; Rafiqul Islam, M.; Diyana Abdul Rahman, F.; Tahir, M. Anomaly Detection in ICS Datasets with Machine Learning Algorithms. *Comput. Syst. Sci. Eng.* **2021**, *37*, 33–46. https://doi.org/10.32604/csse.2021.014384.
123. Alwathaf, Y.; Abdel Kawy, W.A.; Aljawzi, A.A.; Eweys, O.A. Application of Integrated Urban Water Security Index (IUWSI) in Sana'a, Yemen (Case study). *Urban Water J.* **2023**, *21*, 33–44. https://doi.org/10.1080/1573062x.2023.2256722.
124. Wang, X.; Han, Y.; Leung, V.C.M.; Niyato, D.; Yan, X.; Chen, X. Convergence of Edge Computing and Deep Learning: A Comprehensive Survey. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 869–904. https://doi.org/10.1109/comst.2020.2970550.
125. Dasgupta, D.; Akhtar, Z.; Sen, S. Machine learning in cybersecurity: A comprehensive survey. *J. Def. Model. Simul. Appl. Methodol. Technol.* **2020**, *19*, 57–106. https://doi.org/10.1177/1548512920951275.
126. Zhang, J.; Pan, L.; Han, Q.L.; Chen, C.; Wen, S.; Xiang, Y. Deep Learning Based Attack Detection for Cyber-Physical System Cybersecurity: A Survey. *IEEE/CAA J. Autom. Sin.* **2022**, *9*, 377–391. https://doi.org/10.1109/jas.2021.1004261.
127. Qiu, T.; Chi, J.; Zhou, X.; Ning, Z.; Atiquzzaman, M.; Wu, D.O. Edge Computing in Industrial Internet of Things: Architecture, Advances and Challenges. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 2462–2488. https://doi.org/10.1109/comst.2020.3009103.
128. Deng, S.; Zhao, H.; Fang, W.; Yin, J.; Dustdar, S.; Zomaya, A.Y. Edge Intelligence: The Confluence of Edge Computing and Artificial Intelligence. *IEEE Internet Things J.* **2020**, *7*, 7457–7469. https://doi.org/10.1109/jiot.2020.2984887.
129. Kravchik, M.; Shabtai, A. Efficient Cyber Attack Detection in Industrial Control Systems Using Lightweight Neural Networks and PCA. *IEEE Trans. Dependable Secur. Comput.* **2022**, *19*, 2179–2197. https://doi.org/10.1109/tdsc.2021.3050101.
130. Alazab, M.; Awajan, A.; Alazzam, H.; Wedyan, M.; Alshawi, B.; Alturki, R. A Novel IDS with a Dynamic Access Control Algorithm to Detect and Defend Intrusion at IoT Nodes. *Sensors* **2024**, *24*, 2188. https://doi.org/10.3390/s24072188.
131. Kaur, B.; Shrestha, N.K.; Daggupati, P.; Rudra, R.P.; Goel, P.K.; Shukla, R.; Allataifeh, N. Water Security Assessment of the Grand River Watershed in Southwestern Ontario, Canada. *Sustainability* **2019**, *11*, 1883. https://doi.org/10.3390/su11071883.
132. Jin, M.; Wang, S.; Ma, L.; Chu, Z.; Zhang, J.Y.; Shi, X.; Chen, P.Y.; Liang, Y.; Li, Y.F.; Pan, S.; et al. Time-LLM: Time Series Forecasting by Reprogramming Large Language Models. *arXiv* **2023**, arXiv:2310.01728. https://doi.org/10.48550/arxiv.2310.01728.
133. Zhu, G.; Liu, D.; Du, Y.; You, C.; Zhang, J.; Huang, K. Toward an Intelligent Edge: Wireless Communication Meets Machine Learning. *IEEE Commun. Mag.* **2020**, *58*, 19–25. https://doi.org/10.1109/mcom.001.1900103.
134. Octavianti, T.; Staddon, C. A review of 80 assessment tools measuring water security. *WIREs Water* **2021**, *8*, e1516. https://doi.org/10.1002/wat2.1516.
135. Mubarak, S.; Hadi Habaebi, M.; Rafiqul Islam, M.; Balla, A.; Tahir, M.; A. A. Elsheikh, E.; M. Suliman, F. Industrial Datasets with ICS Testbed and Attack Detection Using Machine Learning Techniques. *Intell. Autom. Soft Comput.* **2022**, *31*, 1345–1360. https://doi.org/10.32604/iasc.2022.020801.
