Journal of Science, Innovation and Creativity
Vol. 5 No. 2 (2026): ISSN (Online): 2958-454X
DOI: https://doi.org/10.58721/0tt4fn47
CAUS as a Data-Quality Layer for IoT Flood Monitoring in Uganda: Design and Preliminary Simulation

CAUS as a Data-Quality Layer for IoT Flood Monitoring in Uganda: Design and Preliminary Simulation

Kosea Erasto Muwanga, Daniel Otanga & Rupra Singh Satwinder
Masinde Muliro University of Science and Technology, Kenya

Article History
Received: 2026-04-22
Revised: 2026-09-04
Accepted: 2026-09-16
Published: 2026-11-10

Keywords
Confidence scoring
Data
Flood monitoring
Internet of Things

How to cite:
Muwanga, K. E., Otanga, D., & Satwinder, R. S.(2026). CAUS as a Data-Quality Layer for IoT Flood Monitoring in Uganda: Design and Preliminary Simulation. *Journal Science, Innovation and Creativity, 5*(2), 396-406.

Copyright © 2026

Abstract
This article examines the Context Aggregation, Validation and Unification System (CAUS) as a data-quality layer in an Internet of Things (IoT) flood monitoring architecture for data-scarce settings in Uganda. Guided by design-science research and ITIL 4 service-management principles, it synthesises a review of IoT flood frameworks, CHIRPS rainfall analysis for Bwaise, Kalerwe and Mbarara (2010–2024), available Kampala-area ERA5 Ag comparisons, and a cloud-hosted proof of concept using synthetic sensor observations. The rainfall dataset contained 5,479 daily records per location. Bwaise and Kalerwe had identical extracted CHIRPS series; Mbarara recorded a lower mean daily rainfall (2.58 mm) but a higher maximum (53.49 mm). Daily CHIRPS–ERA5 Ag correlations were approximately 0.40–0.43 and monthly correlations 0.75–0.77 for the available Kampala-area comparison. The recommends a meythodological approach to reproducible CAUS thresholds, fusion weights, imputation settings or a confidence equation. This article therefore distinguishes reported implementation from a proposed, parameterised specification covering validation, source-aware fusion, short-gap imputation, calibration metadata and a weighted quality index. The simulation demonstrated workflow execution, not field accuracy or calibrated confidence. CAUS is presented as a design and testing blueprint; controlled fault-injection, independent sensor calibration and field validation are required before operational warning use.

Introduction

Flood early-warning systems depend on timely observations of rainfall, water-level and the local conditions that shape runoff. In urban settlements, the value of these observations is conditioned by drainage capacity, land cover, topography and the speed at which water accumulates. The challenge is especially important in East African settings where monitoring coverage, maintenance resources, communication infrastructure and access to quality-controlled ground observations may be uneven. Regional climate and disaster-risk assessments have repeatedly emphasised the need to strengthen hydrometeorological services and people-centred early-warning capacity (World Meteorological Organisation [WMO], 2024). In Uganda, urban flood exposure is compounded by rapid settlement growth and drainage constraints; Kampala's low-lying neighbourhoods, including Bwaise and Kalerwe, have been identified as flood-prone settings in the underlying study (Abebe, 2013; Lwasa et al., 2014).

Internet of Things (IoT) architectures can support continuous and geographically distributed monitoring through connected sensors, communication networks and cloud analytics (Atzori et al., 2010; Gubbi et al., 2013). Recent reviews of flood early-warning systems emphasise that effective systems combine monitoring technology with forecasting, institutional coordination, preparedness and accessible risk communication; sensor deployment alone does not constitute an

effective warning service (Tao et al., 2024; Painter et al., 2025; Thinh, 2025). IoT-based flood monitoring reviews also identify sensor selection, network reliability, edge processing and multi-channel dissemination as interdependent design concerns, particularly in developing-country settings (Pahuriray & Cerna, 2025). However, connectivity does not guarantee trustworthy observations. Sensors may produce missing, delayed, duplicated, drifting or implausible values; communication failures can create apparent sensor outages; and external gridded datasets may disagree because they use different inputs, spatial resolutions and estimation methods. Recent IoT data-quality research has proposed real-time anomaly screening for heterogeneous infrastructures and highlighted the importance of sensor placement, calibration, validation, fusion and uncertainty management in environmental networks (Arnau Muñoz et al., 2024; Okafor et al., 2024).

