International Journal of Engineering and Information Systems (IJEAIS)
ISSN: 2643-640X
Vol. 10 Issue 10, October - 2026, Pages: 75-81

Federated and Explainable Edge–Fog Intelligence for Continuous Glucose Monitoring and Ambient Assisted Care in Diabetes: A Review

Dr. Virendra Tank

Shri Mahaveer College, Jaipur, Rajasthan, India
vtank87@gmail.com 0009-0003-9126-0982

Abstract—*Diabetes mellitus requires lifelong, data-driven self-management, and the sensors that make this possible (continuous glucose monitors, wearables and home ambient sensors) now produce a continuous stream of sensitive data. Sending all of it to the cloud creates delays, privacy exposure and bandwidth costs, while purely on-device processing is limited by memory and energy. This review examines how a layered edge–fog–cloud design, combined with personalized machine learning, federated learning and explainable AI, can support glucose forecasting, diabetes risk prediction and ambient assisted care at home. It formulates the main technical problems, organizes the methods into a layered taxonomy, maps each problem to its solution, consolidates evidence reported in the literature, and discusses open issues such as non-IID data, clinical validation, safety, energy use and regulation.*

Index Terms—*diabetes, continuous glucose monitoring, fog computing, edge intelligence, federated learning, explainable AI, ambient assisted living, remote patient monitoring.*

I. INTRODUCTION

Diabetes is one of the largest chronic-disease burdens worldwide. The International Diabetes Federation estimated that about 537 million adults aged 20–79 were living with diabetes in 2021 and projected 783 million by 2045 [1]. Good outcomes depend on frequent glucose measurement, timely insulin or medication adjustment, diet and physical activity, which clinical guidelines increasingly support with diabetes technology such as continuous glucose monitoring (CGM) [2]. For older adults living alone, the same monitoring must also cover falls, inactivity and daily-routine changes, which is the goal of ambient assisted living (AAL) [3].

Cloud-centric designs process everything in remote data centers. For health monitoring this raises three concerns: round-trip delay and dependence on connectivity for time-critical alerts, continuous upload of identifiable health data, and recurring bandwidth and storage cost. Fog computing places computation on gateways and local servers between devices and the cloud [4], and edge computing pushes part of it onto the devices themselves [5]. Smart e-health gateways have already been shown to host local analytics, data filtering and protocol translation for healthcare IoT [6].

At the same time, machine learning has become standard in diabetes research, from risk stratification on clinical records [7] to deep-learning glucose forecasting [8]. The open question is where the learning and inference should run, and how models can be personalized and improved across many homes without centralizing raw data. Federated learning (FL) [9] and explainable AI [10] provide parts of the answer. This paper reviews these ideas for diabetes care and ambient monitoring. Its contributions are (i) a problem formulation, (ii) a layered architecture and method taxonomy, (iii) a problem-solution mapping, (iv) a consolidated summary of reported evidence, and (v) open challenges and future directions.

II. PROBLEM STATEMENT

Consider N users, each with a data stream $D_k = \{(x_t, y_t)\}$ of sensor features x_t (glucose, heart rate, activity, home context) and outcomes y_t (future glucose, hypoglycemia event, diabetes onset). The goal is to learn a personalized predictor f_k for each user that is accurate, timely, private, resource-efficient and explainable. The following problems make this difficult.

P1: Inter-patient variability and data scarcity. Glucose dynamics differ strongly between people and over time, and each person has limited labeled data, so one global model fits individuals poorly [11].

P2: Latency and reliability. Hypoglycemia and falls need alerts within seconds or minutes; a cloud-only chain fails when the internet link is slow or down.

P3: Privacy and regulation. Glucose, location and behavior data are highly sensitive, and centralizing them increases breach risk and legal exposure [12].

P4: Resource constraints. Wearables and gateways have limited memory, compute and battery, so large models must be compressed [13].

P5: Non-IID, heterogeneous data. Devices, sampling rates, populations and lifestyles differ across sites, which hurts collaborative training [14].