136. Murshed, M.G.S.; Murphy, C.; Hou, D.; Khan, N.; Ananthanarayanan, G.; Hussain, F. Machine Learning at the Network Edge: A Survey. *ACM Comput. Surv.* **2021**, *54*, 1–37. https://doi.org/10.1145/3469029.
137. Perales Gómez, Á.L.; Fernández Maimó, L.; Huertas Celdrán, A.; García Clemente, F.J. MADICS: A Methodology for Anomaly Detection in Industrial Control Systems. *Symmetry* **2020**, *12*, 1583. https://doi.org/10.3390/sym12101583.
138. Robles-Durazno, A.; Moradpoor, N.; McWhinnie, J.; Russell, G. A supervised energy monitoring-based machine learning approach for anomaly detection in a clean water supply system. In *Proceedings of the 2018 International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*; IEEE: New York, NY, USA, 2018; pp. 1–8. https://doi.org/10.1109/cybersecpods.2018.8560683.
139. Perales Gómez, Á.L.; Fernández Maimó, L.; Huertas Celdrán, A.; García Clemente, F.J.; Cadenas Sarmiento, C.; Del Canto Masa, C.J.; Méndez Nistal, R. On the Generation of Anomaly Detection Datasets in Industrial Control Systems. *IEEE Access* **2019**, *7*, 177460–177473. https://doi.org/10.1109/access.2019.2958284.
140. Junejo, K.N. Predictive safety assessment for storage tanks of water cyber physical systems using machine learning. *Sādhanā* **2020**, *45*, 61. https://doi.org/10.1007/s12046-020-1290-y.
141. Tian, Z.; Zhuo, M.; Liu, L.; Chen, J.; Zhou, S. Anomaly detection using spatial and temporal information in multivariate time series. *Sci. Rep.* **2023**, *13*, 4400. https://doi.org/10.1038/s41598-023-31193-8.
142. Chen, J.; Ran, X. Deep Learning With Edge Computing: A Review. *Proc. IEEE* **2019**, *107*, 1655–1674. https://doi.org/10.1109/jproc.2019.2921977.
143. Yoong, C.H.; Heng, J. Framework for Continuous System Security Protection in SWaT. In Proceedings of the 2019 3rd International Symposium on Computer Science and Intelligent Control, Amsterdam, Netherlands, 25–27 September 2019; pp. 1–6. https://doi.org/10.1145/3386164.3387297.
144. Mujeeb Ahmed, C.; Umer, M.A.; Binte Liyakkathali, B.S.S.; Jilani, M.T.; Zhou, J. Machine Learning for CPS Security: Applications, Challenges and Recommendations. In *Studies in Computational Intelligence*; Springer International Publishing: Cham, Switzerland, 2020; pp. 397–421. https://doi.org/10.1007/978-3-030-57024-8_18.

145. Bozdal, M.; Ileri, K.; Ozkahraman, A. Comparative analysis of dimensionality reduction techniques for cybersecurity in the SWaT dataset. *J. Supercomput.* **2023**, *80*, 1059–1079. https://doi.org/10.1007/s11227-023-05511-w.
146. Gulzar, Q.; Mustafa, K. Hybrid Cyber-Attack Detection Model on Cyber-Physical Systems Using Machine Learning Techniques. In *Lecture Notes in Networks and Systems*; Springer Nature: Singapore, 2024; pp. 197–214. https://doi.org/10.1007/978-981-99-6547-2_16.
147. Semwal, P.; Handa, A. Cyber-Attack Detection in Cyber-Physical Systems Using Supervised Machine Learning. In *Handbook of Big Data Analytics and Forensics*; Springer International Publishing: Cham, Switzerland, 2022; pp. 131–140. https://doi.org/10.1007/978-3-030-74753-4_9.
148. Raza, A.; Memon, S.; Nizamani, M.A.; Hussain Shah, M. Machine Learning-Based Security Solutions for Critical Cyber-Physical Systems. In *Proceedings of the 2022 10th International Symposium on Digital Forensics and Security (ISDFS)*; IEEE: New York, NY, USA, 2022; pp. 1–6. https://doi.org/10.1109/isdfs55398.2022.9800811.
149. Chen, X.; Polycarpou, M.M. A Ranked Contamination Diagnosis Method for Water Distribution Systems Using Spatio-Temporal Graph Neural Networks. In *Proceedings of the 2025 6th International Conference on Control and Fault-Tolerant Systems (SysTol)*; IEEE: New York, NY, USA, 2025; pp. 337–342. https://doi.org/10.1109/systol66549.2025.11267313.
