International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 3
Article Received: *25 September 2026* ***Revised:*** *12 October 2026* ***Accepted:*** *22 November 2026*

A Critical Review of Metaheuristic Approaches For Fog Node Placement in the Fog–Cloud Continuum

Apoorva Sharma

Assistant Professor, Department of Computer Science, Poornima University, Jaipur (Rajasthan)

Email: apoorva.sharma1@poornima.edu.in

Abstract

Fog computing extends cloud capabilities towards end devices by introducing geographically distributed compute, storage, and networking resources. The placement of fog nodes is a foundational infrastructure decision because it influences device coverage, network connectivity, communication delay, energy consumption, deployment expenditure, and resilience. The resulting optimisation problem is combinatorial and becomes increasingly difficult when heterogeneous capacities, multiple objectives, and dynamic workloads are considered. This paper presents a critical structured narrative review of metaheuristic approaches relevant to fog node placement and closely related edge/fog placement problems. It distinguishes physical fog-node placement from fog service placement and edge-server placement, develops a common location–assignment model, and organises the literature into evolutionary, swarm-intelligence, hybrid, and learning-assisted approaches. A six-dimensional evaluation framework is proposed, covering solution quality, computational effort, scalability, multi-objective capability, adaptability, and reproducibility. The reviewed evidence indicates that Pareto-based evolutionary methods are appropriate when trade-offs among delay, energy, and deployment cost must be exposed; swarm methods have been applied directly to coverage- and connectivity-oriented fog deployment; hybrid methods combine complementary search behaviours; and learning-assisted approaches are promising for adaptive resource management, although much of their evidence concerns service placement rather than physical node siting. The literature does not support a universal ranking of algorithm families because studies use different decision variables, constraints, benchmarks, and performance indicators. The review therefore recommends a reproducible benchmark protocol, explicit separation of placement problem types, repeated-run statistical analysis, and evaluation under heterogeneous and time-varying conditions. This paper is a literature synthesis and reports no new experimental results.

Keywords: fog computing; fog node placement; edge computing; metaheuristic optimisation; swarm intelligence; multi-objective optimisation; IoT; benchmarking

1. Introduction

The growth of Internet of Things (IoT) applications has increased demand for computation close to data sources. Cloud data centres provide scalable processing and storage, but communication with distant centralised facilities may introduce delay, transport traffic, and dependence on wide-area connectivity. Fog and edge computing complement the cloud by distributing computational resources across intermediate locations closer to users and devices. The benefits of this architecture depend not only on the availability of fog resources but also on where those resources are deployed and how they are connected. The conceptual IoT–fog–cloud arrangement considered in this review is illustrated in Fig. 1.

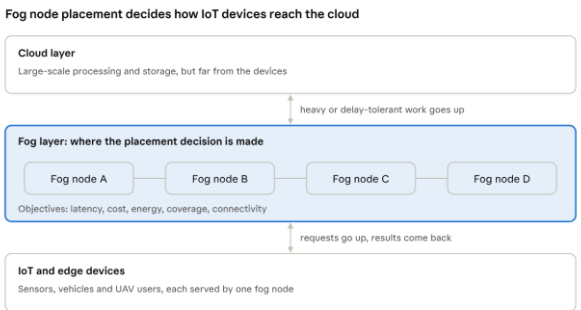

Figure 1. Three-layer IoT–fog–cloud architecture. Fog node placement determines which devices each node serves and how fog nodes connect. Source: Original conceptual figure from the source manuscript; retained in this revision.

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 3
Article Received: *25 September 2026* ***Revised:*** *12 October 2026* ***Accepted:*** *22 November 2026*

Fog node placement is an infrastructure-level decision: it selects a subset of candidate physical or logical sites at which fog nodes are installed. It affects which devices can be served, the topology connecting fog nodes, the capacity available at each location, and deployment cost. It must be distinguished from fog service placement, which maps application services or components to already available nodes, and from edge-server placement, which selects edge stations or facilities in a related but not identical infrastructure model. Existing surveys document broad variation in edge/fog placement methods, parameters, and constraints [1], while recent surveys separately examine fog service placement [2] and microservice placement [3].

These distinctions are essential when interpreting algorithmic evidence. A method that optimises software-service allocation over a fixed set of fog servers does not directly demonstrate effectiveness in selecting physical fog-node locations. Edge-server placement may provide transferable location-allocation formulations, but its assumptions about candidate sites, radio access, demand distribution, and infrastructure costs must be examined before its results are generalised to fog deployments.

Direct evidence on fog node deployment includes the use of the Marine Predators Algorithm (MPA) for jointly improving device coverage and network connectivity [4]. UAV-assisted fog deployment has also been studied for rescue scenarios, where adaptive placement is required under changing network conditions [5]. Related work applies NSGA-II to fog service placement [6], whale optimisation to IoT service placement [7], and grey wolf optimisation to autonomic service placement [8]. Hybrid GA–PSO methods have been evaluated for edge-server placement [9], while multi-objective evolutionary optimisation has been used for facility location in vehicular fog networks [10]. These studies are relevant to the broader placement landscape, but their decision scopes are not interchangeable.

