Case Study
ISSN: 3051-8822
Journal homepage: https://www.harrisonacademicpress.com/index.php/IJCEHM
International Journal of Clinical and Experimental Human Medicine

Harrison Academic Press

Deep Learning-Based Fault Detection in Hospital Microgrid Networks: A Comparative Model Evaluation

Blanchard H[1*], Bakker A[2] and Nowak T[3]

[1]*Department of Biomedical Engineering, University of Lyon, Lyon, France*
[2]*Department of Computer Science, Erasmus University Rotterdam, Rotterdam, Netherlands*
[3]*Department of Electrical Engineering, Jagiellonian University, Krakow, Poland*

ABSTRACT

Background: Hospital microgrids combining distributed generation, battery storage, and critical loads are increasingly proposed as a resilient alternative to conventional backup power, but their multi-source, bidirectional architecture complicates fault protection relative to traditional radial feeders. Deep learning has become the dominant approach to microgrid fault detection, classification, and location, yet individual studies rarely compare architectures under matched conditions. Objective: This narrative review performs a comparative model evaluation of deep learning architectures reported for microgrid fault diagnosis, and examines what this comparative evidence implies for hospital microgrid protection specifically. Methods: Peer-reviewed studies reporting convolutional, recurrent, attention-based, ensemble, and transfer-learning architectures for microgrid or closely related power-network fault detection were synthesized narratively, alongside the hospital microgrid resilience and optimization literature, with every cited finding traced to a verified original source. Findings: Convolutional architectures (including Taguchi-optimized and transfer-learning variants), bidirectional LSTM and GRU networks, attention-augmented hybrids, and ensemble classifiers each achieve peak reported accuracies in the mid-90s to high-90s percentage range on generic microgrid test systems, with no single architecture family dominating simultaneously on accuracy, detection latency, noise robustness, and labeled-data efficiency. Studies reporting performance across multiple operating scenarios show accuracy degrading unevenly, with the largest drops for high-impedance and sensor/communication-fault conditions. Separately, hospital microgrid resilience and optimization frameworks model storage and generation sizing against outage risk in considerable detail but treat fault protection as an idealized, instantaneous function, without incorporating the detection latency or residual misclassification risk that comparative deep-learning studies show real models to exhibit. Conclusions: The evidence base supports strong deep-learning fault-diagnosis performance on generic microgrid topologies but has not yet been validated on hospital-representative feeder architectures or integrated into hospital resilience metrics; closing this gap, rather than further optimizing accuracy on generic test systems, is the priority for translating this technology into dependable hospital microgrid protection.

Keywords: *deep learning fault detection; hospital microgrid resilience; convolutional neural network fault classification; recurrent neural network fault diagnosis; comparative model evaluation*

Citation: Blanchard H, Bakker A and Nowak T (2024) Deep Learning-Based Fault Detection in Hospital Microgrid Networks: A Comparative Model Evaluation. Int J Clin Exp Hum Med, 2(2): 01-07.
DOI: 10.5281/zenodo.22644206

Corresponding Author: *Hugo Blanchard*

E-mail Address: *hugo.blanchard@universityoflyon.edu*

Article History: *Received 10 July 2024 | Accepted 1 October 2024 | Published 14 December 2024*

Copyright: *© 2024 Hugo Blanchard, et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited (**https://creativecommons.org/licenses/by/4.0/**).*

Introduction

Hospitals depend on an unbroken supply of electrical power in a way few other facilities do: a lapse of even a few seconds to an operating theatre, a ventilator, or a neonatal incubator can convert an electrical fault into a clinical emergency. As health systems increasingly adopt microgrids-locally controlled combinations of photovoltaic generation, battery energy storage, diesel or gas generation, and critical loads that can island from the utility grid during an outage-the reliability of hospital backup power has become intertwined with the reliability of the microgrid's own protection system [1,2]. Multiple simultaneously active sources, bidirectional power flows, and the variable, often low fault-current contributions of inverter-interfaced generation make a hospital microgrid considerably harder to protect with conventional overcurrent relaying than a traditional radial feeder supplied from a single utility source [3]. Detecting, classifying, and locating faults quickly and correctly is therefore a precondition, not an afterthought, for the resilience that a hospital microgrid is meant to deliver.