Several method families are relevant to this problem. Rule-based checks are transparent and computationally inexpensive, but fixed limits may miss context-dependent faults or adapt poorly to changing conditions. Statistical and machine-learning anomaly detectors can identify more complex patterns, but their performance depends on representative training data, labelled faults and the handling of concept drift. Recent work on IoT data-quality assurance demonstrates the use of unsupervised anomaly detection to screen heterogeneous sensor packets, while environmental-network studies identify calibration, validation, fusion and data-gap management as essential quality controls (Arnau Muñoz et al., 2024; Okafor et al., 2024). Missing-data methods range from simple interpolation and carry-forward approaches to model-based and multivariate imputation; their suitability depends on the missingness mechanism, temporal structure and the risk of smoothing extremes (Zhang & Thorburn, 2022; Chatterjee & Ahmed, 2022). In flood monitoring, an imputed value must remain distinguishable from a direct observation because the difference may matter during an extreme event.

Confidence and uncertainty information are also important. A quality score can help downstream users distinguish complete, timely and plausible records from records affected by missingness or processing. Yet a numerical score is not automatically a probability. A probabilistic interpretation requires comparison with independent labels and evidence of calibration. Design-science research similarly requires an explicit account of the artefact, its intended utility and the evaluation evidence supporting claims about it (Hevner et al., 2004; Gregor & Hevner, 2013; Peffers et al., 2007). Venable et al. (2016) further distinguish artificial evaluation in controlled settings from naturalistic evaluation in operational contexts.

This article examines CAUS (Context Aggregation, Validation and Unification System), designed as the data-quality and integration layer within a wider IoT–analytics–GIS–alerting architecture. The research question is: **How can a context-aware data-quality layer be positioned within an IoT flood-monitoring architecture to preserve the provenance and reliability status of environmental observations in data-scarce settings?** The gap addressed is not the absence of anomaly detection algorithms generally, but the limited articulation in the study's reviewed framework set of how observation-level quality, source context, missing-value status and service-management processes should be carried together into downstream analytics.

The article's contribution is therefore bounded: it presents a thesis-grounded architecture and preliminary software demonstration of CAUS, and identifies the evidence needed for a valid technical evaluation. It does not claim a new fault-detection algorithm or verified fault-recovery performance. The remainder of the article describes the research design, CAUS specification as supported by the source study, rainfall and proof of concept results, and implications for future validation.

Research context and conceptual grounding

Flood monitoring and data constraints in Uganda

The underlying study considered Bwaise and Kalerwe in Kampala and Mbarara as a comparative location. It used gridded precipitation products because a consistent, accessible set of coordinate-specific gauge observations was not available in the research dataset. CHIRPS combines satellite infrared information with station observations and provides a long daily record that is useful for regional rainfall analysis (Funk et al., 2015). ERA5 is a global atmospheric reanalysis that combines model forecasts and observations through data assimilation; it is a model-based reanalysis product rather than a local rain gauge (Hersbach et al., 2020). Agreement between these products can inform contextual uncertainty, but cannot substitute for sensor fault labels or gauge validation.

Sensor quality, anomaly detection and imputation

The sensor-data literature commonly distinguishes completeness, validity, consistency, accuracy and timeliness. A value may be within a physical range but still be stale; a value may be statistically unusual but represent a genuine extreme. Consequently, quality checks should be interpreted with sensor type, location, sampling interval and environmental context in view. Rule-based range and rate checks offer interpretability, while temporal or multivariate anomaly detection may identify patterns that fixed rules cannot. Recent environmental IoT studies explicitly identify sensor placement, calibration, data validation, fusion, environmental interference and uncertainty as connected determinants of data-quality (Okafor et al., 2024). Machine-learning approaches can supplement, but should not replace, transparent physical and operational checks unless they are evaluated using representative labelled data (Arnau Muñoz et al., 2024).

