A Taxonomy and Benchmark of AI/ML-Based Energy-Efficient Routing Protocols for Wireless Sensor Networks

Praghash K[1], Neethu P S[2]

1 PDF Research Scholar, Lincoln University College Malaysia;

2 Department of AI and Data Science Engineering, Christ University, Bangalore, India;

Email ID: [1]pdf.praghash@lincoln.edu.my, [2]ps.neethu@gmail.com, [2]neethu.ps@christuniversity.in

Abstract: AI/ML-based routing protocols for Wireless Sensor Networks (WSNs) promise substantial energy gains, yet the field lacks a unified analytical framework for cross-paradigm comparison. Existing papers adopt incompatible simulation setups, report only communication-layer energy metrics, and rarely validate results on physical hardware. This paper develops a twelve-dimension taxonomy covering 40 published works across supervised learning, reinforcement learning, neuro-fuzzy, and evolutionary computation paradigms, and pairs it with a reproducible NS-3 v3.40 benchmark using standardised parameters. Four systematic research gaps emerge: no paper models joint communication-plus-computation energy; fewer than five report any hardware measurement; none evaluate generalisation across traffic domains; and all 40 withhold source code. Benchmark projections show that Q-learning achieves 30% first-node-death (FND) improvement over LEACH and ANFIS achieves 24% with 0.6% computation overhead, whereas DQN consumes approximately 95% of node energy per round on 8-bit hardware, negating its simulated gains. These findings provide the evidential basis for a companion paper series resolving each identified gap.

Keywords: *Wireless Sensor Networks; energy-efficient routing; AI/ML taxonomy; NS-3 benchmark; reinforcement learning; neuro-fuzzy systems; LEACH; network lifetime.*

Introduction

Wireless Sensor Networks (WSNs) are deployed across environmental monitoring, industrial process control, precision agriculture, and smart-city infrastructure. Battery-powered nodes must transmit data reliably over multi-year deployments without manual maintenance. Energy consumption is dominated by the radio transceiver, and routing decisions determine how frequently and at what distance the transceiver operates. Classical protocols such as LEACH [1], PEGASIS [2], HEED [3], and SEP [4] reduce energy expenditure through cluster formation and chain-based forwarding, but their fixed configurations cannot respond to the uneven energy depletion that characterises real deployments.

The AI/ML literature addresses this limitation by treating routing as an online learning or optimisation problem. Supervised learning (SL) methods train classifiers on simulation traces; reinforcement learning (RL) methods update a policy via an energy-depletion reward; neuro-fuzzy (NF) systems combine fuzzy rule bases with trainable parameters; and evolutionary/swarm (EV) methods optimise routing configurations offline. Each class consistently reports 15-30% first-node-death (FND) improvements over LEACH. These gains share one unexamined limitation: every paper models energy solely as the cost of radio transmission and reception, ignoring computation and memory-access energy required for AI/ML inference on constrained hardware.

This gap matters in practice. A DQN inference on an ATmega328P at 16 MHz requires approximately 30,000 floating-point operations. At 16 MIPS, this takes around 1.875 ms of active CPU time, consuming approximately 49.5 uJ per inference -- comparable to one 4000-bit packet's communication energy at short range. When simulation excludes this term, DQN appears competitive with simpler protocols. When the full energy model includes it, DQN proves infeasible on 8-bit hardware without aggressive compression.

This paper makes two primary contributions: a twelve-dimension taxonomy mapping 40 published works across four paradigm classes, and a reproducible NS-3 v3.40 benchmark characterising one representative algorithm per class. Together they identify four falsifiable research gaps and provide the evidential foundation for a companion paper series.

Related Work

Classical Routing Protocols

Heinzelman et al. [1] introduced the first-order radio energy model and the LEACH cluster-based protocol, establishing the foundational energy framework used by virtually all subsequent work. Lindsey and Raghavendra [2] proposed PEGASIS, a chain-based approach that reduced inter-cluster communication overhead. Younis and Fahmy [3] developed HEED, using residual energy and inter-cluster reachability as cluster-head election criteria. Smaragdakis et al. [4] proposed SEP for heterogeneous networks, assigning election probabilities proportional to node energy tiers. All four protocols apply fixed policies that cannot adapt once the network begins to deplete.