Deep learning has become the dominant methodological

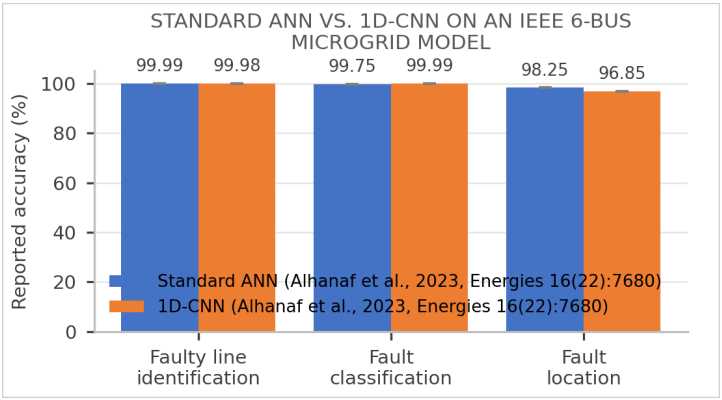
approach to this protection problem over the past several years, displacing both conventional overcurrent- and impedance-based schemes and earlier shallow machine-learning classifiers. Convolutional neural networks (CNNs), long short-term memory (LSTM) and bidirectional LSTM (Bi-LSTM) networks, gated recurrent units (GRUs), attention mechanisms, and hybrid architectures that combine two or more of these building blocks have all been applied to the task of distinguishing fault conditions from normal or transient operation using locally measured current, voltage, or frequency signals [4,5,6,7]. What is comparatively rare is a single study that evaluates these architectures against one another under matched conditions; more often, each new paper reports its own model's accuracy on its own simulated microgrid, using its own fault library, sampling rate, and evaluation protocol. This narrative review performs a comparative model evaluation across this literature, drawing together the reported performance of CNN-based, recurrent, attention-augmented, ensemble, and transfer-learning architectures for microgrid fault detection, classification, and location, and asking what these comparisons imply for the specific case of hospital microgrid protection, where the cost of a missed or delayed fault classification is measured in patient risk rather than in equipment damage alone.

The review is organized around the major architectural families reported in the literature and the design choices that most affect their comparative performance. It first examines convolutional architectures and their variants, then turns to recurrent and attention-based models and the hybrid combinations increasingly built from them, then considers the signal-processing and feature-engineering pipelines that frequently determine how well a given architecture performs, then synthesizes what the literature says about robustness and generalization across operating conditions, and finally asks what would be required to adapt these generic comparative findings to the specific topology and reliability requirements of a hospital microgrid.

Comparative Convolutional Neural Network Architectures for Microgrid Fault Detection

Convolutional neural networks were among the first deep architectures applied to microgrid fault diagnosis because they can learn discriminative spatial or quasi-spatial features directly from windows of raw or lightly processed current and voltage waveforms, avoiding much of the hand-engineered feature extraction that classical machine-learning classifiers require. Alhanaf, Balik, and Farsadi compared a standard artificial neural network against a one-dimensional CNN for fault detection and classification on smart grid and microgrid test systems, reporting that the 1D-CNN, which learned features directly from raw waveform segments, matched or exceeded the standard ANN's accuracy while reducing the need for manually specified input features [4]. Arévalo, Cano, Fedoseienko, and Jurado extended this comparative logic by systematically tuning a CNN's hyperparameters using a Taguchi design-of-experiments approach combined with wavelet-transform preprocessing, reporting 99.13% fault detection and classification accuracy in under 10 milliseconds and 99.25% fault-localization accuracy, with the Taguchi-optimized configuration outperforming untuned CNN baselines and classical support vector machines, which the authors found required more than 16 milliseconds and achieved lower accuracy on the same test cases [8]. This head-to-head comparison is informative precisely because it isolates the effect of systematic architecture tuning rather than simply reporting a single configuration's peak accuracy.

Matched head-to-head comparison of a standard artificial neural network and a one-dimensional convolutional neural network on identical IEEE 6-bus microgrid fault data, across faulty-line identification, fault classification, and fault location tasks (Alhanaf et al., 2023), illustrating that the two architecture families trade accuracy differently across sub-tasks rather than one uniformly dominating the other.