P6: Trust and clinical safety. Clinicians and patients need to understand why a risk alert was raised, and average error metrics can hide clinically dangerous errors [15].

III. METHODOLOGY: ARCHITECTURE AND TAXONOMY

This is a narrative review organized around a layered architecture. Peer-reviewed methods, public datasets and clinical guidance are grouped by the layer where they operate (Fig. 1) and by the stage of the data-to-decision workflow (Fig. 2).

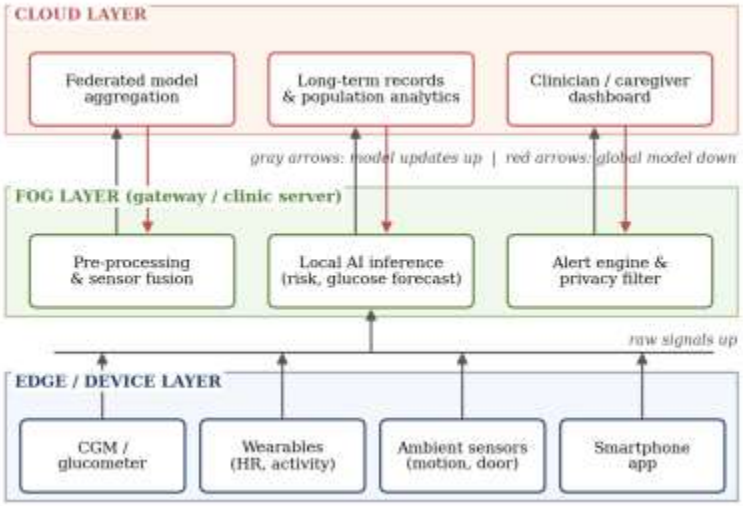


Fig. 1. Layered edge–fog–cloud architecture for diabetes monitoring and ambient assisted care.

A. Layered Architecture

The **edge layer** contains CGM sensors, glucometers, wearables, smartphones and ambient sensors (motion, door contacts, bed or pressure sensors). Devices perform light operations such as filtering, compression and simple threshold checks. The **fog layer**, typically a home gateway or clinic server, performs sensor fusion, runs the personalized model, applies privacy filters and issues alerts, and can keep working when the cloud link is lost. The **cloud layer** aggregates models, stores long-term records, supports population analytics and offers clinician dashboards. Simulators such as iFogSim [16] allow resource placement and latency of such hierarchies to be studied before deployment.

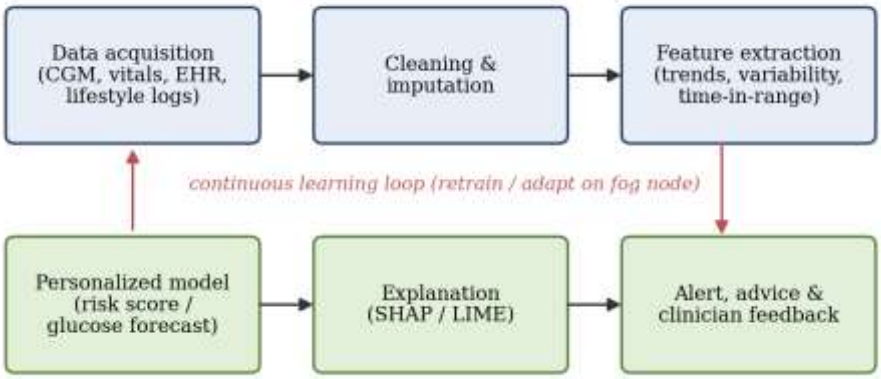


Fig. 2. Data-to-decision workflow for personalized prediction with an explanation step and a continuous learning loop.