Imputation can preserve continuity for downstream analysis, but should not overwrite the original missingness or imply that an estimate was measured. Environmental-monitoring research has compared a range of near-real-time imputation approaches and shown that method choice depends on the data and monitoring task (Zhang & Thorburn, 2022). The CAUS design therefore treats imputation as a controlled data-handling function, not as proof that the underlying sensor fault has been repaired.

Design-science and service-management

The source study adopted Design-Science Research Methodology (DSRM), which structures work around problem identification, objectives, design and development, demonstration, evaluation and communication (Peffers et al., 2007). The artefact is CAUS within the wider flood monitoring architecture. Its evaluation should separate technical functionality from utility and operational effectiveness. The FEDS framework recommends aligning evaluation purpose and setting with the claims being made (Venable et al., 2016).

The architecture also draws on ITIL-aligned service-management. Service-level monitoring, monitoring and event management, incident management, problem management, risk management and continual improvement provide an operational route for recording disruptions, identifying recurring causes and revising service controls. These practices are governance mechanisms; they do not themselves validate sensor readings. Recent IT service-management studies describe the use of ITIL-aligned continual improvement and Plan–Do–Check–Act cycles to improve service quality and operational control (Baptista & Barata, 2024). ITIL 4 frames continual improvement as an iterative process that links improvement activity to organisational goals (AXELOS, 2019).

Recent developments in flood-warning and sensor calibration research

Recent flood early-warning literature increasingly treats the warning-chain as a socio-technical service: hazard observation and forecasting must be connected to institutional responsibility, impact-based messages, preparedness and community response. Painter et al. (2025) identify timeliness, human capacity and system integration as recurring characteristics of effective flood early-warning systems, while also noting gaps in participatory science, bottom-up approaches and

preparedness financing. A 2025 review of urban flood resilience similarly highlights the growing role of real-time IoT and remote-sensing observations, while identifying fragmented coordination, model generalisability and inequitable risk communication as continuing barriers (Thinh, 2025). For CAUS, this means that data-quality outputs should be traceable and interpretable by the people and agencies responsible for deciding whether to issue or escalate a warning.

Calibration is equally important because a sensor that passes a plausibility check may still be biased or drifting. Okafor et al. (2024) discuss calibration and validation as part of environmental IoT data-quality management. Hossain et al. (2025) investigate drift detection and calibration using a variational-autoencoder approach, illustrating that calibration status and drift should be managed over the operational life of a sensor rather than treated as a one-time installation activity. These studies support recording calibration dates, reference instruments, correction coefficients, uncertainty and maintenance history as part of CAUS provenance. They do not provide universal thresholds for rainfall or water-level sensors; such limits must be derived from the selected instrument specifications and local field conditions.

ITIL is relevant at the service governance layer rather than as a sensor-validation algorithm. In a flood monitoring service, monitoring and event management should identify deviations in sensor availability, data freshness and gateway operation; incident management should coordinate response to outages; problem management should investigate recurring causes such as water ingress, power loss or drift; and continual improvement should use incident records and post-event reviews to revise thresholds, maintenance schedules and communication procedures. Recent ITSM case studies describe continual improvement and PDCA as mechanisms for making service controls measurable and repeatable (Baptista & Barata, 2024)

Method

Research design and evidence streams

This article is a secondary scholarly account of the completed doctoral study, not a new field experiment. Five evidence streams informed the account: (1) a review of eight IoT-based flood-prediction frameworks against ITIL-related service-management domains; (2) descriptive analysis of CHIRPS daily rainfall for Bwaise, Kalerwe and Mbarara; (3) comparison of CHIRPS with available Kampala-area ERA5 Ag precipitation series; (4) design of the wider resilient IoT architecture and CAUS; and (5) a cloud-hosted proof of concept using synthetic environmental observations.

The framework review identified the extent to which the selected systems described service governance, data-quality management, confidence information, incident response and continual improvement. The article does not reproduce a complete item-by-item scoring matrix because the manuscript and extracted source material do not provide a sufficiently auditable scoring table and criterion-level results. Accordingly, the article does not use that review as quantitative evidence that all existing frameworks lack a particular capability. It is treated as a design input that informed the proposed requirements.

