Metaheuristic Optimization Review

ISSN:3066-280X Vol: 07, No. 02, 2026 PP. 28–50

Article DOI: https://doi.org/10.54216/MOR.070203

REVIEW ARTICLE

Intelligent Energy Management: A Comprehensive Review of Artificial Intelligence, Machine Learning, and Metaheuristic Optimization Approaches

Faustino D. Reyes[1,*]

[1] ICT Bahrain Polytechnic,PO Box 33349, Isa Town, Bahrain

Email: faustino.reyes@polytechnic.bh

Received: May 18, 2026 Revised: July 01, 2026 Accepted: August 29, 2026 ⋆ Corresponding author

ABSTRACT

Intelligent energy management is increasingly important for improving energy efficiency, reliability, economic performance, and sustainability across buildings, smart homes, industrial facilities, microgrids, and interconnected energy systems. This review examines recent advances integrating machine learning, deep learning, reinforcement learning, Internet of Things technologies, predictive analytics, feature selection, and metaheuristic optimization. These methods support energy-consumption forecasting, demand response, renewable-energy coordination, load scheduling, anomaly detection, building optimization, smart-home control, industrial energy management, and microgrid operation. The reviewed studies indicate that intelligent approaches can improve forecasting accuracy, reduce energy costs, optimize resource allocation, and enhance real-time decision-making under uncertain conditions. Metaheuristic algorithms further improve performance through parameter optimization, feature selection, and multi-objective scheduling. However, practical deployment remains challenged by data quality, computational complexity, cybersecurity, privacy, interoperability, generalizability, transparency, infrastructure cost, and communication reliability. Future research should prioritize explainable, scalable, secure, and adaptive intelligent energy-management frameworks for efficient, resilient, flexible, and low-carbon energy system operation.

Keywords: Agentic artificial intelligence ▪ Autonomous agents ▪ Large language models ▪ Multi-agent systems ▪ Human-centered governance

1. INTRODUCTION

The continuing transformation of modern energy systems has substantially increased the complexity associated with the production, distribution, consumption, monitoring, and coordination of electrical energy. Conventional energy infrastructures were predominantly designed around centralized generation, comparatively predictable consumption profiles, and operational strategies governed by predefined rules and deterministic control procedures. Contemporary energy environments differ considerably from this traditional configuration. Increasing electricity demand, the widespread deployment of distributed resources, the growing penetration of renewable generation, the electrification of transportation and industrial activities, and the continuous expansion of digitally connected devices have created energy systems characterized by greater variability, decentralization, and operational uncertainty. At the same time, modern energy infrastructures generate extensive streams of heterogeneous information through smart meters, sensing devices, communication networks, automated controllers, and digitally monitored equipment. The resulting environment creates an important requirement for management mechanisms that can transform continuously generated data into useful operational knowledge rather than merely record historical consumption. Intelligent energy management has consequently emerged as a significant technological direction for coordinating energy resources through the integration of computational learning, data analysis, optimization, sensing, and automated decision-support mechanisms.

The development of intelligent energy management is closely associated with the broader evolution of connected computational environments in which physical resources are monitored and coordinated through distributed sensing and communication technologies. The increasing use of connected devices enables energy-management platforms to obtain detailed information regarding demand conditions, equipment operation, environmental characteristics, resource availability, and system status. However, greater connectivity does not automatically lead to more effective management because the resulting infrastructure must allocate computational and communication resources efficiently while handling devices with heterogeneous capabilities and service requirements. Research concerned with resource allocation in connected-device environments demonstrates that large-scale interconnected systems can present complex optimization problems and that nature-inspired metaheuristic approaches provide flexible mechanisms for addressing resource-allocation requirements under heterogeneous operating conditions [1]. From the perspective of intelligent energy management, this computational principle is important because future energy infrastructures increasingly depend on interconnected sensing, communication, and control components whose efficient coordination directly influences the quality and responsiveness of management decisions.