B. AI Models for Risk Prediction and Glucose Forecasting

Table I lists the main model families. Classical models such as logistic regression and tree ensembles (for example XGBoost [17]) suit tabular risk prediction from clinical and lifestyle variables, and their small size makes them easy to run on a gateway [7]. Recurrent networks, especially long short-term memory (LSTM) models [18], are common for glucose time series, and recent work covers hybrid and attention-based forecasting [8]. A diabetes risk score can be written as

$$p = \sigma(\mathbf{w}^T\mathbf{x} + b) = 1 / (1 + exp(-(\mathbf{w}^T\mathbf{x} + b))) \quad (1)$$

and forecasting accuracy at horizon h is usually reported as the root mean square error over N samples,

$$RMSE_h = \sqrt{\ [\ (1/N) \sum_i (\hat{g}_i - g_i)^2\ ]} \quad (2)$$

where $\hat{g}_i$ and g_i are predicted and measured glucose. Because RMSE can hide dangerous errors, the Clarke error grid [15] is used to check whether errors fall into clinically acceptable zones.

TABLE I. Main AI model families used in diabetes analytics

Family	Typical task	Fog suitability
Logistic regression, trees, XGBoost	Risk score from tabular records	High: small, fast
LSTM / GRU	Glucose forecasting from CGM series	Medium: needs compression
Hybrid / attention models	Multi-sensor forecasting	Low-medium: cloud or distilled
Activity-recognition CNN / RNN	Fall and routine detection (AAL)	Medium: edge or fog
Federated personalized models	Cross-site learning without raw data	High: designed for fog nodes

C. Federated and Personalized Learning

In FL, each fog node trains on local data and sends only model updates to an aggregator [9], as shown in Fig. 3. In federated averaging, the global weights are the data-weighted mean of the local weights from the K participating nodes:

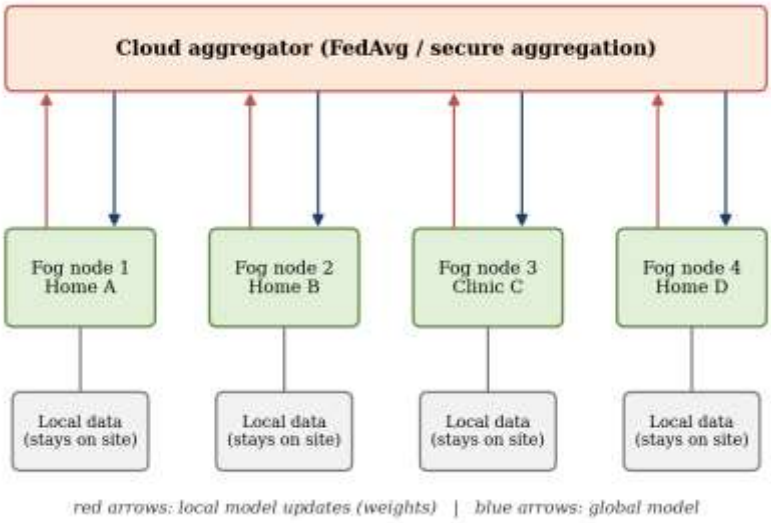


Fig. 3. Federated learning across home and clinic fog nodes: raw data stays local and only model updates are exchanged.

$$w_{t+1} = \sum_{k=1}^{K} (n_k / n)\ w_{t+1}^k \quad (3)$$

where n_k is the sample count of node k and $n = \sum n_k$. When data are non-IID, FedProx adds a proximal term $(\mu/2)\|w - w_t\|^2$ to each local objective to limit drift [19]. Personalization can then be done by fine-tuning the global model on each user's recent data, or by keeping a shared feature extractor and a local output layer. Surveys describe the underlying concepts and open issues [20, 14], and a review of digital health stresses its promise for clinical data that cannot leave the institution [12].

D. Privacy, Explainability and Model Compression

Model updates can still leak information, so FL is usually combined with secure aggregation [21] and differential privacy [22], which add noise to bound what any single record can reveal, at some cost in accuracy. Post-hoc explanation methods such as LIME [23] and SHAP [10] identify the features behind a prediction (for example recent glucose trend or meal pattern) so that the fog node can attach a reason to each alert. For constrained nodes, quantization [13] and knowledge distillation [24] reduce model size and inference cost.