DSRM activities and evaluation criteria

Table 1: Mapping of design-science activities to the study

DSRM activity	Application in this study	Evidence and evaluation boundary
Problem identification and motivation	Sensor unreliability, fragmented local observations and data-quality risks in flood monitoring were identified.	Literature and framework review; establishes the problem context, not CAUS performance.
Define objectives	Specify an architecture that aggregates context, validates observations, preserves missingness and communicates quality information.	Design requirements; no independent test of each requirement reported.
Design and development	Develop the layered IoT framework and conceptual CAUS functions.	Architecture and proof of concept workflow.
Demonstration	Run synthetic rainfall, water-level, humidity and temperature observations through a cloud-hosted workflow.	Demonstrates workflow execution under configured simulation conditions.
Evaluation	Review reported system-level indicators and rainfall-product comparisons against their actual evidential scope.	No labelled fault metrics, baseline comparison, imputation error or calibration reported.
Communication	Document the architecture, findings, limitations and next validation requirements.	This article is a thesis-derived communication of the artefact and its evidence.

The appropriate evaluation criteria for a future CAUS-specific test are fault-class precision, recall and F1; false-alarm rate; imputation error on withheld observations; calibration or discrimination of confidence scores; processing latency and throughput; and resource use on representative edge hardware. These are proposed criteria, not completed results in the present study.

Study locations and rainfall data

The CHIRPS analysis covered 1 January 2010 to 31 December 2024 and contained 5,479 daily observations for each of the three locations. ERA5 Ag v2.0 Daily was used as a complementary gridded product where an available Kampala-area series permitted comparison. The study dataset did not contain equivalent coordinate-specific ERA5 Ag coverage for Mbarara. The available comparison therefore refers to the Kampala-area series only; it is not a three-site comparison. Bwaise and Kalerwe had identical extracted CHIRPS series, so they do not constitute two independent rainfall time series in this analysis.

CAUS architecture and operation

CAUS is positioned between IoT sensing/edge acquisition and cloud analytics. The conceptual architecture (Figure 1) contains sensing, edge/gateway, CAUS, analytics, GIS/context, and alerting layers. ITIL-aligned governance and continual improvement processes provide feedback across the pipeline.

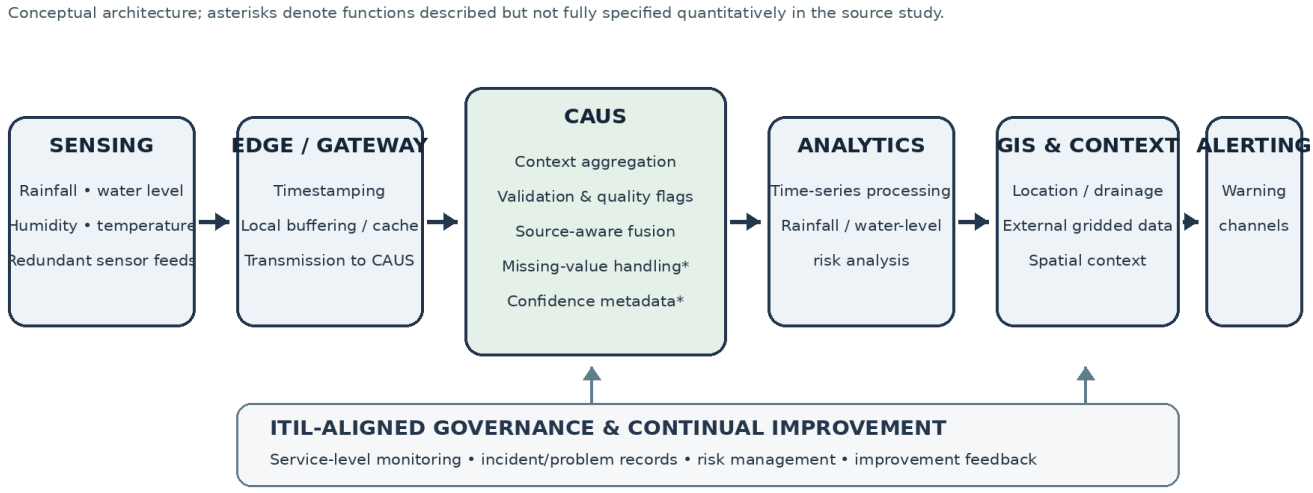


Figure 1: Conceptual architecture of the proposed IoT flood monitoring framework and the position of CAUS