The central gap addressed here is therefore methodological and comparative: the literature contains multiple algorithms and related placement formulations, but evidence is fragmented across problem definitions, objective functions, simulation settings, baselines, and reporting practices. A review that treats all placement papers as directly comparable risks overstating the evidence. This paper instead evaluates the strength and transferability of each study and proposes a common protocol for future comparisons. This distinction is consistent with prior surveys that report heterogeneous placement definitions and evaluation practices [1], [2], [19].

1.1 Contributions

Defines the boundary between physical fog-node placement, fog service placement, and edge-server placement.

Formulates a general location–assignment model with capacity, coverage, connectivity, and budget constraints.

Classifies relevant approaches into evolutionary, swarm-based, hybrid, and learning-assisted families.

Compares representative studies by decision scope, objectives, reported evaluation evidence, and limitations.

Proposes a reproducible experimental protocol for comparing PSO, GWO, MPA, GA/NSGA-II, and hybrid variants.

2. Review Methodology

This study adopts a targeted structured narrative-review design. It is not presented as a PRISMA systematic review, because a fully auditable screening log, duplicate-removal record, and database-specific yield counts were not maintained for the source manuscript. The purpose of the method is to make the search scope, query concepts, eligibility rules, and synthesis procedure explicit without implying a systematic-review process that was not performed. The reporting distinction is deliberate: PRISMA 2020 is a reporting guideline for systematic reviews, whereas this manuscript is explicitly a targeted narrative synthesis and does not claim PRISMA-compliant screening [19].

2.1 Review questions

RQ1: Which metaheuristic families have been applied to physical fog-node placement and closely related edge/fog placement problems?

RQ2: Which objective functions and constraints are used, and how do they differ between node placement and service placement?

RQ3: What evaluation environments, datasets or traces, baselines, and performance indicators are reported?

RQ4: What limitations prevent reliable cross-study comparison, and what benchmark protocol should guide future work?

2.2 Information sources and search strategy

The source set was assembled from bibliographic records and publisher-hosted material associated with IEEE Xplore, Scopus-indexed journal records, ScienceDirect,

SpringerLink, Wiley Online Library, MDPI, EAI Publications, PubMed Central, arXiv, and institutional repositories. Publisher pages and DOI landing pages were prioritised for bibliographic verification; preprints and repository copies were used to identify early versions or accessible technical descriptions. The search concepts combined infrastructure terms, placement terms, and optimisation terms. The use of publisher records and DOI landing pages is intended to prioritise version-of-record metadata, while preprints are labelled separately.

The following representative search strings define the search strategy and can be rerun in each database (with syntax adapted to the platform):

("fog node" OR "fog nodes" OR "fog deployment") AND (placement OR deployment OR location) AND (metaheuristic OR optimisation OR optimization)

("fog computing" OR "fog-cloud" OR "edge-fog") AND ("node placement" OR "server placement" OR "facility location")

("fog service placement" OR "IoT service placement" OR "application placement") AND (GA OR NSGA-II OR PSO OR GWO OR WOA OR MPA OR hybrid)

("edge server placement" OR "vehicular fog" OR "UAV fog") AND (genetic OR swarm OR evolutionary OR reinforcement learning)

("fog placement" OR "edge placement") AND (latency OR energy OR coverage OR connectivity OR cost OR load balancing)

The review window was 2021–2026, matching the scope of the source manuscript. Earlier foundational works were retained only where needed to explain a model or algorithmic lineage. The search cut-off for this revision is 26 November 2026. Because database-specific search exports and complete hit counts were not available in the source material, no numerical claim is made about the total number of records retrieved or screened.

2.3 Eligibility criteria

Criterion	Included	Excluded or treated as contextual
Publication period	2021–2026; earlier foundational work where essential	Outside the window without a clear foundational role
Topic	Fog node siting/deployment; edge/fog server placement; fog service/application placement where algorithmically relevant	General cloud scheduling without a placement decision
Method	Evolutionary, swarm, hybrid, mathematical optimisation, or learning-assisted placement/resource decisions	Pure architecture papers without an optimisation method
Evidence	Peer-reviewed journal/conference records; identifiable preprints where publication status is stated	Untraceable records or incomplete metadata presented as final publications
Problem scope	Decision variables and objectives can be identified	Studies whose placement meaning cannot be distinguished from generic scheduling

Studies were not treated as equally direct evidence. Each retained work was assigned a scope label: (i) direct fog-node placement, (ii) edge/vehicular/UAV infrastructure placement, (iii) fog service/application placement, or (iv) adjacent resource management. This scope label governs how strongly its results are used in the comparative discussion.