150. Bonilla, C.A.; Brentan, B.; Montalvo, I.; Ayala-Cabrera, D.; Izquierdo, J. Assessing the Impacts of Failures on Monitoring Systems in Real-Time Data-Driven State Estimation Models Using GCN-LSTM for Water Distribution Networks. *Water* **2024**, *17*, 46. https://doi.org/10.3390/w17010046.
151. Li, Z.; Tian, Y.; Peng, S. Detection and localization of pipeline leakage of water distribution networks based on graph convolutional networks. *Comput.-Aided Civ. Infrastruct. Eng.* **2025**, *40*, 6656–6677. https://doi.org/10.1111/mice.70173.
152. Liu, R.; Zayed, T.; Xiao, R. Explainable Machine-Learning Leak Identification Framework for Water Distribution Networks. *J. Comput. Civ. Eng.* **2025**, *39*, 04025090. https://doi.org/10.1061/jccee5.cpeng-6119.
153. Jun, S.; Jung, D. Exploration of deep learning leak detection model across multiple smart water distribution systems: Detectable leak sizes with AMI meters. *Water Res. X* **2025**, *29*, 100332. https://doi.org/10.1016/j.wroa.2025.100332.
154. Ishikura, D.; Hara-Yamamura, H.; Igarashi, R.; Otani, K.; Yola, S.E.; Hanamoto, S.; Yamamoto-Ikemoto, R.; Yamamura, H.; Honda, R. Fluorescent fingerprints-based anomaly detection in drinking water quality and identification of contributing features by explainable AI. *Discov. Water* **2025**, *5*, 116. https://doi.org/10.1007/s43832-025-00318-7.
155. Changwang, S.; Changxi, S.; Shaowei, H.; Fuqu, P.; Zhichao, Z.; Yuxiao, Y.; Haifen, Z. Intelligent diagnosis of water supply pipe leakage based on multi-modal information fusion and improved heterogeneous temporal graph neural network. *Dev. Built Environ.* **2026**, *25*, 100825. https://doi.org/10.1016/j.dibe.2025.100825.
156. Mileiko, S.; Karim, H.; Balsamo, D. Machine Learning Approaches for Leak Detection in Water Distribution Systems: A Comparative Study. In *Proceedings of the 2025 IEEE Sensors Applications Symposium (SAS)*; IEEE: New York, NY, USA, 2025; pp. 1–6. https://doi.org/10.1109/sas65169.2025.11105198.
157. Madala, P. Predictive Maintenance in Water Treatment Plants: A Machine Learning-Based Fault Diagnosis and Optimization Approach. In *Proceedings of the 2025 6th International Conference on Electronics and Sustainable Communication Systems (ICESC)*; IEEE: New York, NY, USA, 2025; pp. 1532–1538. https://doi.org/10.1109/icesc65114.2025.11212605.
158. Arava, A.; Palleti, V.R. Real Time Leak Detection in Urban Water Networks using a Hybrid Data Driven Approach: Insights from the SWaDNet Testbed. In *Proceedings of the 2025 Eleventh Indian Control Conference (ICC)*; IEEE: New York, NY, USA, 2025; pp. 69–74. https://doi.org/10.1109/icc69100.2025.11372308.
159. Liu, X.; Gong, Z.; Zhang, X. Research on Anomaly Detection in Wastewater Treatment Systems Based on a VAE-LSTM Fusion Model. *Water* **2025**, *17*, 2842. https://doi.org/10.3390/w17192842.
160. Wu, X.; Bi, J.; Wang, G.; Wang, Z.; Li, Y.; Zhang, J.; Yuan, H.; Zhang, J.; Chang, X. Shapelet Temporal Evolution Graph Network for Water Quality Anomaly Detection. In *Proceedings of the 2025 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*; IEEE: New York, NY, USA, 2025; pp. 2827–2832. https://doi.org/10.1109/smc58881.2025.11342887.
161. Zhao, F.; Xia, W.; Zhang, Z.; Zhang, S.; Tao, L.; Wang, J.; Li, N. Unsupervised Water Pipe Network Leakage Detection Method Based on Discrete-Time Dynamic Graph Network. In *Proceedings of the 2025 IEEE 2nd International Conference on Energy and Electrical Engineering (EEE)*; IEEE: New York, NY, USA, 2025; pp. 1–5. https://doi.org/10.1109/eee64897.2025.11162765.