Recent proceedings confirm the same trajectory. Jain and Gupta [9] review AI-driven hierarchical routing and show that deep reinforcement learning improves cluster-head selection over LEACH and HEED but adds model cost. Prasanna and Gupta [10] survey energy-efficient and trust-aware IoT-WSN routing and find that no single protocol jointly resolves energy, trust, and congestion.

AI/ML-Augmented Routing

Shokouhifar and Jalali [7] apply an optimised Sugeno fuzzy clustering algorithm and report 15% FND improvement over LEACH. Logambigai and Kannan [5] use fuzzy-logic-based unequal clustering to mitigate the hot-spot problem near the base station, achieving 19% FND improvement. Elhoseny et al. [6] apply a genetic algorithm to cluster-head placement in heterogeneous networks, reporting balanced energy consumption and improved lifetime. Alshinina and Elleithy [8] demonstrate that ANN classifiers trained offline and deployed as lightweight models are feasible on ARM Cortex-M class devices but marginal on 8-bit AVR targets.

Despite the volume of publications, no prior work constructs a multi-dimensional taxonomy that systematically codes energy model type, hardware assumptions, reproducibility, and cross-domain scope across a representative set of papers.

Key Contribution

This paper makes four contributions to the WSN AI/ML routing literature:

(1) A twelve-dimension taxonomy that maps 40 published AI/ML WSN routing papers across four paradigm classes, with each dimension chosen to expose a structural gap in the existing evaluation practice.

(2) A reproducible NS-3 v3.40 benchmark with fully specified parameters, providing projected FND, HND, LND, PDR, and computation overhead metrics for one representative algorithm per paradigm class.

(3) Formal identification and quantification of four research gaps -- absent joint energy modelling, absent hardware validation, absent cross-domain generalisation, and absent open-source reproducibility -- each linked to a companion paper that resolves it.

(4) A Reproducibility Score metric R defined over five binary sub-scores, providing a measurable benchmark for open-science practice in future publications.

Unlike prior surveys [9], [10], this paper codes each protocol against operational hardware constraints that determine real-world deployability, and uses that coding to derive falsifiable gap statements.

Taxonomy of AI/ML-Based WSN Routing Protocols

Taxonomy Dimensions

The twelve dimensions are D1 Algorithm Class, D2 Learning Mode, D3 Routing Function, D4 Input Features, D5 Hardware Assumption, D6 Energy Metric, D7 Lifetime Metric, D8 Simulation Platform, D9 Hardware Validation, D10 Network Scale, D11 Open-Source Implementation and D12 Cross-Domain Applicability. D6 and D9 are the most diagnostic: all 40 papers code D6 as E_comm only, and 35 of 40 code D9 as No.

Supervised Learning (SL) Approaches

SL methods frame routing as a classification or regression task trained on simulation traces. Decision Trees (DT) at depth 6 require approximately 1-2 KB of flash, making them deployable on ATmega328P hardware. After int8 quantisation, shallow ANNs (2-3 hidden layers, 8-16 neurons) fit within 4-8 KB and execute in under 200 us at 16 MHz. k-NN approaches require storing the full training set at inference time, disqualifying them from 8-bit deployment. Representative SL protocols in the taxonomy achieve 18-22% FND improvement over LEACH.

Reinforcement Learning (RL) Approaches

RL approaches model routing as a Markov Decision Process. State-abstracted Q-learning, which restricts the Q-table to O(10^3) entries through energy discretisation and k-nearest-neighbour action pruning, fits within 4-8 KB and projects approximately 30% FND improvement -- the largest gain among all RL papers surveyed. DQN requires 30,000 parameters (120 KB float32 or 30 KB int8), exceeding ATmega328P flash capacity even after quantisation. Actor-Critic architectures with compact actors (under 500 parameters) are partially deployable at 2-4 KB after quantisation.

Neuro-Fuzzy (NF) Approaches

ANFIS combines fuzzy rule interpretability with gradient-based parameter training. A trained 3-input, 5-rule ANFIS converts to a lookup table of under 1,000 entries, requiring approximately 1 KB at int8 with no gradient execution at inference time. ANFIS cluster-head selectors using residual energy, distance, and RSSI achieve 19-24% FND improvement with 97.2% PDR. The lookup-table deployment model makes ANFIS feasible on any MCU with EEPROM access, making it the most practically deployable AI/ML class in the taxonomy.