IV. SOLUTIONS PROVIDED BY EDGE–FOG AI

Table II connects each problem of Section II with the solution offered by the layered, federated and explainable design.

TABLE II. Mapping of problems to solutions

Problem	Solution	Key references
P1 Variability, scarce data	Personalized fine-tuning of a federated global model	[11], [9]
P2 Latency, reliability	Local inference and alerting on the fog node; offline operation	[4], [6]
P3 Privacy, regulation	Keep raw data local; secure aggregation; differential privacy	[21], [22], [12]
P4 Resource limits	Quantization, distillation, small tree models	[13], [24]
P5 Non-IID data	FedProx and personalization layers	[19], [14]
P6 Trust, safety	SHAP / LIME explanations; error-grid evaluation	[10], [23], [15]

Fast, local decisions. Because the fog node sits in the home or clinic, inference and alert generation do not wait for the cloud. Only periodic summaries and model updates travel upstream, saving bandwidth. FedAvg itself was designed to cut communication; it was reported to need 10–100 times fewer communication rounds than synchronized stochastic gradient descent in the original experiments [9].

Personal yet collaborative models. A global model learns patterns shared across many users, and a few days of local data are then enough to adapt it to the individual, which addresses the scarcity of per-patient data [11].

Trustworthy alerts. Attaching a feature-based explanation and a confidence level to each alert helps clinicians judge whether to act, and a tiered escalation policy avoids alarm fatigue (Fig. 4).

V. AMBIENT ASSISTED HEALTHCARE MONITORING

AAL systems combine physiological data with context from the home to support independent living [3]. For diabetes, ambient information adds value in several ways: reduced night-time movement or repeated bathroom visits may accompany nocturnal hypoglycemia or hyperglycemia, long inactivity may indicate a fall or illness, and missed meal-time routines may signal that insulin timing needs review. Human activity recognition from inertial sensors is well established, with public datasets such as the smartphone-based dataset of Anguita *et al.* [25]. In a fog design, these activity inferences run locally, and only event summaries are shared with caregivers. Fig. 4 shows a typical alert pipeline in which the risk model runs on the fog node and escalation depends on severity, with local alerting retained if the cloud link is lost.

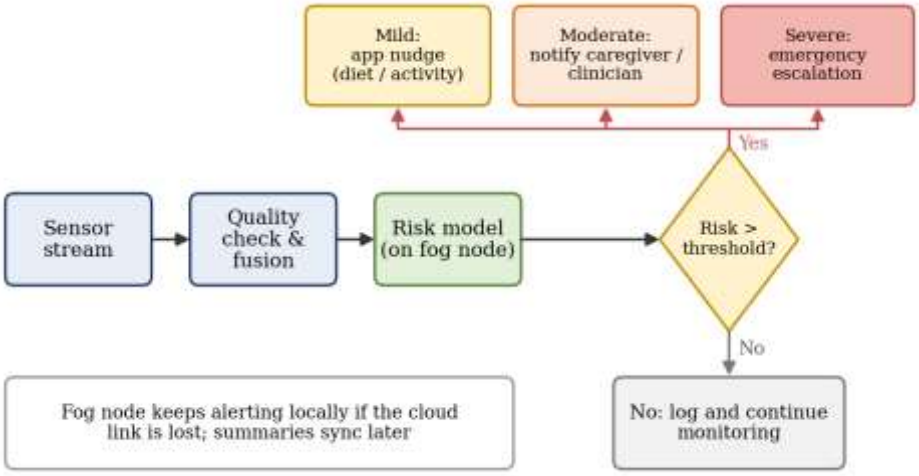


Fig. 4. Real-time alert pipeline on a fog node with severity-based escalation.

VI. RESULTS AND DISCUSSION

This section consolidates evidence reported in the cited literature; no new experiments were run. Table III summarizes representative findings, and Fig. 5 gives a qualitative comparison of the three tiers, which is the author's synthesis of the trade-offs described in [4, 5, 6] rather than a measured result.