2.4 Screening and data extraction

Screening proceeded in three conceptual passes: title and abstract relevance; inspection of the available full text, publisher record, or repository description; and eligibility assessment against the criteria above. For each study, the synthesis template records publication details, placement scope, algorithm, decision variables,

objectives, constraints, evaluation platform or trace if explicitly reported, comparator algorithms, reported metrics, and stated limitations. Missing items are recorded as not reported rather than inferred.

2.5 Synthesis and limitations of the review method

The evidence was synthesised narratively and grouped by search paradigm and placement scope. No pooled effect size or cross-paper numerical ranking was computed because the studies do not share a common benchmark, workload, topology, or metric normalisation. The review is limited by its targeted rather than exhaustive search, the uneven availability of full texts, and the fact that some 2026 records may be preprints or early online publications. These limitations should be considered when interpreting the taxonomy and comparative judgements. Accordingly, the synthesis avoids a pooled effect estimate and follows the caution raised by evaluation surveys regarding cross-study comparability [19], [20].

3. Problem Scope and Formal Model

3.1 Placement problem distinctions

Problem type	Primary decision	Typical objectives	Evidence status in this review
Fog node placement	Which candidate sites receive physical/logical fog nodes?	Coverage, connectivity, deployment cost, latency, energy, resilience	Direct target problem
Fog service placement	Which existing fog/cloud node hosts each service or component?	Response time, resource utilisation, service cost, SLA, energy	Related but not equivalent
Edge-server placement	Where should edge servers be installed in an access/network infrastructure?	Latency, transport load, QoS/QoE, capacity, cost	Transferable location-allocation evidence
Dynamic/UAV/vehicular placement	Where should mobile or temporary fog resources be positioned over time?	Coverage, connectivity, delay, energy, network lifetime	Specialised dynamic placement case

This separation prevents a common comparison error: an algorithm may perform well in service assignment while the underlying physical node locations are fixed. The algorithm family may still be relevant, but the evidence must be described as transfer evidence rather than direct node-placement validation.

3.2 General location–assignment formulation

Let $D=\{1,\ldots,|D|\}$ denote end devices and $S=\{1,\ldots,|S|\}$ denote candidate fog sites. Let $x_s \in \{0,1\}$ indicate whether a fog node is activated at site s, and $y_{ds} \in \{0,1\}$ indicate whether device d is assigned to site s. Let C_s be the service capacity of site s, λ_d the offered demand of device d, a_{ds} indicate whether site s can cover device d, and c_s the installation cost of site s. Let b_d be the demand-specific latency budget and L_{ds} the estimated end-to-end latency between device d and site s.

A basic feasibility model is:

Each device is assigned once: $\Sigma(s \in S)\ y_{ds} = 1,\ \forall d \in D.$

Assignment requires an active node: $y_{ds} \le x_s,\ \forall d \in D,\ s \in S.$

Capacity is respected: $\Sigma(d \in D)\ \lambda_d\ y_{ds} \le C_s\ x_s,\ \forall s \in S.$

Coverage is respected: $y_{ds} \le a_{ds},\ \forall d \in D,\ s \in S.$

Installation budget: $\Sigma(s \in S)\ c_s\ x_s \le B$, where B is the deployment budget.

Optional node-count bound: $\Sigma(s \in S)\ x_s \le K.$

The generic multi-objective vector can be written as minimise $F(x,y)=[f_{latency}, f_{cost}, f_{energy}, f_{imbalance}]$ subject to the feasibility constraints. For coverage-oriented variants, the objective may instead maximise the proportion of devices covered by active nodes. Connectivity can be represented through a graph $G=(V,E)$ induced by active fog nodes and feasible

communication links; the model may require all active nodes to belong to one connected component or penalise disconnected components. Multi-objective delay–energy formulations for vehicular fog facility location provide a concrete example of such trade-offs [10].

3.3 Objective-function definitions

A representative latency objective is f_latency = Σ(d∈D)Σ(s∈S) w_d L_ds y_ds, where w_d is the importance or request rate of device d. If processing delay is modelled separately, L_ds may include access, transmission, queueing, and processing components. A deployment-cost objective is f_cost = Σ(s∈S)c_s x_s + Σ(d,s) m_ds y_ds, where m_ds represents assignment or operational cost. An energy objective may combine node activation and workload-dependent energy: f_energy = Σ_s(P_s^idle x_s + P_s^dyn u_s), where u_s is the utilisation-dependent component. A load-imbalance objective can be expressed as the variance of normalised node utilisation ρ_s, for example f_imbalance = (1/|S'|)Σ(s∈S')(ρ_s−ρ̄)^2 over active sites S'.

These expressions are generic modelling templates, not a claim that every reviewed study uses the same equations. Their purpose is to expose the decision variables and objective components that a common benchmark should standardise. In a practical model, latency, energy, and cost must be normalised before weighted aggregation; alternatively, Pareto-based methods can preserve trade-offs without selecting weights in advance.