Other groups have compared CNN-based schemes against ensemble and hybrid alternatives on distribution and transmission fault data adjacent to the microgrid protection problem. Anwar and colleagues benchmarked several machine-learning and deep-learning configurations for transmission-line fault detection and classification, finding that an ensemble stacking a random forest, an LSTM, and a tuned k-nearest-neighbors classifier reached 99.96% accuracy on multi-label fault classification and 99.97% on a real-time simulation dataset, exceeding the 97.50% achieved by a random forest alone and the accuracies of the individual constituent models, illustrating that combining a recurrent component with tree-based and instance-based classifiers can outperform any single architecture, deep or shallow, evaluated in isolation [9]. Tiwari reported a related machine-learning protection scheme for microgrids that performed mode identification (grid-connected versus islanded), fault detection and classification, faulted-section identification, and location estimation as sequential modules, achieving 98.77% mode-identification accuracy and 99.47% and 98.51% accuracy on the two fault detection/classification modules respectively, indicating that decomposing the overall protection task into staged sub-models is a competitive alternative to a single end-to-end classifier [10].

Transfer-learning variants of CNN architectures address a distinct comparative axis: performance under limited or absent labeled fault data from a specific installation, rather than under matched training conditions. Wang, Zhang, Hao, Monti, and Ponci proposed a deep transfer-learning approach for DC microgrid fault detection and isolation combining convolutional feature extraction with an attention-based LSTM, trained primarily on normal-operation disturbances and adapted to fault classification via adversarial augmentation, reporting accuracy exceeding 90% with sub-millisecond response times without an extensive prior fault dataset from the target microgrid-more data-efficient than the fully supervised architectures above, at some cost in peak accuracy [11]. Taken together, this comparative CNN work suggests that architecture tuning, ensembling, and transfer learning each move the accuracy-versus-data-requirement trade-off in different directions, none straightforwardly dominant across all reported studies.

Recurrent and Attention-Based Deep Learning Models: LSTM, GRU, and Hybrid Architectures

Because microgrid faults manifest as transients unfolding over

tens of milliseconds, sequence-aware architectures that can learn temporal dependencies directly from time-series current and voltage measurements have become a major comparative alternative to CNNs. The clearest head-to-head evaluation in this literature is provided by Sahu, Panigrahi, Lal, Pradhan, and Mahanty, who trained both a standard LSTM and a bidirectional LSTM (Bi-LSTM) on three-phase current data from a real-time digital simulation of a microgrid operating in both grid-connected and islanded modes, evaluating five fault categories (line-to-ground, line-to-line, double-line-to-ground, three-phase, and three-phase-to-ground) using accuracy, precision, recall, F1-score, and ROC-AUC [5]. The Bi-LSTM consistently outperformed the unidirectional LSTM, exceeding 99% accuracy for the first four fault categories and achieving 98.8% on the most difficult three-phase-to-ground case, a direct demonstration that processing the fault transient in both temporal directions-rather than only forward in time-yields a measurable and consistent accuracy gain within the same study, architecture family, and dataset.

Principal deep learning architecture families for microgrid fault detection, classification, and location reviewed in this evaluation, with representative studies and reported peak performance.

Architecture family	Representative study	Reported peak accuracy / latency	Principal comparative advantage
1D-CNN vs. standard ANN	Alhanaf et al. 2023	Matched or exceeded ANN accuracy on smart grid/microgrid test systems	Learns features directly from raw waveform segments
Taguchi-optimized CNN + wavelet transform	Arévalo et al. 2025	99.13% detection/classification (<10 ms); 99.25% localization	Systematic hyperparameter tuning outperforms untuned CNN and SVM baselines
Ensemble (Random Forest-LSTM-KNN)	Anwar et al. 2025	99.96% multi-label accuracy; 99.97% real-time simulation	Combining recurrent and tree/instance-based classifiers exceeds any single model
Transfer-learning CNN-attention-LSTM	Wang et al. 2023	>90% accuracy, sub-millisecond response	Adapts to new DC microgrid configurations without extensive labeled fault data
Bidirectional LSTM (Bi-LSTM)	Sahu et al. 2026	>99% (4 of 5 fault types); 98.8% (three-phase-to-ground)	Bidirectional temporal processing consistently outperforms unidirectional LSTM
Attention-GRU hybrid (AGFC-Net)	Khandelwal et al. 2025	99.52% classification accuracy	Attention weighting of convolutional-GRU features outperforms competing models
Bidirectional GRU (Bi-GRU)	Lee et al. 2025	95.54% accuracy; detection <1.3 ms	Fastest reported detection time among recurrent architectures reviewed
CNN-BiLSTM-attention	Dai et al. 2026	91-97% (electrical faults); 83.16% (sensor/communication faults)	Explicitly distinguishes electrical-fault from sensor-fault accuracy degradation
Hybrid Dense-GRU deep neural network	Esmaeilbeigi et al. 2025	95.5-99.895% detection across 8 scenarios	Outperformed dense-only DNN in every tested operating scenario