TABLE III. Representative evidence reported in the literature

Topic	Reported evidence	Source
Federated training cost	10–100× fewer communication rounds than synchronized SGD in the original FedAvg experiments	[9]
Glucose forecasting	RNN with variance estimation on OhioT1DM: RMSE about 18.9 mg/dL (30 min) and 31.4 mg/dL (60 min)	[26]
Benchmark data	OhioT1DM: 12 people with type 1 diabetes, about eight weeks of CGM, insulin and wearable data each	[27]
Risk prediction data	Pima dataset: 768 records, 8 clinical attributes; widely used but small, so results should be read with caution	[28], [7]
Model compression	8-bit integer inference reduces model size roughly four-fold versus 32-bit floats	[13]
Fog e-health gateway	Local filtering, analysis and protocol handling at the gateway reduces dependence on cloud processing	[6]

	Cloud	Fog	Edge
Responsiveness (latency)	Low	High	High
Bandwidth saving	Low	High	High
Data privacy (locality)	Low	High	High
Compute capacity	High	Medium	Low
Global / population learning	High	Medium	Low
Resilience when offline	Low	High	Medium

Fig. 5. Qualitative comparison of cloud, fog and edge tiers for health monitoring (author's synthesis).

Observations. (1) The fog tier offers the best balance for health monitoring: it is responsive, keeps data local and has enough compute for personalized models, whereas the cloud is best kept for aggregation and population analytics (Fig. 5). (2) Personalization matters, since errors depend heavily on the individual and the horizon; the rise in RMSE from 30 to 60 minutes in Table III shows how quickly forecasts degrade [26]. (3) Federated training generally approaches centralized accuracy when data are similar across sites, but non-IID data can slow or destabilize it, which motivates FedProx-style methods [19, 14]. (4) Most published evidence is retrospective or simulated; prospective trials of fog-based diabetes systems in real homes remain scarce. (5) Accuracy should always be paired with clinical safety measures such as error-grid analysis [15] and with explanation quality.

VII. CHALLENGES AND FUTURE DIRECTIONS

Clinical validation and regulation. Systems that drive treatment decisions need prospective studies, transparent reporting and regulatory approval. Public benchmarks are small, so larger multi-site, multi-ethnic datasets are needed.

Non-IID data and fairness. Personalized FL methods that stay stable across devices and populations, and checks that performance does not drop for under-represented groups, are open problems [14].

Stronger privacy at low cost. Combining secure aggregation and differential privacy with limited energy and accuracy loss on gateways is still difficult [21, 22].

Energy-aware and adaptive deployment. Scheduling training and inference across devices and fog nodes under battery and network limits can be studied with simulators such as iFogSim [16], followed by pilot deployments.

Explainability that clinicians use. Explanation methods should be evaluated with clinicians and patients, not only by technical fidelity [10, 23].

Multimodal and closed-loop care. Future systems may combine CGM, activity, diet, sleep and ambient data, and connect forecasts to decision support or automated insulin delivery under strict safety constraints.

Security and trust. Gateways are attack targets; secure boot, authentication, tamper-evident logs and defenses against poisoned model updates are required.

Large language models for patient support. Language models running locally or at the fog tier could turn alerts and trends into plain-language guidance, but they need safeguards against incorrect advice.

VIII. CONCLUSION

Diabetes care and ambient monitoring generate continuous, private data and need quick, trustworthy decisions. A layered edge–fog–cloud architecture with personalized and federated learning, explainable outputs and compressed models addresses the problems of latency, privacy, scarcity and trust that limit cloud-only systems. The evidence so far is encouraging but largely retrospective. Progress now depends on prospective clinical validation, robust handling of heterogeneous data, efficient privacy protection and clinician-centered explanation.

REFERENCES

[1] International Diabetes Federation, *IDF Diabetes Atlas*, 10th ed. Brussels, Belgium: IDF, 2021.

[2] American Diabetes Association Professional Practice Committee, "Standards of Care in Diabetes—2024," *Diabetes Care*, vol. 47, Suppl. 1, 2024.