3.4 Complexity and modelling assumptions

The model combines binary site activation and assignment decisions with capacity, coverage, and network-topology constraints. Such location–allocation formulations are combinatorial and may become computationally intractable as candidate sites and devices increase. Exact optimisation remains valuable for small instances and for generating lower bounds or reference solutions. Metaheuristics are attractive for larger instances because they can explore discrete and non-convex search spaces, but they do not guarantee global optimality and their results depend on encoding, constraint handling, parameterisation, and random seed. The overall search-and-evaluation cycle followed by population-based placement methods is summarised in Fig. 2. Location-allocation formulations and the distinction between infrastructure placement and application placement are also discussed in the placement literature [1], [2], [12]. The combinatorial nature of placement is consistent with the location-allocation formulations studied for edge/fog infrastructures [1], [12].

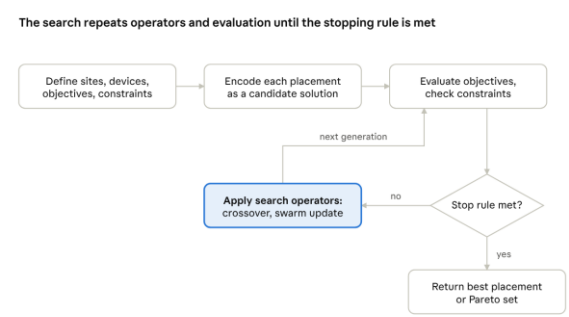

Figure 2. Generic metaheuristic workflow for fog node placement: candidate placements are encoded, evaluated, improved through search operators, and iterated until a stopping criterion is met. Source: Original conceptual figure from the source manuscript; retained in this revision.

4. Taxonomy of Metaheuristic Approaches

The reviewed methods are classified along three independent dimensions: search paradigm, objective handling, and placement decision scope. This avoids classifying a paper solely by its algorithm name when its decision problem may differ substantially. The resulting algorithm-family taxonomy is presented in Fig. 3. The taxonomy is aligned with the algorithmic diversity reported in fog-placement surveys and genetic-optimisation reviews [1], [2], [16], [17].

4.1 Search paradigm

Family	Representative methods	Technical rationale	Typical caution
Evolutionary	GA, NSGA-II, MOEA/D	Population evolution, selection, crossover, mutation; Pareto ranking and diversity preservation support conflicting objectives	Encoding and constraint repair can dominate performance; computational cost may be high

Swarm intelligence	PSO, GWO, WOA, MPA	Collective search guided by social, leader, or movement-inspired updates; adaptable to continuous and encoded discrete spaces	Premature convergence, parameter sensitivity, and discretisation effects
Hybrid	GA+PSO, NSGA-II+SMPSO, PSO+CRO	Combines complementary exploration, exploitation, or local-search mechanisms	More parameters and difficulty attributing performance gains
Learning-assisted	DRL with search, Q-learning-based selection, game-theoretic search	Learns policies or algorithm/operator selection under changing states and workloads	Training cost, state/action design, reproducibility, and generalisation

4.2 Objective handling

Single-objective formulations commonly combine multiple criteria using a weighted sum. This is computationally convenient but requires weights and can conceal trade-offs. Pareto-based multi-objective methods return a set of non-dominated solutions and are more informative when operators must choose between latency, energy, and cost. Decomposition-based methods divide a many-objective problem into scalar subproblems. The selected approach should match the decision context and the number of objectives, rather than being treated as an intrinsic superiority of one algorithm family.

4.3 Decision scope

The decision scope may be offline site selection, dynamic re-placement under changing demand or mobility, or joint placement and allocation. Offline planning is suitable for static or slowly changing demand forecasts. Dynamic placement introduces state transitions, migration or repositioning costs, and the need to avoid oscillatory decisions. Joint placement-allocation models are richer but should report which decisions are optimised simultaneously and which are fixed.

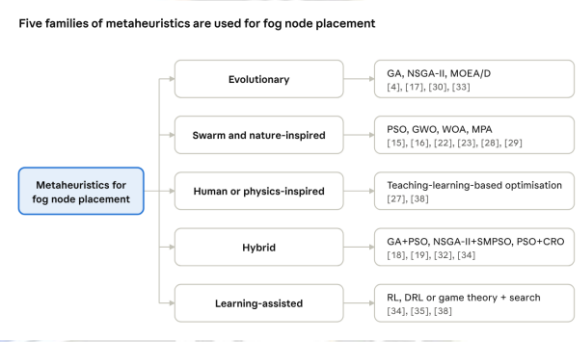

Figure 3. Taxonomy of metaheuristic families for fog node placement and related placement studies. Source: Original conceptual figure from the source manuscript; retained in this revision.

5. Critical Review of Representative Studies

The following synthesis distinguishes direct fog-node deployment from adjacent placement evidence. The available publisher records and abstracts support the problem scope, principal objectives, and some evaluation details. Where a simulator, dataset, or baseline is not stated in the accessible record, it is explicitly marked as not established here rather than inferred. Direct fog-node deployment, UAV-assisted fog deployment, service placement, and edge-server placement are therefore evaluated as separate evidence classes [4]–[9], [10]–[14].