The source study describes five CAUS activities: (1) context aggregation, combining observations with sensor identity, time, location and relevant external data; (2) validation, assessing completeness, plausibility and consistency; (3) data fusion, using multiple available sources while retaining source identity; (4) missing-value handling, reconstructing selected gaps from historical or contextual data where the implemented method permits; and (5) confidence assignment, attaching a quality indicator for downstream interpretation.

Reproducibility of the CAUS approach is possible through a researcher-guided implementation procedure, using the proposed architecture and methodological framework as a reference. The researcher guides the implementation of the key CAUS processes, namely context aggregation, data validation, source integration, missing-value handling and confidence assignment. This guidance provides a consistent basis for subsequent researchers to follow and replicate the CAUS workflow. To support transparency and consistency, the implementation procedures, parameters, assumptions and configurations adopted should be clearly documented. This will enable independent researchers to reproduce the approach, compare implementation outcomes and evaluate its performance under similar experimental conditions.

The reported mean confidence score of 0.88 is treated as a simulation output, not a probability of correctness. The study does not document the calculation sufficiently to reproduce it or establish calibration. Likewise, the target baseline of 0.85 is reported as a design target. It is therefore not interpreted here as a validated acceptance threshold.

Proof of concept implementation

The proof of concept used Python 3.11 and GitHub Actions. Synthetic rainfall, water-level, humidity and temperature observations were generated, with selected missing-data, variability and communication delays introduced. The CAUS workflow performed data-quality checking, fusion, missing-value handling and confidence assignment before downstream analytical processing. The recorded outputs included simulation_log.csv and latest_summary.md. The source study does not report the number of independent runs, random seeds, sensor count, sampling interval, exact missingness rate or pattern, noise distributions, delay distribution, or a controlled no-CAUS/conventional-validation baseline. Those omissions prevent independent replication of the reported performance indicators.

Rainfall analysis

Daily and monthly CHIRPS and ERA5 Ag series were temporally harmonised and compared using Pearson correlation. The analysis was exploratory. It did not include rank correlation, bias, MAE, RMSE, probability of detection, false-alarm ratio, critical success index or confidence intervals. Nor did it use an independent gauge dataset. These metrics are therefore not added retrospectively. The 25 mm/day and 60 mm/three-day thresholds are retained only as exploratory indicators used by the source study; they were not independently validated against an event-based flood-impact catalogue and should not be treated as universal warning thresholds.

Results

Rainfall characteristics

The extracted CHIRPS dataset contained 5,479 daily records for each site. Bwaise and Kalerwe had identical descriptive statistics, while Mbarara had a lower mean and a higher maximum daily rainfall. The identical urban series limit the independence of site-level comparison and may reflect shared gridded representation or extraction. They should not be interpreted as evidence that the two neighbourhoods experience identical local rainfall.

Table 2: Descriptive statistics of daily CHIRPS rainfall, 2010–2024

Location	Daily records	Mean (mm)	SD (mm)	Minimum (mm)	75th percentile (mm)	Maximum (mm)
Bwaise	5,479	3.49	5.18	0.00	5.71	45.67
Kalerwe	5,479	3.49	5.18	0.00	5.71	45.67
Mbarara	5,479	2.58	5.21	0.00	3.57	53.49

Rainfall threshold indicators

Using the exploratory thresholds in the source study, daily rainfall greater than 25 mm occurred 24 times in Bwaise, 24 times in Kalerwe and 38 times in Mbarara. Three-day cumulative rainfall greater than 60 mm was reported three times in each Kampala location and four times in Mbarara. These are threshold exceedances, not verified flood events. The analysis did not establish the sensitivity or specificity of the thresholds against observed flood impacts.