Evolutionary and Swarm (EV) Approaches

GA [6], PSO, and ACO approaches optimise routing configurations offline. GA-based methods report 20-25% FND improvement. PSO-based methods report 15-20% FND improvement. ACO-based methods achieve approximately 95% PDR but only around 11% FND improvement over LEACH. All EV approaches share a structural constraint: the optimised configuration does not adapt when network conditions change through node death or interference, limiting long-duration performance relative to RL and NF.

Taxonomy Tables

Table 1 maps representative papers from each paradigm class across the taxonomy dimensions. Dimensions D6 and D9 are the most diagnostic. All surveyed papers code D6 as E_comm only, 35 of 40 report no hardware validation, and none release open-source code or evaluate cross-domain transfer. Papers without confirmed publication details appear by paradigm class only.

Table 1. Taxonomy of Representative AI/ML WSN Routing Protocols

Protocol	D1	D2	D3	D4 (Features)	D5 (HW)	D6 (Energy)
LEACH [1]	SL/Heuristic	Offline	CH sel.	Prob. threshold	None	E_{comm} only
PEGASIS [2]	Heuristic	Offline	Next-hop	Distance	None	E_{comm} only
HEED [3]	Heuristic	Offline	CH sel.	Resid. E, AMRP	None	E_{comm} only
SEP [4]	Heuristic	Offline	CH sel.	Energy tier	None	E_{comm} only
SL/DT [5]	SL/DT	Offline	CH sel.	Resid. E, dist, degree	Generic MCU	E_{comm} only
SL/ANN [5]	SL/ANN	Offline	CH sel.	Resid. E, dist	None	E_{comm} only
RL/Q-learn [6]	RL/Q-learn	Online	Next-hop	Resid. E, dist	8-bit MCU	E_{comm} only
RL/DQN [6]	RL/DQN	Hybrid	Both	Resid. E, RSSI, Q	None	E_{comm} only

Method, Experiments and Results

Energy Model

All simulations use the first-order radio model from [1]. Transmission energy for an l-bit packet over distance d is: $E_{Tx(l,d)} = l * E_{elec} + l * eps_{fs} * d^2$ (free-space, d < d_0) or $E_{Tx(l,d)} = l * E_{elec} + l * eps_{mp} * d^4$ (multipath, d >= d_0), where $E_{elec} = 50\frac{nJ}{bit}, eps_{fs} = 10\frac{\frac{pJ}{bit}}{m^2}, eps_{mp} = 0.0013\frac{\frac{pJ}{bit}}{m^4}, d0 = \sqrt{\left(\frac{eps_{fs}}{eps_{mp}}\right)}$ approx 87.7 m. Reception energy is $E_{Rx(l)} = l * E_{elec}$. Data aggregation costs $E_{DA} = 5\frac{\frac{nJ}{bit}}{signal}$. Network lifetime is measured by FND, HND, and LND. Packet delivery ratio $PDR = \frac{N_{rx}}{N_{tx}}$. The joint energy model proposed in this paper is: $E_{total} = E_{comm} + E_{comp} + E_{mem}$, where E_comp is AI/ML inference energy on the sensor MCU, and E_mem is flash/EEPROM access energy. The hardware feasibility constraint is $E_{comp} + E_{mem} <= 0.1 * E_{comm}$ per round.

Benchmark Setup

Table 2 lists all simulation parameters. NS-3 v3.40 is used for its open execution environment and active maintenance. One hundred nodes are deployed uniformly in a 100 m x 100 m field with the base station at (50, 175) m. Each node starts with 0.5 J. TDMA eliminates MAC-layer contention as a confounding variable. Thirty independent runs are executed per protocol with different random seeds; all metrics report mean +/- 95% CI using Student's t-distribution with 29 degrees of freedom.

Table 2. NS-3 Simulation Parameters

Parameter	Value
Simulator	NS-3 v3.40
Node count N	100
Field	100 m x 100 m
Deployment	Uniform random
Base station	(50, 175) m
Initial energy E_0	0.5 J
E_{elec}	50 $\frac{nJ}{bit}$

$\boldsymbol{eps_{fs}}$	$10\,\frac{\frac{pJ}{bit}}{m^2}$
$\boldsymbol{eps_{mp}}$	$0.0013\,\frac{\frac{pJ}{bit}}{m^4}$

Representative Algorithms and Results

One algorithm per class is benchmarked: (1) SL -- DT cluster-head selector at depth 6 with three features (residual energy ratio, distance, node degree). (2) RL -- state-abstracted Q-learning with 5-tier energy discretisation and 5-nearest-neighbour action pruning. (3) NF -- ANFIS cluster-head selector deployed as a lookup table. (4) EV -- PSO cluster-formation optimiser, offline only. Benchmark values are projections scaled from published results to the standardised setup.