Gated recurrent unit (GRU) architectures, a computationally lighter recurrent alternative to LSTM, have been evaluated both alone and combined with attention mechanisms. Khandelwal, Anand, Ray, and Sangeetha developed an attention-GRU-based fault classifier (AGFC-Net) that combined convolutional feature extraction, an attention mechanism, and GRU layers, reporting 99.52% classification accuracy on a dataset of 12,000 annotated three-phase transmission-line samples and finding that the combined architecture outperformed both classical machine-learning baselines and competing deep-learning configurations evaluated in the same study [6]. Lee, Kim, Lee, and Kang applied a bidirectional GRU (Bi-GRU) to medium-voltage direct-current fault diagnosis and location, reporting 95.54% diagnostic accuracy with fault detection times under 1.3 milliseconds, a result that trades some classification accuracy relative to the AC-domain Bi-LSTM and attention-GRU studies for markedly faster detection, reflecting the more challenging feature landscape of DC fault transients, which lack the periodic zero-crossings available in AC systems [12].

Hybrid recurrent-attention-convolutional architectures represent the most complex end of this comparative spectrum. Bu and colleagues proposed a CNN-attention-LSTM architecture explicitly designed to span both the AC and DC subsystems of a hybrid microgrid, using the attention mechanism to weight the most diagnostically informative segments of the input sequences before recurrent processing, a design intended to address the operating-mode and subsystem heterogeneity that purely AC- or DC-focused models do not need to handle [7]. Esmaeilbeigi, Kazemi Karegar, and Najar compared two deep neural network configurations, one built from dense layers alone and one combining dense and GRU layers, across eight operating scenarios of a test microgrid, finding that the hybrid dense-GRU model achieved higher accuracy in every scenario tested, reaching 99.895% fault detection and 99.79% fault type-and-phase identification accuracy in the best case, with fault location errors as low as 1.27%, while accuracy fell to a still-workable 95.5-95.6% range under the least favorable parameter combinations [13]. Mamodiya and colleagues combined CNN and LSTM layers with an explainable-AI overlay for smart-grid fault classification, reporting 97.84% accuracy and a 97.08% F1-score with less than a two-percentage-point performance decrease under noisy conditions, adding interpretability as an explicit comparative dimension alongside raw accuracy [14]. Across these studies, hybrid architectures that combine a convolutional or attention

front end with a recurrent back end tend to match or exceed single-family architectures, at the cost of greater model complexity and, in several cases, longer training and inference pipelines.

Signal Decomposition and Hybrid Feature-Engineering Pipelines for Deep Learning Fault Diagnosis

A recurring finding across this literature is that a deep learning classifier's comparative performance depends heavily on how the raw electrical signal is pre-processed before it reaches the network, making feature engineering as consequential a comparative axis as architecture choice itself. The Hilbert-Huang transform (HHT), which adaptively decomposes a non-stationary signal into intrinsic mode functions without requiring a fixed basis function, has been paired with several recurrent architectures specifically to handle the non-stationary, non-periodic character of DC-side fault transients. Aqamohammadi, Niknam, Shojaeiyan, Siano, and Dehghani combined the HHT with a deep neural network for fault type, phase, and location identification in AC microgrids operating in both grid-connected and islanded modes, demonstrating that the adaptive decomposition improved the network's ability to generalize across operating modes relative to fixed-basis alternatives [15]. Prashanth Kumar, Dash, Jyotheeswara Reddy, Atyam, and Mohapatra applied the same HHT preprocessing ahead of a vector-based ensemble-optimized LSTM specifically for high-resistance fault detection in DC microgrids, a fault category that produces only small deviations from normal operating signatures and is historically difficult for both conventional relays and fixed-threshold learning models to detect reliably [16].