162. Bakhtawar, B.; Fares, A.; Zayed, T. AIoT-Driven Leak Detection in Real Water Networks Using Hydrophones. *Water Resour. Manag.* **2024**, *39*, 2551–2566. https://doi.org/10.1007/s11269-024-04077-3.
163. Sharanya, U.G.; Birabbi, K.M.; Sahana, B.; Kumar, D.M.; Sharmila, N.; Mallikarjunaswamy, S. Design and Implementation of IoT-based Water Quality and Leakage Monitoring System for Urban Water Systems Using Machine Learning Algorithms. In *Proceedings of the 2024 Second International Conference on Networks, Multimedia and Information Technology (NMITCON)*; IEEE: New York, NY, USA, 2024; pp. 1–5. https://doi.org/10.1109/nmitcon62075.2024.10698922.
164. Singh, A.; Maheshwari, A.; Singh, S. Digital Twin Framework for Leakages Detection in Large-scale Water Distribution Systems: A Case Study of IIT-Jodhpur Campus. *IFAC-PapersOnLine* **2024**, *57*, 280–285. https://doi.org/10.1016/j.ifacol.2024.05.048.

165. Raman, G.; Mathur, A. Fusing Design and Machine Learning for Anomaly Detection in Water Treatment Plants. *Electronics* **2024**, *13*, 2267. https://doi.org/10.3390/electronics13122267.
166. Cheng, M.; Li, J.; Wang, C.; Ye, C.; Chang, Z. Graph Laplace Regularization-based pressure sensor placement strategy for leak localization in the water distribution networks under joint hydraulic and topological feature spaces. *Water Res.* **2024**, *257*, 121666. https://doi.org/10.1016/j.watres.2024.121666.
167. Cui, W.; Liu, W.; Wang, T. Hierarchical Self-Attention Graph Pooling Networks for Leakage Detection of Water Distribution Networks. In *Proceedings of the 2024 IEEE 13th Data Driven Control and Learning Systems Conference (DDCLS)*; IEEE: New York, NY, USA, 2024; pp. 2012–2017. https://doi.org/10.1109/ddcls61622.2024.10606924.
168. Xu, Z.; Liu, H.; Fu, G.; Zheng, R.; Zayed, T.; Liu, S. Interpretable deep learning for acoustic leak detection in water distribution systems. *Water Res.* **2025**, *273*, 123076. https://doi.org/10.1016/j.watres.2024.123076.
169. Peng, H.; Xu, Z.; Huang, Q.; Qi, L.; Wang, H. Leakage Detection in Water Distribution Systems Based on Logarithmic Spectrogram CNN for Continuous Monitoring. *J. Water Resour. Plan. Manag.* **2024**, *150*, 04024015. https://doi.org/10.1061/jwrmd5.wreng-6276.
170. Rolle, R.P.; Rodrigues, W.C.; Tomazini, L.R.; Monteiro, L.N.; Godoy, E.P. Leveraging graph-based leak localization in water distribution networks. In *Proceedings of the 2024 IEEE International Workshop on Metrology for Industry 4.0 & IoT (MetroInd4.0 & IoT).*; IEEE: New York, NY, USA, 2024; pp. 192–197. https://doi.org/10.1109/metroind4.0iot61288.2024.10584129.
171. Park, H.; Oh, Y.; Kim, K.; Kim, T.; Hyung, J.; Koo, J. Optimal leakage detection and classification of the water distribution network based on the machine learning approach. *Water Supply* **2024**, *24*, 2158–2171. https://doi.org/10.2166/ws.2024.122.
172. Mu, T.; Zhang, C.; Huang, M.; Ning, B.; Wang, J. Partitioning Leakage Detection in Water Distribution Systems: A Specialized Deep Learning Framework Enhanced by Spatial–Temporal Graph Convolutional Networks. *ACS ES&T Water* **2024**, *4*, 3453–3463. https://doi.org/10.1021/acsestwater.4c00285.
173. Syed, T.A.; Muhammad, M.A.; AlShahrani, A.A.; Hammad, M.; Naqash, M.T. Smart Water Management with Digital Twins and Multimodal Transformers: A Predictive Approach to Usage and Leakage Detection. *Water* **2024**, *16*, 3410. https://doi.org/10.3390/w16233410.