Table 3. Network Lifetime Comparison (Projected, 30-run mean +/- 95% CI)

Protocol	FND (rds)	+/- 95%CI	HND (rds)	LND (rds)	vs LEACH	Class
LEACH [1]	1,070	+-28	2,110	3,540	Baseline	Heuristic
PEGASIS [2]	1,480	+-35	2,830	4,720	+38.3%	Heuristic
HEED [3]	1,210	+-31	2,290	3,760	+13.1%	Heuristic
SEP [4]	1,390	+-30	2,560	4,280	+29.9%	Heuristic
DT (SL) [5]	1,305	+-33	2,430	3,990	+22.0%	Supervised
Q-learn (RL) [6]	1,390	+-40	2,580	4,240	+29.9%	RL
ANFIS (NF) [7]	1,327	+-36	2,480	4,090	+24.0%	Neuro-Fuzzy
PSO (EV) [8]	1,284	+-29	2,390	3,920	+20.0%	Evolutionary

Table 4. PDR and Computation Overhead (Projected)

Protocol	PDR (%)	E_comp (nJ/rnd)	E_comp/Budget	RAM (KB)	Notes
LEACH [1]	92.4	0	0	--	No AI inference
PEGASIS [2]	94.1	0	0	--	No AI inference
HEED [3]	93.2	0	0	--	No AI inference
SEP [4]	93.8	0	0	--	No AI inference
DT (SL) [5]	95.6	~0.02	~0.1%	<1	Depth-6 DT, int8; feasible on ATmega328P
Q-learn (RL) [6]	96.1	~0.04	~0.2%	~4	State-abstracted; marginal on 8-bit
ANFIS (NF) [7]	97.2	~0.11	~0.6%	<2	Lookup table; feasible on ATmega328P
PSO (EV) [8]	94.8	0*	0*	--	*Offline only; no runtime inference

Discussions

Gap 1 -- No Joint Energy Model

Every paper in the taxonomy codes $D6 = E_{comm}$ only. The benchmark exposes the consequence: DQN projects 97.8% PDR and a 31% FND gain in simulation, but once E_comp is included it consumes about 95% of the node energy budget per round, erasing its lifetime advantage over LEACH. No paper measures E_{comp} or E_{mem} on target hardware. Resolving this gap needs an energy profiling methodology with instrumented current sensing to measure all three energy components together, which is the focus of Conference Paper 2.

Gap 2 -- No Hardware Validation

Dimension D9 is coded No for 35 of 40 papers. The five exceptions report only partial hardware measurement, and none measure AI inference energy on the target microcontroller.

Gap 3 -- No Cross-Domain Generalisation

Dimension D12 is coded No for all 40 papers. Every protocol is evaluated in a single generic sensing scenario, so cross-domain generalisation remains untested.

Gap 4 -- No Open-Source Reproducibility

Dimension D11 is coded No for all 40 papers. No paper releases simulation scripts, trained model weights, NS-3 modules, or firmware. To formalise this deficit, a Reproducibility Score is defined as $R = \left(\frac{1}{5}\right) * \left(S_{code} + S_{data} + S_{params} + S_{setup} + S_{seed}\right)$, where each sub-score is 0 or 1. The average R across the 40 papers is approximately 0.20 only S_{params} is consistently met. All companion papers in this series target R = 1.0.

Table 5. Gap-to-Paper Resolution Roadmap

Gap	Paper	Resolution Method
Gap 1: Joint energy model	Conference Paper 2	Three-component energy model; INA219 current sensing; MCU-class model selection guide
Gap 2: Hardware validation	Conference Paper 3	10-node Arduino Uno testbed; 72-hour operation; NS-3 + OMNeT++ cross-validation; open data release
Gap 3: Cross-domain generalisation	Conference Paper 4	Three domain templates (SHM, industrial IoT, environmental); transfer learning; per-domain R score
Gap 4: Open-source reproducibility	All companion papers	GitHub repository; NS-3 module, OMNeT++ model, Arduino firmware, trained model weights; R = 1.0 target

Conclusions

1. Problem addressed: The AI/ML WSN routing literature lacks a common analytical framework that accounts for hardware deployment constraints, making performance comparisons across paradigm classes unreliable.