A central requirement of intelligent energy management is the ability to derive useful information from complex operational data. Energy consumption, renewable generation, equipment condition, environmental variables, temporal behavior, and user-related factors can exhibit nonlinear relationships that are difficult to describe through fixed analytical expressions. Machine-learning techniques provide a data-driven mechanism for identifying such relationships and developing predictive representations directly from observed information. Nevertheless, prediction in realistic environments becomes more demanding when data patterns vary, measurements contain noise, or operational conditions change over time. Hybrid computational approaches that combine machine learning with metaheuristic optimization have demonstrated the ability to address complex detection and prediction problems involving dynamically varying data [2]. Although individual application domains differ, the underlying methodological principle is directly relevant to intelligent energy management: learning mechanisms can provide predictive knowledge, while optimization can refine model configuration or decision processes so that analytical outputs become more appropriate for changing operating environments.

The increasing dimensionality of digitally monitored systems creates an additional analytical problem because a larger number of measured variables does not necessarily provide a proportionate improvement in predictive performance. Intelligent energy systems may collect historical demand values, environmental measurements, temporal indicators, equipment states, distributed-generation information, storage conditions, occupancy-related variables, and several other contextual features. Some of these variables can provide highly informative descriptions of system behavior, whereas others may be redundant, weakly related to the prediction target, or affected by unnecessary noise. Including all available inputs can therefore increase computational requirements and complicate model training without guaranteeing improved analytical performance. Feature selection provides a mechanism for reducing this unnecessary dimensionality by identifying compact subsets of informative variables. Multi-objective metaheuristic feature-selection research has shown that feature selection can simultaneously involve competing requirements such as maximizing predictive relevance while minimizing the number of retained variables [3]. Such a formulation is particularly suitable for intelligent energy management, where computational efficiency and predictive quality may need to be balanced when analytical models are expected to operate within resource-constrained or near-real-time environments.

The effectiveness of feature-selection procedures, however, should not be interpreted solely according to the number of variables removed from a dataset. The usefulness of a selected feature subset depends strongly on the characteristics of the target problem, the analytical method applied after selection, and the operational objectives against which the resulting model is evaluated. Consequently, intelligent energy-management research requires evaluation strategies that consider the relationship between feature-selection behavior and the actual requirements of the intended application rather than assuming that one universally optimal subset or method exists. Methodological research examining feature-selection and clustering techniques according to project-specific requirements emphasizes the importance of evaluating analytical procedures in relation to the objectives and conditions of the problem for which they are intended [4]. This perspective is relevant to energy-management applications because different environments may prioritize different characteristics: a building forecasting model may emphasize predictive accuracy, an embedded controller may require low computational complexity, and a large-scale energy platform may require scalability across heterogeneous data sources.

The search for informative feature subsets itself becomes increasingly difficult as the dimensionality of the available dataset grows. For a dataset containing numerous candidate variables, the number of possible feature combinations can increase rapidly, making exhaustive evaluation computationally impractical. Metaheuristic optimization offers an alternative approach by searching the feature space iteratively rather than evaluating every possible combination. The investigation of a hybrid Butterfly–Grey Wolf Optimizer for high-dimensional feature selection demonstrates the broader capability of population-based metaheuristic strategies to address complex feature-selection search spaces [5]. Within intelligent energy management, this capability creates opportunities for constructing predictive models that operate on more compact information representations while preserving the variables most relevant to the target management objective. This role is conceptually distinct from operational energy scheduling: in feature selection, the optimizer deter-