CHIRPS–ERA5 Ag comparison

For the available Kampala-area series, daily Pearson correlations were approximately 0.40–0.43 and monthly correlations approximately 0.75–0.77. The higher monthly association indicates stronger correspondence in broad temporal variation after aggregation. It does not establish agreement in rainfall amount, accuracy against gauges or the superiority of either product. No equivalent comparison was available for Mbarara.

Software proof of concept results

Table 3: Reported system-level indicators from the synthetic proof of concept

Indicator	Reported value	Interpretation and limitation
Simulated system uptime	99.6%	Availability in the configured cloud workflow; not field-system uptime.
Mean confidence score	0.88	Simulation score; formula and calibration are not documented.
Severity-pattern match	92%	Match to expected patterns in synthetic data; not independently labelled flood-prediction accuracy.
Data processed per cycle	Approximately 1,000 points	Reported simulation volume; no throughput benchmark or hardware profile provided.
Mean detection time	Approximately 1.8 minutes	System-level simulated timing; event definition and fault-specific detection are not documented.
Mean repair/resolve time	Less than 2 minutes	Reported workflow timing; operational repair events and denominator are not specified.

The proof of concept demonstrates that an automated workflow could generate synthetic observations, process them through the integrated software pipeline and produce summary artefacts. It does not establish CAUS-specific precision, recall, F1, false-alarm rate, imputation error or confidence calibration. The 92% synthetic severity-pattern match is not a CAUS fault-detection metric and should not be used to claim field predictive accuracy.

Summary of evidence

The results support three limited findings. First, CAUS was represented as a functional component in an automated proof of concept pipeline. Second, the rainfall analysis showed differences in descriptive rainfall characteristics and stronger cross-product temporal association at monthly than daily scale for the available Kampala-area comparison. Third, the simulation produced system-level indicators under synthetic conditions. None of these findings independently verifies CAUS fault-detection or recovery effectiveness.

Discussion

What the study contributes

The study's contribution is an architectural specification that places data-quality and provenance handling between sensing and predictive analytics. This responds to a practical weakness in many monitoring pipelines: the downstream model may receive values without sufficient information about whether they were observed, delayed, imputed or drawn from a secondary data source. Existing surveys show that IoT sensor verification includes anomaly classification, detection, correction and evaluation, but these tasks require explicit methods and test data (Teh et al., 2020; Nguyen et al. (2023)). CAUS brings these concerns together with environmental context and service-management feedback. Its distinctiveness in this study is the proposed integration and positioning, not a demonstrated new detection algorithm.

The research question is answered at the design-level: CAUS can be positioned as an intermediate data-quality layer that preserves context and reliability metadata before analytics. The hypothesis that this design improves fault-detection or recovery performance is **not tested** by the reported evidence and therefore cannot be accepted or rejected. The architectural proposition is supported by the successful execution of the software workflow, but operational effectiveness remains unestablished.

Interpreting confidence and imputation

The mean confidence score of 0.88 is not sufficient evidence that CAUS reliably distinguishes correct from faulty observations. A confidence index becomes useful when its components, interpretation and thresholds are documented. If intended as a probability, it must be calibrated using independent labelled records and evaluated with appropriate measures, such as reliability diagrams or Brier score. If it is a heuristic index, it should be named and interpreted as such. The

source study does not provide the formula or weights, so the present article does not reconstruct them.

Imputation should likewise be treated as continuity support rather than physical sensor repair. The original missingness, imputation flag, method, time window and source observations should remain available. During extreme rainfall or rapidly rising water levels, a system should avoid presenting estimated values as equivalent to direct observations. This distinction is consistent with environmental-monitoring work on near-real-time missing-data handling, which emphasises method choice and evaluation against known observations (Handling Missing-Data in Near Real-Time Environmental-Monitoring, 2022).