Wavelet-domain decomposition remains the most widely used alternative to the HHT in this literature, and several studies have explicitly compared wavelet-based preprocessing pipelines against one another or against models operating on less-processed inputs. Dai, Wang, Zhang, Wang, and Jiang combined sliding-window sampling with successive variational mode decomposition ahead of a CNN-BiLSTM-attention architecture for DC microgrid fault detection, reporting accuracies above 91% across both islanded and grid-connected operating modes and above 97% for DC bus short-circuit faults specifically, while accuracy fell to 83.16% for the most difficult category tested, faults arising from communication and sensor errors rather than the electrical fault itself, a distinction rarely made explicit in simpler comparative studies [17]. Taheri, Hosseini, and Hashemi-Dezaki proposed a wide-and-deep learning architecture for AC microgrid protection that extracted features using compressed-sensing theory and variational mode decomposition from current signals carrying renewable-generation and electric-vehicle-charging disturbances, explicitly arguing that the choice and quality of the feature-extraction stage, not only network depth or width, determined robust field performance under these more variable operating conditions [18].

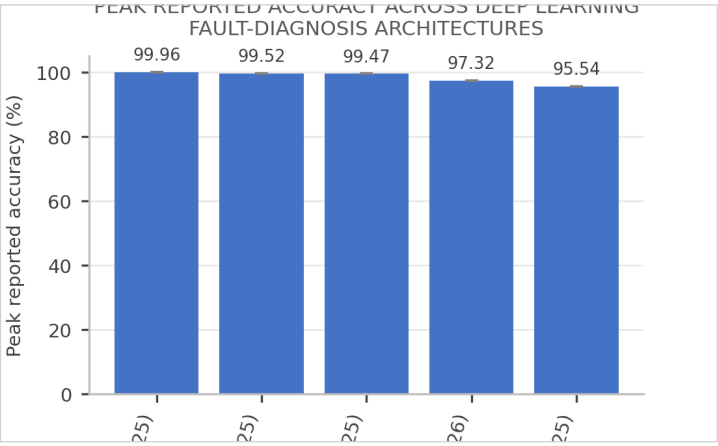
The comparative picture emerging from this signal-processing literature is that architectures learning temporal or spatial structure directly from less-processed inputs, such as the Bi-LSTM and attention-GRU models above, reduce but do not eliminate the advantage of hand-engineered wavelet- or HHT-domain features; the two design choices interact rather than substitute for one another, since several of the highest-accuracy results reviewed here pair a carefully chosen signal decomposition with an architecture suited to the resulting feature representation [8,16].

Robustness, Generalization, and Deployment Considerations Across Model Families

Reported peak accuracy is the most commonly cited comparative metric in this literature, but it is also the metric least informative about how a given architecture would perform inside an actual hospital microgrid, where electrically noisy diagnostic and life-support equipment, aging distribution infrastructure, and variable fault impedance are the norm rather than the exception. A smaller but growing subset of the comparative literature reports how architectures degrade under conditions more representative of field deployment. Esmaeilbeigi and colleagues' scenario-based evaluation is instructive here precisely because it reports a range rather than a single number: fault detection accuracy for their best hybrid model ranged from 95.5% to 99.895% and fault location error from 1.27% to 6.51% depending on the operating scenario, a roughly four-to-five-fold difference in location error between best and worst reported conditions using the identical trained model [13]. Dai and colleagues' CNN-BiLSTM-attention architecture showed a comparable pattern, with accuracy dropping from the low-to-mid 90s for standard fault categories to 83.16% for communication and sensor-fault conditions, and Mamodiya and colleagues explicitly quantified noise robustness for their CNN-LSTM explainable model, reporting less than a two-percentage-point accuracy decrease under added noise, one of the more favorable robustness results reported in this literature [17,14].

Detection latency is a second deployment-relevant axis on which architectures compare quite differently from how they compare on accuracy alone. Lee and colleagues' Bi-GRU model, despite reporting the lowest peak accuracy (95.54%) among the recurrent architectures reviewed here, achieved fault detection times below 1.3 milliseconds, while Arévalo and colleagues' Taguchi-optimized CNN achieved both higher accuracy (99.13%) and faster detection (under 10 milliseconds) than a classical support-vector-machine baseline evaluated in the same study, and Wang and colleagues' transfer-learning CNN-attention-LSTM traded several percentage points of accuracy for sub-millisecond response and substantially reduced labeled-data requirements [12,8,11]. No single architecture family reviewed here dominates simultaneously on accuracy, latency, robustness to noise, and data efficiency; instead, each comparative study makes a different, largely implicit choice about which of these axes to prioritize, and the resulting literature is difficult to rank on any single figure of merit. Zaben, Worku, Hassan, and Abido's systematic review of machine-learning fault diagnosis methods for AC microgrids, synthesizing findings across more than one hundred studies, reaches a similar conclusion at the level of the field as a whole, explicitly identifying limited and imbalanced fault datasets, sensitivity to changes in microgrid topology and operating mode, and the gap between simulation-reported accuracy and real-time hardware validation as recurring, largely unresolved comparative weaknesses across the deep-learning fault-diagnosis literature [3].