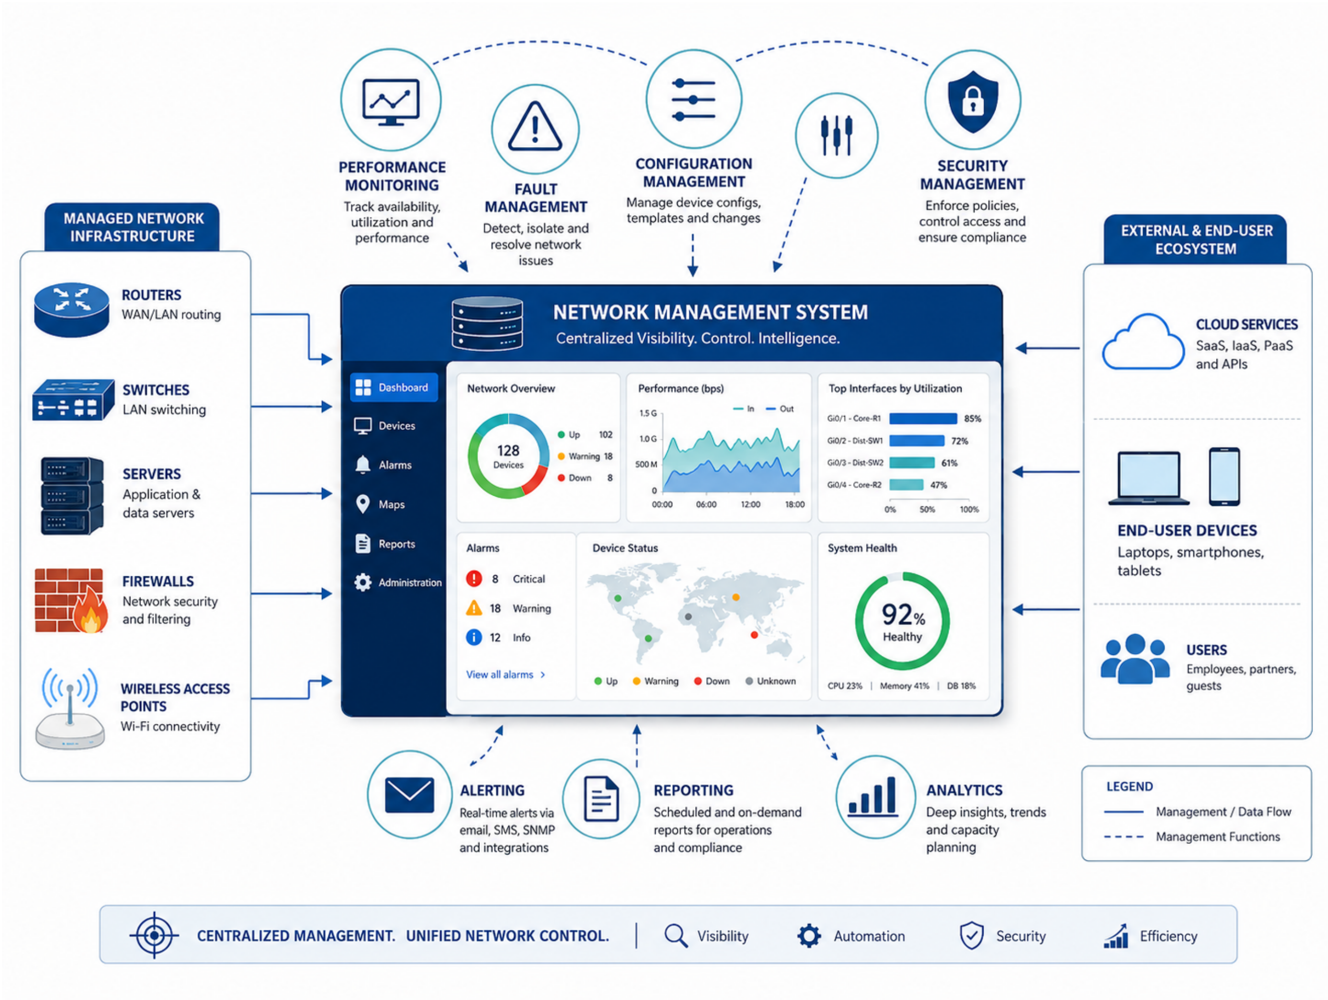

Figure 1. Conceptual architecture of a Network Management System integrating network infrastructure.

mines which information should enter the learning process, whereas in operational optimization the decision variables can represent actual energy resources or control actions.

Predictive modeling represents another fundamental component of intelligent energy management because many operational decisions depend on information concerning conditions that have not yet occurred. Estimates of future energy demand, renewable-resource availability, equipment behavior, consumption variation, or other relevant variables can allow management systems to prepare actions in advance rather than react only after an undesirable condition becomes observable. The performance of predictive models, however, can depend significantly on their configuration and on the interaction between learning architecture and dataset characteristics. The integration of a machine-learning model with metaheuristic optimization for environmental time-series prediction provides evidence that optimization can be embedded directly within predictive frameworks to improve model configuration for complex nonlinear forecasting tasks [6]. In intelligent energy-management architectures, this principle provides an important connection between prediction and optimization because an optimizer can contribute to improving the analytical model before that model is used to support operational decisions.