Study	**Problem scope and method**	**Objectives / evidence reported**	**Critical interpretation**
Naouri et al. (2024) [4]	Direct fog-node placement; MPA	Maximises device coverage and fog-node connectivity; tests multiple network	Most direct evidence for the target problem. Weighted aggregation uses a user-

		configurations; compares against other methods as described by the article	defined weight; generalisation depends on topology and parameter settings.
Naouri et al. (2024) [5]	UAV fog deployment for rescue; adaptive WOA plus node-swapping	Dynamic rescue/evacuation setting; coverage/connectivity and network lifetime subproblems; rapid convergence and low computational demand are reported	Highly relevant dynamic infrastructure case, but UAV mobility and rescue assumptions limit direct transfer to static campus/urban fog deployment.
Niu et al. (2022) [6]	Fog service placement; NSGA-II	Constrained multi-objective optimisation of resource utilisation, network latency, and service-placement cost	Useful evidence for Pareto-based service allocation, not physical node siting. Must be labelled adjacent evidence.
Ghobaei-Arani & Shahidinejad (2022) [7]	IoT service placement in fog; WOA	Cost-efficient service placement with heterogeneous service QoS requirements	Supports WOA for service placement; does not alone establish superiority for physical fog-node deployment.
Salimian et al. (2021) [8]	Autonomic IoT service placement; GWO	GWO-based autonomic placement in fog; detailed objective/baseline data should be checked in full text before numerical comparison	Relevant swarm-intelligence evidence, but service-placement scope must remain explicit.
Hussain et al. (2023) [10]	Vehicular fog facility location; SONG (NSGA-II + SMPSO)	Pareto trade-off between service delay and energy; real vehicular traces; HV, IGD, and CPU delay-gap indicators; compared with NSGA-II and SMPSO	Strong multi-objective evaluation practice. It concerns vehicular fog facility location and should not be conflated with generic static fog siting.
Han et al. (2024) [9]	Edge-server placement; hybrid GA–PSO	Edge-server placement; hybrid GA and PSO	Useful transferable edge-placement evidence; physical fog-node equivalence is not established.
Ye et al. (2023) [11]	Edge-server deployment; GA	Balances server workload and communication delay; Journal of Cloud Computing article 148	Related infrastructure placement evidence; report as edge-server deployment.

Lähderanta et al. (2021) [12]	Edge server placement; capacitated location allocation	Latency-focused placement; evaluates MEC and fog scenarios with a large-scale real-world dataset; open-source R implementation	Valuable problem-specific and real-data comparator; not itself evidence of a metaheuristic family.
Wang et al. (2025) [13]	Resource management across edge/fog/cloud; DRL framework	Distributed and centralised DRL resource-management capabilities	Learning-assisted resource management is relevant to dynamic adaptation, but its placement decision scope must be checked before treating as node-siting evidence.
Sulimani et al. (2022) [14]	Decentralised fog service placement; reinforcement optimisation	iFogSim evaluation; service latency, network usage, and computing usage; comparisons with established algorithms	Reports a simulator and metrics; service placement, not physical node deployment.
Sahu et al. (2025) [15]	Edge resource scheduling; cellular Potts + PSO	Energy and latency optimisation in edge computing	Adjacent scheduling/resource-management evidence; exclude from direct node-placement counts.

The study-level evidence demonstrates why a single cross-paper performance table is not defensible without normalisation. For example, the MPA study [4] optimises coverage and connectivity in a physical deployment graph, whereas NSGA-II service placement [6] optimises resource utilisation, network latency, and placement cost over existing fog resources. SONG [10] uses real vehicular traces and Pareto indicators for delay–energy facility location. These are complementary contributions, not directly comparable experimental runs.

5.1 Algorithm-wise interpretation

Evolutionary methods are particularly appropriate when a decision-maker needs a diverse Pareto set for conflicting objectives. NSGA-II-based service placement and the SONG vehicular-fog model demonstrate the use of non-dominated sorting and diversity mechanisms [6], [10]. However, service-placement findings should not be described as direct evidence of physical fog-node siting. Representative fog and vehicular-fog applications include NSGA-II-based service placement and hybrid Pareto optimisation [6], [10], [16].

Swarm methods have a more direct connection to coverage/connectivity-oriented deployment. MPA has been applied to physical fog-node deployment [4], while adaptive WOA has been used for UAV fog deployment in rescue scenarios [5]. WOA and GWO also appear in service-placement studies [7], [8]. These results support the relevance of swarm methods but do not establish a universal ranking among PSO, GWO, WOA, and MPA.

Hybrid methods aim to balance global exploration and local refinement. SONG combines NSGA-II and SMPSO and evaluates Pareto quality against its constituent algorithms on vehicular traces [10]. GP4ESP combines GA and PSO for edge-server placement [9]. Hybridisation can improve search behaviour, but any gain must be tested against tuned standalone baselines under equal computational budgets. Examples include SONG and GA–PSO edge-server placement [9], [10].