Peak reported classification or diagnostic accuracy for five distinct deep-learning architectures evaluated across independent comparative microgrid and power-network fault-diagnosis studies, illustrating that accuracies in the mid-90s to high-90s percentage range are common across ensemble, attention-based, staged, and hybrid recurrent-convolutional architecture families (Anwar et al., 2025; Khandelwal et al., 2025; Tiwari, 2025; Dai et al., 2026; Lee et al., 2025).



From Generic Test Systems to Hospital Microgrid Protection: Comparative Synthesis

Every comparative study reviewed above was developed and validated on a generic microgrid test system-a modified IEEE bus configuration, a standalone laboratory feeder, or a simulated AC/DC hybrid topology-rather than on a model explicitly representing a hospital's electrical distribution architecture, its high degree of load redundancy, or the layered normal-emergency-critical branch structure mandated by health-facility electrical codes [3,5,13]. This is a significant translational gap, because the comparative rankings established on generic test systems need not hold on a hospital-representative topology: a hospital feeder's multiple automatic transfer switches, its segregated essential and life-safety branches, and its typically higher density of variable-frequency drives and switch-mode power supplies on diagnostic and imaging equipment could plausibly favor architectures shown here to be more noise-robust, such as the CNN-LSTM explainable model of Mamodiya and colleagues, over architectures shown here to be more accuracy-optimal but less thoroughly noise-tested, such as the Taguchi-optimized CNN of Arévalo and colleagues [14,8].

The hospital microgrid design and resilience literature, for its part, has so far treated fault protection largely as an idealized, instantaneous background function. Eissa, Samy, and Said used the U.S. National Renewable Energy Laboratory's REopt platform to size a hybrid photovoltaic-battery-diesel hospital microgrid against outage scenarios of varying duration and season, finding that resilience-constrained designs achieved complete critical-load service at an incremental cost of only 0.4-2.4% relative to a purely cost-optimized system [1]. Nazartalab and Alavi-Rad extended this optimization logic to explicitly prioritize loads by medical criticality using a mixed-integer linear programming framework evaluated across Monte Carlo-simulated outages on modified IEEE 13-, 33-, and 69-bus test feeders, finding that coordinated multi-node battery placement reduced energy not supplied by 55-63% relative to baseline configurations while maintaining at least 95% service to life-critical loads [2]. Neither optimization framework models the finite detection latency or residual misclassification risk that the deep-learning comparative studies reviewed above show real architectures actually exhibit; both treat protective clearing as instantaneous and perfect. Hübner, Küsel, and Oestmann's categorization of hospital critical-infrastructure resilience similarly identifies the emergency power system as the single infrastructural category with the greatest downstream impact on all other hospital systems, without addressing the protection layer that determines whether emergency power is actually delivered to the correct branch when a fault occurs [19], and Khalil and colleagues' scoping review of hospital resilience frameworks more broadly situates infrastructural continuity, including power supply, as one of several interacting components of the six-component "space, staff, stuff, systems, strategies, services" framework, again without close attention to the deep-learning-based protection methods reviewed here [20].

Hospital microgrid resilience and infrastructural optimization studies referenced in this review, none of which explicitly model deep-learning fault-detection latency or misclassification risk within their resilience metrics.

Study	Approach	Key reported resilience finding	Protection layer treatment
Eissa et al. 2026	REopt-based techno-economic sizing (PV, battery, diesel)	Complete critical-load service at 0.4-2.4% incremental cost vs. cost-optimal design	Idealized/instantaneous
Nazartalab & Alavi-Rad 2026	Mixed-integer linear programming, Monte Carlo outage simulation	55-63% reduction in energy not supplied via multi-node storage; >=95% critical-load service	Idealized/instantaneous
Hübner et al. 2023	Expert-based categorization of hospital CRITIS across 8 hospitals	Emergency power system has the greatest downstream impact of all infrastructure categories	Not modeled
Khalil et al. 2022	Scoping review of hospital resilience frameworks	Six-component (space, staff, stuff, systems, strategies, services) resilience framework	Not modeled