Data preparation can further influence predictive quality when multiple input variables exhibit different levels of association with the target quantity. Correlation-based analysis can provide an initial mechanism for characterizing relationships among candidate variables, while optimization can subsequently determine model or feature configurations according to the intended predictive objective. A prosumer electricity-forecasting study combines feature engineering with backward wrapper-based selection tailored to LightGBM [7]. The methodological significance of such integration for intelligent energy management lies in its structured combination of input analysis, feature selection, and learner-specific prediction. Energy datasets frequently contain variables that differ considerably in relevance, and preprocessing strategies capable of identifying meaningful relationships before optimization may reduce unnecessary search effort and provide more interpretable analytical inputs.

Practical intelligent systems must also remain effective when operating conditions differ from those represented during initial model development. Energy infrastructures are inherently heterogeneous: buildings vary in physical characteristics, industrial facilities operate under different production schedules, distributed resources exhibit different generation profiles, and energy consumption changes according to environmental and behavioral conditions. Consequently, models evaluated under a narrow operating regime may not necessarily maintain equivalent performance when deployed under different conditions. Research combining metaheuristic feature selection with machine learning for scalable diagnostic modeling has explicitly considered performance across diverse operating conditions [8]. This methodological concern is directly relevant to intelligent energy management because the value of a predictive or diagnostic model depends not only on its performance within a single experimental configuration but also on its ability to remain useful when the physical system experiences realistic operational variation.

The increasing intelligence and connectivity of energy infrastructures also make cybersecurity an integral consideration in the development of management architectures. Smart energy systems depend on data exchange among sensors, communication nodes, control platforms, and distributed resources. Decisions generated by an intelligent management framework can therefore be influenced by the reliability and integrity of the information transmitted through the underlying communication environment. A compromised monitoring or communication layer may introduce misleading measurements, interfere with device coordination, or reduce confidence in automated control decisions. Research applying optimization-based feature selection and machine-learning intrusion detection specifically within wireless mesh networks associated with smart grids demonstrates the importance of incorporating intelligent security mechanisms into digitally connected energy infrastructure [9]. Energy-management intelligence should consequently be considered together with the integrity of the digital environment that supplies the information upon which operational decisions are based.

This security requirement becomes particularly significant as intelligent energy infrastructures increasingly resemble wider industrial connected-device ecosystems in which physical processes, communication technologies, and artificial intelligence components operate collaboratively. Industrial environments demonstrate how interconnected devices and automated analytical systems create expanded opportunities for real-time monitoring while simultaneously increasing the number of potential points through which malicious or abnormal activity may affect system operation. Hybrid artificial-intelligence frameworks developed for intrusion detection in industrial connected-device environments illustrate the increasing role of intelligent detection mechanisms in protecting data-driven operational infrastructures [10]. For intelligent energy-management systems deployed in industrial or highly connected environments, secure information exchange is therefore not merely an external information-technology requirement; it forms part of the reliability conditions necessary for trustworthy data-driven energy decision-making.

The combined development of predictive learning, feature selection, metaheuristic optimization, connected sensing, and secure communication demonstrates that intelligent energy management is not defined by a single computational model. Instead, it represents a layered technological framework in which different components perform distinct analytical and operational functions. Data-acquisition technologies provide information regarding the state of the physical environment; preprocessing mechanisms transform raw observations into suitable representations; feature-selection procedures identify informative variables; predictive models estimate relevant system behavior; optimization mechanisms search for appropriate configurations or decisions; and communication infrastructures enable information to move among distributed physical and computational components. The overall effectiveness of intelligent energy management therefore depends on how these functions are coordinated within a coherent management architecture.

This distinction is important because improvements in one computational stage cannot automatically compensate for limitations in another. Increasing the complexity of a predictive model is unlikely to provide meaningful operational benefit when input measurements are unreliable or dominated by irrelevant variables. Similarly, sophisticated optimization cannot guarantee effective energy decisions when the predictive information supplied to the optimization process does not adequately characterize system conditions. A highly connected sensing infrastructure can generate substantial volumes of data, but those data have limited management value unless suitable analytical mechanisms transform them into actionable information. Intelligent energy management should therefore be considered as an integrated decision process rather than as a collection of independently optimized computational techniques.