174. Wu, W.; Pan, X.; Kang, Y.; Xu, Y.; Han, L. To Feel the Spatial: Graph Neural Network-Based Method for Leakage Risk Assessment in Water Distribution Networks. *Water* **2024**, *16*, 2017. https://doi.org/10.3390/w16142017.
175. Bi, J.; Yuan, M.; Yuan, H.; Wang, Z.; Qiao, J. Water Quality Anomaly Detection with Dual Sliding Windows and Convolutional LSTM. In *Proceedings of the 2024 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*; IEEE: New York, NY, USA, 2024; pp. 3559–3564. https://doi.org/10.1109/smc54092.2024.10831202.
176. Mücke, N.T.; Pandey, P.; Jain, S.; Bohté, S.M.; Oosterlee, C.W. A Probabilistic Digital Twin for Leak Localization in Water Distribution Networks Using Generative Deep Learning. *Sensors* **2023**, *23*, 6179. https://doi.org/10.3390/s23136179.
177. Brahmbhatt, P.; Maheshwari, A.; Gudi, R.D. Digital twin assisted decision support system for quality regulation and leak localization task in large-scale water distribution networks. *Digit. Chem. Eng.* **2023**, *9*, 100127. https://doi.org/10.1016/j.dche.2023.100127.
178. Chen, S.; Wang, Y.; Zhang, W.; Zhang, H.; He, Y. Leak Detection in Water Supply Network Using a Data-Driven Improved Graph Convolutional Network. *IEEE Access* **2023**, *11*, 117240–117249. https://doi.org/10.1109/access.2023.3326470.
179. Moradpoor, N.; Maglaras, L.; Abah, E.; Robles-Durazno, A. The Threat of Adversarial Attacks Against Machine Learning-based Anomaly Detection Approach in a Clean Water Treatment System. In *Proceedings of the 2023 19th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*; IEEE: New York, NY, USA, 2023; pp. 453–460. https://doi.org/10.1109/dcoss-iot58021.2023.00077.
180. Chen, X.; Li, L.; Peng, X.; Yang, D. A Fault Diagnosis Method Based on Canonical Correlation Analysis and Transfer Learning in Wastewater Treatment Plant. In *Proceedings of the 2022 37th Youth Academic Annual Conference of Chinese Association of Automation (YAC)*; IEEE: New York, NY, USA, 2022; pp. 908–913. https://doi.org/10.1109/yac57282.2022.10023760.
181. Zhang, C.; Alexander, B.J.; Stephens, M.L.; Lambert, M.F.; Gong, J. A convolutional neural network for pipe crack and leak detection in smart water network. *Struct. Health Monit.* **2022**, *22*, 232–244. https://doi.org/10.1177/14759217221080198.
182. Aggarwal, S.; Sehgal, S. Automation of Water Distribution System by Prediction of Water Consumption and Leakage Detection Using Machine Learning and IoT. In *Proceedings of the 2022 12th International Conference on Cloud Computing, Data Science & Engineering (Confluence)*; IEEE: New York, NY, USA, 2022; pp. 345–350. https://doi.org/10.1109/confluence52989.2022.9734192.
183. Örn Garðarsson, G.; Boem, F.; Toni, L. Graph-Based Learning for Leak Detection and Localisation in Water Distribution Networks*. *IFAC-PapersOnLine* **2022**, *55*, 661–666. https://doi.org/10.1016/j.ifacol.2022.07.203.
184. Parvathy, A.; Leela Kasyap, G.; Venkata Abhinav, D.; Surya Sai, A.; Sriranjani, R.; Hemavathi, N. Hybrid Machine Learning based False Data Injection Attack Detection and Mitigation Model for Waste Water Treatment Plant. In *Proceedings of the 2022 International Conference on Automation, Computing and Renewable Systems (ICACRS)*; IEEE: New York, NY, USA, 2022; pp. 674–680. https://doi.org/10.1109/icacrs55517.2022.10029198.

185. Janvier, N.; Gilbert, N.; J.M.V, M.; Clarisse, K.; Sophonie, H. IOT Enabled Enhanced Method for Water Leakage Detection and Localization in Water Distribution Systems. In *Proceedings of the 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N)*; IEEE: New York, NY, USA, 2022; pp. 1380–1384. https://doi.org/10.1109/icac3n56670.2022.10074125.
186. Zhang, W.; Yang, X.; Li, J. Sensor Placement for Leak Localization in Water Distribution Networks Based on Graph Convolutional Network. *IEEE Sens. J.* **2022**, *22*, 21093–21100. https://doi.org/10.1109/jsen.2022.3208415.

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