Learning-assisted methods are attractive when workload and topology change over time. ReinFog provides a DRL framework for resource management across edge/fog/cloud [13], while decentralised reinforcement optimisation has been studied for fog service placement using iFogSim [14]. Such work is promising for adaptive

orchestration, but training overhead, policy generalisation, and decision-scope differences remain important limitations. Representative edge/fog resource and service-management studies include [13], [14].

6. Evaluation Framework and Comparative Analysis

6.1 Evaluation criteria

Criterion	Recommended measures	Reporting requirement
Solution quality	Coverage ratio; connectivity ratio; mean/95th-percentile latency; deployment cost; energy; load imbalance	Report objective definitions, units, and feasibility rate
Convergence and runtime	Wall-clock runtime; objective vs. iteration; evaluations to target	Use equal stopping rules and computational budgets
Scalability	Performance as devices, sites, demand, and constraints increase	Use multiple instance sizes and disclose hardware
Multi-objective quality	Hypervolume (HV), IGD/IGD+, spread/diversity, epsilon indicator	State reference point/front and normalisation procedure
Dynamism	Recovery time after demand/mobility change; migration/repositioning cost; SLA violations	Use time-varying workloads and mobility traces
Reproducibility	Random seeds, parameter settings, source code, topology/workload data, statistical tests	Provide repeated runs and confidence intervals

6.2 Evidence-based comparison

The available studies report heterogeneous evidence. The MPA fog-node deployment work evaluates coverage and connectivity across generated network configurations [4]. The UAV fog study targets dynamic rescue deployment using adaptive WOA [5]. SONG reports real vehicular traces and multi-objective indicators including hypervolume and inverted generational distance, and compares its hybrid method with NSGA-II and SMPSO [10]. The capacitated location-allocation study uses a large-scale real-world dataset and provides an open-source implementation [12]. The decentralised reinforcement-optimisation study reports iFogSim-based evaluation of latency, network usage, and computing usage [14]. These are useful evidence types, but the underlying tasks and metrics differ. This variation in task scope and metrics is also documented in surveys of fog application placement and edge/fog placement [1], [2], [19].

Accordingly, the family-level comparison below is qualitative and is justified by the type of evidence available, not by a pooled numerical score.

Approach	Evidence-supported strengths	Limitations / evidence gaps	Appropriate use case
Evolutionary / NSGA-II	Native Pareto search; explicit multi-objective evidence in service placement and vehicular fog facility location [6], [10]	May be costly; some studies optimise services rather than physical sites; results depend on encoding and population settings	Offline planning with multiple conflicting objectives

Swarm (MPA/WOA/GWO/PSO)	Direct MPA evidence for coverage/connectivity [4]; adaptive WOA for UAV fog deployment [5]; related WOA/GWO service-placement studies [7], [8]	Premature convergence and parameter sensitivity; many results are not directly comparable; some are service-level	Coverage/connectivity and rapid search, subject to benchmark validation
Hybrid	Combines complementary search operators; SONG reports HV/IGD comparisons [10]; GA–PSO applied to edge placement [9]	Extra parameters and computational overhead; improvements need equal-budget tuned baselines	Complex or multi-objective placement where single methods show weaknesses
Learning-assisted	Can adapt to changing states; DRL and reinforcement methods have been evaluated in edge/fog resource management [13], [14]	Training cost, reproducibility, policy transfer, and physical node-placement evidence remain limited	Dynamic service/resource management; potential future dynamic node placement

6.3 Statistical and benchmarking recommendations

A future benchmark should execute each stochastic algorithm over at least 30 independent seeds where computationally feasible, report mean, standard deviation, median, and interquartile range, and use non-parametric statistical tests when distributional assumptions are not justified. Pairwise comparisons should include effect sizes and adjusted p-values when multiple algorithms are tested. For multi-objective algorithms, HV and IGD/IGD+ should be computed using a common normalisation range and a clearly defined reference point. Runtime comparisons should use identical hardware, stopping criteria, and function-evaluation budgets.

No numerical winner is declared in this review. The original studies do not share a common topology, workload, objective scaling, or computational budget; any cross-paper ranking would therefore be methodologically unsound.

7. Discussion and Research Gaps

7.1 Research gaps

Problem-scope inconsistency: papers often use the term placement for node siting, service mapping, application deployment, and task allocation. Future studies should define the decision variables explicitly.

Benchmark fragmentation: many studies use independently generated topologies and workloads, making cross-paper comparison difficult.

Limited joint modelling: latency, energy, cost, coverage, connectivity, reliability, security, and load balance are rarely integrated in one validated physical node-placement model.

Insufficient dynamic evaluation: mobility, workload bursts, node failures, and repeated re-placement are not consistently represented.

Limited reproducibility: complete parameter settings, random seeds, code, and benchmark instances are not always available.

Algorithmic fairness: comparisons can be biased when algorithms receive unequal tuning effort, stopping budgets, or baseline implementations.