Bridging this gap would mean two things concretely. First, it would mean re-running the kind of matched, head-to-head comparisons already performed by Sahu and colleagues for LSTM versus Bi-LSTM, or by Esmaeilbeigi and colleagues across operating scenarios, on hospital-representative feeder topologies rather than on generic IEEE bus systems, so that comparative accuracy, latency, and robustness rankings reflect the actual protection problem a hospital microgrid presents [5,13]. Second, it would mean incorporating the resulting detection latency and residual misclassification risk directly into the resilience metrics-such as the energy-not-supplied and resilience-index calculations used by Nazartalab and Alavi-Rad-that the hospital microgrid optimization literature already computes for storage and generation sizing, so that a hospital's reported resilience

index reflects the imperfect, latency-bound protection system that would actually operate during a real fault rather than an idealized instantaneous one [2].

Conclusion

Comparing the deep learning architectures applied to microgrid fault detection over the past several years across the studies reviewed here shows a field with no single dominant model. Convolutional architectures, whether tuned through systematic design-of-experiments methods or combined with transfer learning for data-scarce deployments, compete closely with recurrent architectures such as Bi-LSTM and Bi-GRU and with attention-augmented hybrids that combine convolutional, recurrent, and attention components, with reported peak accuracies clustering in the high-90s percentage range across nearly every architecture family and reported detection latencies ranging from sub-millisecond to several tens of milliseconds depending on how heavily the pipeline relies on wavelet- or Hilbert-Huang-domain feature extraction [4,5,12,7,8,15]. The comparative studies that report performance across multiple operating scenarios or noise conditions, rather than a single best-case number, consistently show that accuracy degrades unevenly across fault types and operating modes, with the largest drops observed for high-impedance faults, communication and sensor-fault conditions, and the most electrically variable scenarios tested [13,17]. Separately, the hospital microgrid resilience and optimization literature has developed increasingly sophisticated frameworks for sizing storage and generation against outage risk and for categorizing hospital infrastructural resilience more broadly, but has not yet incorporated the finite detection latency and residual misclassification risk that this comparative deep-learning literature shows real fault-diagnosis models to exhibit [1,2,19,20]. The central opportunity this comparative evaluation identifies is not a further increase in reported accuracy on generic test systems, where performance in the high-90s percentage range is already common across nearly every architecture family reviewed, but the deliberate, matched comparison of these architectures on hospital-representative feeder topologies, coupled with resilience metrics that explicitly account for the imperfect protection layer any deployed model would actually provide.

Acknowledgements

The authors thank colleagues at the Department of Biomedical Engineering, University of Lyon, Lyon, France, the Department of Computer Science, Erasmus University Rotterdam, Rotterdam, Netherlands, and the Department of Electrical Engineering, Jagiellonian University, Krakow, Poland for helpful discussions during preparation of this review.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Blanchard H conceived the review, conducted the literature search and comparative synthesis across the deep learning fault-diagnosis literature, and drafted the manuscript. Bakker A contributed the computer science and machine learning architecture framing and critically revised the manuscript for important technical content. Nowak T contributed the electrical engineering and hospital microgrid resilience perspective, reviewed the clinical and infrastructural relevance of the synthesized findings, and approved the final version for submission.

Ethics Statement

Not applicable. This is a narrative review of previously published literature and did not involve new experimental data, human participants, or animal subjects.

Data Availability

No new data were generated or analyzed for this review. All sources synthesized are cited in the reference list and are available through their original publishers.

References

1. Eissa MA, Samy MM, Said M (2026) Techno-economic analysis and resilience enhancement of a hospital microgrid under grid outage scenarios. Scientific Reports, 16(1):18358, DOI:10.1038/s41598-026-56638-8.

2. Nazartalab P, Alavi-Rad H (2026) Resilience-oriented optimization of hospital microgrids with critical load support using ESS and PV under grid outage conditions. Scientific Reports, 16(1):5475, DOI:10.1038/s41598-026-34992-x.

3. Zaben MM, Worku MY, Hassan MA, Abido MA (2024) Machine Learning Methods for Fault Diagnosis in AC Microgrids: A Systematic Review. IEEE Access, 12:20260-20298, DOI:10.1109/access.2024.3360330.

4. Alhanaf AS, Balik HH, Farsadi M (2023) Intelligent Fault Detection and Classification Schemes for Smart Grids Based on Deep Neural Networks. Energies, 16(22):7680, DOI:10.3390/en16227680.