[3] P. Rashidi and A. Mihailidis, "A survey on ambient-assisted living tools for older adults," *IEEE J. Biomed. Health Inform.*, vol. 17, no. 3, pp. 579-590, 2013.

[4] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the Internet of Things," in *Proc. 1st MCC Workshop on Mobile Cloud Computing*, 2012, pp. 13-16.

[5] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637-646, 2016.

[6] A. M. Rahmani *et al.*, "Exploiting smart e-Health gateways at the edge of healthcare Internet-of-Things: A fog computing approach," *Future Gener. Comput. Syst.*, vol. 78, pp. 641-658, 2018.

[7] I. Kavakiotis, O. Tsave, A. Salifoglou, N. Maglaveras, I. Vlahavas, and I. Chouvarda, "Machine learning and data mining methods in diabetes research," *Comput. Struct. Biotechnol. J.*, vol. 15, pp. 104-116, 2017.

[8] T. Zhu, K. Li, P. Herrero, and P. Georgiou, "Deep learning for diabetes: A systematic review," *IEEE J. Biomed. Health Inform.*, vol. 25, no. 7, pp. 2744-2757, 2021.

[9] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. AISTATS, PMLR*, vol. 54, 2017, pp. 1273-1282.

[10] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017, pp. 4765-4774.

[11] S. Oviedo, J. Vehi, R. Calm, and J. Armengol, "A review of personalized blood glucose prediction strategies for T1DM patients," *Int. J. Numer. Method. Biomed. Eng.*, vol. 33, no. 6, e2833, 2017.

[12] N. Rieke *et al.*, "The future of digital health with federated learning," *npj Digit. Med.*, vol. 3, art. 119, 2020.

[13] B. Jacob *et al.*, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Proc. IEEE/CVF CVPR*, 2018, pp. 2704-2713.

[14] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50-60, 2020.

[15] W. L. Clarke, D. Cox, L. A. Gonder-Frederick, W. Carter, and S. L. Pohl, "Evaluating clinical accuracy of systems for self-monitoring of blood glucose," *Diabetes Care*, vol. 10, no. 5, pp. 622-628, 1987.

[16] H. Gupta, A. Vahid Dastjerdi, S. K. Ghosh, and R. Buyya, "iFogSim: A toolkit for modeling and simulation of resource management techniques in the Internet of Things, Edge and Fog computing environments," *Softw. Pract. Exper.*, vol. 47, no. 9, pp. 1275-1296, 2017.

[17] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785-794.

[18] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735-1780, 1997.

[19] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. Machine Learning and Systems (MLSys)*, vol. 2, 2020, pp. 429-450.

[20] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, art. 12, 2019.

[21] K. Bonawitz *et al.*, "Practical secure aggregation for privacy-preserving machine learning," in *Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS)*, 2017, pp. 1175-1191.

[22] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3-4, pp. 211-407, 2014.

[23] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 1135-1144.

[24] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.

[25] D. Anguita, A. Ghio, L. Oneto, X. Parra, and J. L. Reyes-Ortiz, "A public domain dataset for human activity recognition using smartphones," in *Proc. 21st European Symp. Artificial Neural Networks (ESANN)*, 2013, pp. 437-442.

[26] J. Martinsson *et al.*, "Blood glucose prediction with variance estimation using recurrent neural networks," *J. Healthc. Inform. Res.*, vol. 4, pp. 1-18, 2020.

[27] C. Marling and R. Bunescu, "The OhioT1DM dataset for blood glucose level prediction: Update 2020," in *Proc. 5th Int. Workshop on Knowledge Discovery in Healthcare Data, CEUR Workshop Proc.*, vol. 2675, 2020, pp. 71-74.

[28] J. W. Smith, J. E. Everhart, W. C. Dickson, W. C. Knowler, and R. S. Johannes, "Using the ADAP learning algorithm to forecast the onset of diabetes mellitus," in *Proc. Symp. Computer Applications in Medical Care*, 1988, pp. 261-265.