7.2 Recommended comparative experimental framework

A focused follow-on study should compare GA/NSGA-II, PSO, GWO, and MPA on the same physical fog-node placement formulation. The first stage should define candidate sites, device demand, node capacities, communication links, coverage radii, and installation costs. The second stage should implement a common discrete encoding and a shared feasibility-repair mechanism. The third stage should evaluate each algorithm under identical population sizes, function-evaluation budgets, and stopping criteria. The fourth

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 3
Article Received: *25 September 2026* ***Revised:*** *12 October 2026* ***Accepted:*** *22 November 2026*

stage should test small, medium, and large instances under both static and time-varying workloads. The fifth stage should report coverage, connectivity, latency, energy, cost, load imbalance, runtime, and feasibility rate. Finally, statistical tests and an open benchmark package should support reproducibility.

Experimental component	Recommended design
Algorithms	GA/NSGA-II, PSO, GWO, MPA; optional hybrid variants only after standalone baselines
Problem variants	Static node siting; capacity-constrained siting; coverage/connectivity; dynamic demand
Instance design	Multiple site/device scales; heterogeneous capacities; sparse/dense topologies
Baselines	Exact/MILP on small instances; greedy/location-allocation heuristic; tuned algorithm baselines
Metrics	Coverage, connectivity, latency, energy, cost, imbalance, runtime, feasibility, HV/IGD
Statistics	≥30 seeds where feasible; confidence intervals; Friedman/Wilcoxon with correction; effect sizes
Reproducibility	Publish code, configurations, seeds, topology generator, workload traces, and hardware details

8. Practical Implications

For static deployment planning with several conflicting objectives, a Pareto-based evolutionary method is a defensible candidate because it exposes trade-offs rather than hiding them in a single weighted score. For coverage- and connectivity-oriented physical deployment, MPA and other swarm approaches deserve direct evaluation because MPA has been applied to this specific problem [4]. For mobile or emergency scenarios, adaptive placement methods such as the UAV-focused WOA approach are relevant [5]. For dynamic service allocation, reinforcement-learning methods may be appropriate, but their success should not be assumed to transfer automatically to physical node siting [13], [14]. The use of Pareto-based methods is supported by multi-objective fog-service and vehicular-fog facility-location studies [6], [10], while direct coverage/connectivity evidence exists for MPA-based fog-node deployment [4].

In practice, algorithm selection should follow the decision problem, not popularity. A deployment operator should first identify whether the task is node siting, service placement, or resource scheduling; then select objectives and constraints; and only then choose an optimisation family. Small instances should be solved exactly where possible to establish reference solutions.

9. Conclusion

This paper has critically reviewed metaheuristic approaches relevant to fog node placement and adjacent edge/fog placement problems. The review separates physical fog-node deployment from service placement and edge-server placement, formalises a general location–assignment model, and classifies the literature into evolutionary, swarm-based, hybrid, and learning-assisted approaches. The evidence indicates that MPA has been directly applied to coverage- and connectivity-oriented fog-node deployment, adaptive WOA has been studied for UAV fog deployment, and evolutionary and hybrid methods have been used in service and vehicular-fog facility-placement settings. Learning-assisted approaches provide promising mechanisms for dynamic resource management, although much of the available evidence concerns service placement rather than physical node siting. The conclusions are bounded by the evidence scope and the evaluation limitations identified in prior fog-placement surveys.

The principal conclusion is not that one algorithm family is universally superior. Rather, algorithm suitability depends on the placement scope, objective structure, topology, workload dynamics, and evaluation protocol. Current evidence is fragmented by non-standard benchmarks, inconsistent baselines, limited statistical reporting, and incomplete reproducibility. A controlled comparison of GA/NSGA-II, PSO, GWO, and MPA under a shared physical fog-node placement model is a logical next step. Such a study should use common instances, equal computational budgets, repeated random seeds, statistical tests, and open artefacts. Because the present article is a literature synthesis, all comparative judgements are qualitative and require validation through controlled experiments. This conclusion is consistent with the heterogeneous objectives and evaluation settings reported across fog-node, service-placement, and vehicular-fog studies.

References

[1] S. A. Al-Asadi and S. O. Al-Mamory, "A survey on edge and fog nodes' placement methods, techniques, parameters, and constraints," IET Networks, vol. 12, no. 5, pp. 197–228, 2023, doi: 10.1049/ntw2.12087.

[2] H. K. Apat, V. Goswami, B. Sahoo, R. K. Barik, and M. J. Saikia, "Fog Service Placement Optimization: A Survey of State-of-the-Art Strategies and Techniques," Computers, vol. 14, no. 3, Art. no. 99, 2025, doi: 10.3390/computers14030099.

[3] S. Pallewatta, V. Kostakos, and R. Buyya, "Placement of Microservices-based IoT Applications in Fog Computing: A Taxonomy and Future Directions," arXiv:2207.05399, 2022.