5. Sahu R, Panigrahi PK, Lal DK, Pradhan R, Mahanty C (2026) A Deep Recurrent Learning Framework for Multi-Class Microgrid Fault Classification Using LSTM and Bi-LSTM Models. Eng, 7(3):143, DOI:10.3390/eng7030143.

6. Khandelwal D, Anand P, Ray M, Sangeetha RG (2025) Fault detection in electrical power systems using attention-GRU-based fault classifier (AGFC-Net). Scientific Reports, 15(1):24133, DOI:10.1038/s41598-025-06493-w.

7. Bu Q, Lyu P, Sun R, Jing J, Lyu Z, Hou S (2025) Fault Diagnosis Method Using CNN-Attention-LSTM for AC/DC Microgrid. Modelling, 6(3):107, DOI:10.3390/modelling6030107.

8. Arévalo P, Cano A, Fedoseienko O, Jurado F (2025) A data-driven approach to microgrid fault detection and classification using Taguchi-optimized CNNs and wavelet transform. Applied Soft Computing, 170:112667, DOI:10.1016/j.asoc.2024.112667.

9. Anwar T, Mu C, Yousaf MZ, Khan W, Khalid S, Hourani AO, Zaitsev I (2025) Robust fault detection and classification in power transmission lines via ensemble machine learning models. Scientific Reports, 15(1):2549, DOI:10.1038/s41598-025-86554-2.

10. Tiwari SP (2025) An efficient protection scheme for critical fault detection in microgrid under uncertain scenarios using deep learning algorithm. Electrical Engineering, 107(7):9009-9023, DOI:10.1007/s00202-024-02861-3.

11. Wang T, Zhang C, Hao Z, Monti A, Ponci F (2023) Data-driven fault detection and isolation in DC microgrids without prior fault data: A transfer learning approach. Applied Energy, 336:120708, DOI:10.1016/j.apenergy.2023.120708.

12. Lee B, Kim Y, Lee H, Kang C (2025) Bidirectional Gated Recurrent Unit Neural Network for Fault Diagnosis and Rapid Maintenance in Medium-Voltage Direct Current Systems. Sensors, 25(3):693, DOI:10.3390/s25030693.

13. Esmaeilbeigi S, Kazemi Karegar H, Najar A (2025) A fault classification scheme based on protective agents for microgrid with parameters impact analysis. Scientific Reports, 15(1):44316, DOI:10.1038/s41598-025-28178-0.

14. Mamodiya U, Sinha D, Kishor I, Jawarneh M, Haj Qasem M, Almegren A (2026) Hybrid deep learning-driven explainable AI framework for fault detection and classification in smart power grids. Scientific Reports, 16:online ahead of print, DOI:10.1038/s41598-026-62771-1.

15. Aqamohammadi AR, Niknam T, Shojaeiyan S, Siano P, Dehghani M (2023) Deep Neural Network with Hilbert-Huang Transform for Smart Fault Detection in Microgrid. Electronics, 12(3):499, DOI:10.3390/electronics12030499.

16. Prashanth Kumar HK, Dash R, Jyotheeswara Reddy K, Atyam NR, Mohapatra B (2025) High resistance fault detection in DC microgrid using Hilbert Huang transform and vector-based ensemble optimized LSTM networks. Scientific Reports, 15(1):16039, DOI:10.1038/s41598-025-00977-5.

17. Dai Y, Wang M, Zhang L, Wang S, Jiang Z (2026) Sliding sampling and successive variational mode decomposition CNN-BiLSTM-attention based fault detection and early warning method for DC microgrid. Scientific Reports, 16:17545, DOI:10.1038/s41598-026-48188-w.

18. Taheri B, Hosseini SA, Hashemi-Dezaki H (2025) Enhanced Fault Detection and Classification in AC Microgrids Through a Combination of Data Processing Techniques and Deep Neural Networks. Sustainability, 17(4):1514, DOI:10.3390/su17041514.

19. Hübner RU, Küsel C, Oestmann JW (2023) Resilience of the critical infrastructure in hospitals: Categorization and quantification as a basis for optimization. Die Anaesthesiologie, 72(10):710-718, DOI:10.1007/s00101-023-01318-9.

20. Khalil M, Ravaghi H, Samhouri D, Abo J, Ali A, Sakr H, Camacho A (2022) What is "hospital resilience"? A scoping review on conceptualization, operationalization, and evaluation. Frontiers in Public Health, 10:1009400, DOI:10.3389/fpubh.2022.1009400.