Implications for Ugandan stakeholders

For local governments and disaster-management agencies, the immediate implication is to specify data governance alongside sensor procurement. A pilot should maintain sensor identifiers, calibration and maintenance histories, timestamp synchronisation, communication logs and explicit quality flags. National meteorological and hydrological agencies should document the provenance, spatial coverage and update characteristics of external datasets. Universities and technology partners can support controlled fault-injection tests and independent validation against calibrated gauges and water-level sensors. Community alert protocols should also distinguish a high hazard estimate based on direct observations from one materially dependent on imputed or uncertain inputs. This emphasis on sensor calibration and traceable quality metadata is consistent with recent environmental IoT monitoring research (Okafor et al., 2024), while the wider warning-chain perspective reflects recent reviews of flood early-warning system design (Painter et al., 2025; Thinh, 2025).

For resource-constrained deployments, local buffering and delayed synchronisation can protect data continuity during connectivity interruptions, but these functions require testing on the actual gateway hardware. The current cloud-hosted proof of concept did not measure energy use, memory pressure, processor load or edge latency. The framework should therefore be treated as a design candidate for local adaptation rather than a proven low-resource deployment.

Limitations

The evidence base has several limitations. The rainfall analysis used gridded CHIRPS and ERA5 Ag products rather than a local gauge network; the ERA5 Ag comparison was limited to available Kampala-area series; and Bwaise and Kalerwe had identical extracted CHIRPS time series. Pearson correlation alone does not characterise agreement or event detection for zero-inflated, skewed rainfall. The rainfall thresholds were not validated against observed flood impacts.

The proof of concept used synthetic observations in GitHub Actions and did not report enough detail about simulation runs, random seeds, sensor counts, sampling intervals, noise, missingness and delay distributions to reproduce the experiment. It did not include a no-CAUS or conventional-validation baseline, labelled fault scenarios, per-fault confusion matrices, withheld-value imputation tests or confidence calibration. The confidence equation and its 0.85 target rationale are not documented in the available source material. The framework review also lacks an auditable criterion-level scoring summary in the article materials. These limitations constrain claims about novelty, comparative performance and operational readiness.

Priority validation programme

A focused next phase should: (1) publish the complete CAUS implementation specification, including validation rules, fusion/conflict handling, imputation windows and confidence score definition; (2) conduct repeatable fault-injection experiments covering missing, delayed, duplicated, stuck, drifting and implausible readings, with documented seeds and baselines; (3) report per-fault precision, recall, F1 and false alarms, and assess imputation error on withheld observations; (4) test confidence discrimination and calibration against independent labels, while profiling latency and resource use on representative edge devices; and (5) undertake a field pilot

using calibrated rainfall and water-level sensors, maintenance logs and observed flood-impact records. Code, configuration, simulation logs and a DOI-backed repository should accompany any subsequent submission once the authors have prepared and deposited those materials.

Conclusion

This article has examined CAUS as the data-quality and integration layer in a proposed IoT-based flood monitoring framework for Uganda. The design positions CAUS between sensing and analytics and assigns it functions for context aggregation, validation, source-aware fusion, missing-value handling and confidence metadata. The study also demonstrates an automated cloud-hosted workflow using synthetic environmental observations.

The evidence does not establish fault-specific detection or recovery performance. The rainfall results are based on gridded products, and the proof of concept indicators are system-level simulation outputs rather than field measurements. The defensible contribution is therefore a design-level integration and a preliminary software demonstration, accompanied by a clear set of requirements for reproducible technical evaluation. Claims of improved fault-detection, calibrated confidence, field resilience or operational flood-warning accuracy should await controlled baseline testing and field validation.

References

Abebe, A. J. (2013). Uncertainty in urban flood modeling: A case study of Bwaise, Kampala. Journal of Flood Risk Management, 6(3), 210–224.

Arnau Muñoz, L., Berná Martínez, J. V., Maciá Pérez, F., & Lorenzo Fonseca, I. (2024). Anomaly detection system for data quality assurance in IoT infrastructures based on machine learning. Internet of Things, 25, Article 101095. https://doi.org/10.1016/j.iot.2024.101095

Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. Computer Networks, 54(15), 2787–2805. https://doi.org/10.1016/j.comnet.2010.05.010

AXELOS. (2019). ITIL foundation: ITIL 4 edition. TSO (The Stationery Office).