[4] A. Naouri, N. A. Nouri, A. Khelloufi, A. Ben Sada, H. Ning, and S. Dhelim, "Efficient fog node placement using nature-inspired metaheuristic for IoT applications," Cluster Computing, vol. 27, pp. 8225–8241, 2024, doi: 10.1007/s10586-024-04409-3.

[5] A. Naouri et al., "Maximizing UAV fog deployment efficiency for critical rescue operations: A multi-objective optimization approach," Future Generation Computer Systems, vol. 159, pp. 255–271, 2024, doi: 10.1016/j.future.2024.05.007.

[6] J. Niu, G. Liu, L. Yu, and J. Wang, "Service Placement Optimization Based on Evolutionary Algorithm in Fog Computing," EAI Endorsed Transactions on Internet of Things, vol. 7, no. 26, Art. no. e5, 2022, doi: 10.4108/eai.22-2-2022.173492.

[7] M. Ghobaei-Arani and A. Shahidinejad, "A cost-efficient IoT service placement approach using whale optimization algorithm in fog computing environment," Expert Systems with Applications, vol. 200, Art. no. 117012, 2022, doi: 10.1016/j.eswa.2022.117012.

[8] M. Salimian, M. Ghobaei-Arani, and A. Shahidinejad, "Toward an autonomic approach for Internet of Things service placement using gray wolf optimization in the fog computing environment," Software: Practice and Experience, vol. 51, no. 8, pp. 1745–1772, 2021.

[9] F. Han, H. Fu, B. Wang, Y. Xu, and B. Lv, "GP4ESP: a hybrid genetic algorithm and particle swarm optimization algorithm for edge server placement," PeerJ Computer Science, vol. 10, Art. no. e2439, 2024, doi: 10.7717/peerj-cs.2439.

[10] M. M. Hussain, A. T. Azar, R. Ahmed, S. U. Amin, B. Qureshi, V. D. Reddy, I. Alam, and Z. I. Khan, "SONG: A Multi-Objective Evolutionary Algorithm for Delay and Energy Aware Facility Location in Vehicular Fog Networks," Sensors, vol. 23, no. 2, Art. no. 667, 2023, doi: 10.3390/s23020667.

[11] H. Ye, B. Cao, J. Liu, P. Li, B. Tang, et al., "An edge server deployment method based on optimal benefit and genetic algorithm," Journal of Cloud Computing, vol. 12, Art. no. 148, 2023, doi: 10.1186/s13677-023-00524-5.

[12] T. Lähderanta, T. Leppänen, L. Ruha, L. Lovén, E. Harjula, M. Ylianttila, J. Riekki, and M. J. Sillanpää, "Edge computing server placement with capacitated location allocation," Journal of Parallel and Distributed Computing, vol. 153, pp. 130–149, 2021, doi: 10.1016/j.jpdc.2021.03.007.

[13] Z. Wang, M. Goudarzi, and R. Buyya, "ReinFog: A Deep Reinforcement Learning empowered framework for resource management in edge and cloud computing environments," Journal of Network and Computer Applications, vol. 242, Art. no. 104250, 2025, doi: 10.1016/j.jnca.2025.104250.

[14] H. Sulimani, A. M. Sajjad, W. Y. Alghamdi, O. Kaiwartya, T. Jan, S. Simoff, and M. Prasad, "Reinforcement optimization for decentralized service placement policy in IoT-centric fog

environment," Transactions on Emerging Telecommunications Technologies, vol. 34, no. 11, Art. no. e4650, 2022, doi: 10.1002/ett.4650.

[15] D. Sahu, Nidhi, S. Prakash, P. Sinha, T. Yang, R. S. Rathore, and L. Wang, "Beyond boundaries: a hybrid cellular Potts and particle swarm optimization model for energy and latency optimization in edge computing," Scientific Reports, vol. 15, Art. no. 6266, 2025, doi: 10.1038/s41598-025-90348-x.

[16] C. Guerrero, I. Lera, and C. Juiz, "Genetic-based optimization in fog computing: Current trends and research opportunities," Swarm and Evolutionary Computation, vol. 72, Art. no. 101094, 2022, doi: 10.1016/j.swevo.2022.101094.

[17] A. Asghari and M. K. Sohrabi, "Server placement in mobile cloud computing: A comprehensive survey for edge computing, fog computing and cloudlet," Computer Science Review, vol. 51, Art. no. 100616, 2024, doi: 10.1016/j.cosrev.2023.100616.

[18] S. A. Temesheva, D. A. Turlykozhayeva, S. N. Akhtanov, N. M. Ussipov, A. A. Zhunuskanov, W. Sun, Q. Xu, and M. Tao, "A Comprehensive Review of Metaheuristic Algorithms for Node Placement in UAV Communication Networks," Sensors, vol. 26, no. 3, Art. no. 869, 2026, doi: 10.3390/s26030869.

[19] M. J. Page et al., "The PRISMA 2020 statement: an updated guideline for reporting systematic reviews," BMJ, vol. 372, Art. no. n71, 2021, doi: 10.1136/bmj.n71.