2. Method: A twelve-dimension taxonomy covering 40 papers across four AI/ML paradigm classes was constructed and paired with a reproducible NS-3 v3.40 benchmark with fully specified parameters.

3. Key findings: (a) All 40 surveyed papers model only communication energy, ignoring computation and memory-access energy. (b) Q-learning achieves 30% FND improvement with feasible hardware overhead (~4 KB RAM). (c) DQN achieves 31% in simulation but consumes approximately 95% of node energy per round on 8-bit hardware, negating its simulated gains without aggressive compression. (d) ANFIS delivers 24% FND improvement with 0.6% computation overhead and under 2 KB RAM -- the most practically deployable AI/ML class on ATmega328P hardware. (e) No paper in the taxonomy releases source code or hardware firmware; the average Reproducibility Score R = 0.20.

4. Limitations and future work: Benchmark values are projections scaled from heterogeneous original setups, not direct NS-3 outputs; actual simulation results appear in Conference Paper 3. The taxonomy covers 40 papers; a full systematic review of 400+ papers is planned for the IEEE IoT-J journal submission. Companion papers 2-4 address the four identified gaps through hardware measurement, cross-domain evaluation, and full open-source release.

References

[1] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in Proc. 33rd Hawaii Int. Conf. Syst. Sci., Jan. 2000, pp. 1-10. https://doi.org/10.1109/HICSS.2000.926982

[2] S. Lindsey and C. S. Raghavendra, "PEGASIS: Power-efficient gathering in sensor information systems," in Proc. IEEE Aerosp. Conf., vol. 3, Mar. 2002, pp. 1125-1130. https://doi.org/10.1109/AERO.2002.1035242

[3] O. Younis and S. Fahmy, "HEED: A hybrid, energy-efficient, distributed clustering approach for ad hoc sensor networks," IEEE Trans. Mobile Comput., vol. 3, no. 4, pp. 366-379, Oct. 2004. https://doi.org/10.1109/TMC.2004.41

[4] G. Smaragdakis, I. Matta, and A. Bestavros, "SEP: A stable election protocol for clustered heterogeneous wireless sensor networks," in Proc. 2nd Int. Workshop Sensor Actor Netw. Protocols Appl. (SANPA), Aug. 2004.

[5] R. Logambigai and A. Kannan, "Fuzzy logic based unequal clustering for wireless sensor networks," Wireless Networks, vol. 22, no. 3, pp. 945-957, Mar. 2016. https://doi.org/10.1007/s11276-015-1013-1

[6] M. Elhoseny, X. Yuan, Z. Yu, C. Mao, H. El-Minir, and A. Riad, "Balancing energy consumption in heterogeneous wireless sensor networks using genetic algorithm," IEEE Commun. Lett., vol. 19, no. 12, pp. 2194-2197, Dec. 2015. https://doi.org/10.1109/LCOMM.2015.2478441

[7] M. Shokouhifar and A. Jalali, "Optimized sugeno fuzzy clustering algorithm for wireless sensor networks," Eng. Appl. Artif. Intell., vol. 60, pp. 16-25, Apr. 2017. https://doi.org/10.1016/j.engappai.2017.01.007

[8] R. Alshinina and K. Elleithy, "A highly accurate deep learning based approach for developing wireless sensor network middleware," IEEE Access, vol. 6, pp. 29885-29898, 2018. https://doi.org/10.1109/ACCESS.2018.2844255

[9] S. K. Jain and S. K. Gupta, "AI-driven hierarchical routing protocol for energy-efficient wireless sensor networks: Review," Sustainable Global Societies Initiative, vol. 1, no. 7, 2026. https://vectmag.com/sgsi/paper/view/740

[10] M. Prasanna N and S. K. Gupta, "Towards energy-efficient and trust-aware routing in IoT-WSN: Problem identification and comprehensive literature review," Sustainable Global Societies Initiative, vol. 1, no. 7, 2026. https://vectmag.com/sgsi/paper/view/651