Baptista, B., & Barata, J. (2024). Continuously improving IT service management in the pharmaceutical industry. Procedia Computer Science, 239, 923–930. https://doi.org/10.1016/j.procs.2024.06.253

Chatterjee, A., & Ahmed, B. S. (2022). IoT anomaly detection methods and applications: A survey. Internet of Things, 19, Article 100568. https://doi.org/10.1016/j.iot.2022.100568

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Rowland, J., Harrison, L., Hoell, A., & Michaelsen, J. (2015). The climate hazards infrared precipitation with stations—A new environmental record for monitoring extremes. Scientific Data, 2, Article 150066. https://doi.org/10.1038/sdata.2015.66

Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. MIS Quarterly, 37(2), 337–355.

Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. Future Generation Computer Systems, 29(7), 1645–1660. https://doi.org/10.1016/j.future.2013.01.010

Hersbach, H., Bell, B., Berrisford, P., et al. (2020). The ERA5 global reanalysis. Quarterly Journal of the Royal Meteorological Society, 146(730), 1999–2049. https://doi.org/10.1002/qj.3803

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105. https://doi.org/10.2307/25148625

Hossain, M. K., Ahmad, I., Habibi, D., & Waqas, M. (2025). Enhancing IoT sensors precision through sensor drift calibration with variational autoencoder. IEEE Internet of Things Journal, 12(7), 8421–8437. https://doi.org/10.1109/JIOT.2024.3503616

Lwasa, S., Koojo, C., Mabasi, J., & Bwambale, B. (2014). Climate change adaptation and urban flood risk management in Kampala, Uganda. *Environment and Urbanization, 26*(2), 431–448.

Nguyen, L. A., Pham, T. K., Lee, S., Yeo, H., & Son, Y. (2023). Comprehensive survey of sensor data verification in Internet of Things. *IEEE Access, 11*, 50560–50577.

https://doi.org/10.1109/ACCESS.2023.3277545

Okafor, N., Ingle, R., Matthew, U., Saunders, M., & Delaney, D. (2024). Assessing and improving IoT sensor data quality in environmental monitoring networks: A focus on peatlands. *IEEE Internet of Things Journal, 11*(24), 40727–40742. https://doi.org/10.1109/JIOT.2024.3454241

Painter, P., Semmens, K., Maxfield, K., Cattoën, C., & Carr, R. H. (2025). Designing effective flood early warning systems: A review of barriers, best practices, and key characteristics. *Journal of Flood Risk Management, 18*(4), Article e70145. https://doi.org/10.1111/jfr3.70145

Pahuriray, A. V., & Cerna, P. D. (2025). IoT-enabled flood monitoring and early warning systems: A systematic review. *International Journal of Computer Science and Mobile Computing, 14*(4). https://doi.org/10.47760/ijcsmc.2025.v14i04.005

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Tao, Y., Tian, B., Adhikari, B. R., Zuo, Q., Luo, X., & Di, B. (2024). A review of cutting-edge sensor technologies for improved flood monitoring and damage assessment. *Sensors, 24*(21), Article 7090. https://doi.org/10.3390/s24217090

Teh, H. Y., Kempa-Liehr, A. W., & Wang, K. I.-K. (2020). Sensor data quality: A systematic review. *Journal of Big Data, 7*, Article 11. https://doi.org/10.1186/s40537-020-0285-1

Thinh, P. N. (2025). Early warning systems for urban flood resilience. *Reviews in Agricultural Science, 13*(4), 61–74. https://doi.org/10.7831/ras.13.4_61

Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: A framework for evaluation in design science research. *European Journal of Information Systems, 25*(1), 77–89. https://doi.org/10.1057/ejis.2014.36

World Meteorological Organisation. (2024). State of the climate in Africa 2023. World Meteorological Organisation.

Zhang, Y., & Thorburn, P. J. (2022). Handling missing data in near-real-time environmental monitoring: A system and a review of selected methods. *Future Generation Computer Systems*, 128, 63–72. https://doi.org/10.1016/j.future.2021.09.033