The need for such integration becomes more apparent as the range of energy-management applications expands. Smart homes require the coordination of flexible residential resources; commercial and institutional buildings require efficient management of diverse electrical and thermal loads; industrial environments must relate energy consumption to changing process conditions; distributed energy systems require coordination between local generation, demand, and storage; and larger interconnected infrastructures must accommodate information exchange among geographically distributed components. Although these environments differ substantially in their physical characteristics, they share a fundamental requirement: large amounts of operational information must be transformed into decisions capable of improving the utilization of available energy resources.

Artificial intelligence therefore contributes to intelligent energy management through several complementary functions rather than through a single universal model. Predictive methods can estimate future behavior, classification mechanisms can distinguish operational states, feature-selection strategies can reduce information dimensionality, and optimization procedures can search for model configurations or operational decisions. The value of integrating these functions lies in the possibility of constructing management systems that can progress beyond descriptive observation toward anticipatory and adaptive operation. Such systems can potentially use historical and real-time information to understand current conditions, estimate likely future states, and determine appropriate responses within the constraints of the managed environment.

Metaheuristic optimization is particularly relevant within this integrated framework because optimization can occur at several different levels. It can be used before model training to identify informative features, during model development to configure analytical parameters, or during system operation to determine resource-allocation and scheduling decisions. These roles should remain conceptually distinct because their decision variables and objectives differ substantially. Feature-selection optimization seeks an informative representation of the available data, model optimization seeks an effective computational configuration, and operational optimization seeks an appropriate physical or managerial action. Recognizing these differences is necessary for evaluating the contribution of metaheuristic algorithms to intelligent energy management without treating all optimization applications as methodologically equivalent.

The motivation for the present review therefore arises from the need to examine intelligent energy management as an integrated computational field in which machine learning, predictive analysis, feature selection, metaheuristic optimization, connected sensing, and secure digital infrastructure perform complementary roles. The review does not treat these technologies merely as an inventory of algorithms. Instead, it considers how they contribute to different stages of the energy-management process and how methodological decisions made at one stage can influence the performance and practicality of subsequent stages. This perspective is particularly important for understanding why improvements in predictive accuracy alone cannot represent the complete effectiveness of an intelligent energy-management framework.

The scope of the review encompasses intelligent computational approaches relevant to energy monitoring, prediction, resource coordination, information selection, optimization, and secure operation. Particular attention is directed toward the transition from conventional static management strategies toward systems capable of extracting knowledge from data and employing that knowledge within computational decision processes. The review further considers how feature selection and model optimization can improve the analytical foundations of energy management while connected sensing and secure communication provide the infrastructure necessary for practical deployment.

The principal contribution of this review is to establish a coherent conceptual foundation for evaluating intelligent energy-management technologies according to the specific function performed by each computational component. By distinguishing data acquisition, feature representation, predictive learning, model optimization, resource decision-making, and digital-system protection, the review provides a structured basis for examining the literature without conflating methodologically different tasks. This organization also enables subsequent analysis to identify relationships among computational approaches while avoiding unnecessary repetition of identical concepts across different application categories.

The remainder of the paper is organized into three principal sections. The Literature Review critically synthesizes the selected studies and examines the methodological development of intelligent computational approaches relevant to energy-management systems. The Discussion interprets the collective methodological implications emerging from the reviewed literature while addressing issues that are not repeated from the study-by-study synthesis. Finally, the Conclusion consolidates the principal observations of the review and presents the overall significance of integrating intelligent prediction, optimization, sensing, and secure computational infrastructures within modern energy-management environments.

2. LITERATURE REVIEW

The literature relevant to intelligent energy management increasingly reflects a methodological movement away from isolated predictive models toward computational architectures in which data reduction, learning, optimization, and decision formation are deliberately coupled. This development is significant because the effectiveness of an intelligent system is determined not only by the individual algorithm used at a particular stage but also by the manner in which computational components interact throughout the analytical pipeline. Recent studies in closely related data-intensive and interconnected environments demonstrate several distinct patterns of integration, including optimizer-assisted data reduction, optimization-guided classifier construction, dynamically coordinated metaheuristic search, hybrid population-based optimization, and ensemble-driven deep-learning architectures. Although many of these studies originate from cybersecurity, distributed computing, industrial connected systems, and cloud environments rather than direct energy-management applications, they provide useful methodological evidence regarding how intelligent computational components can be organized within complex operational systems. Their relevance to intelligent energy management therefore lies primarily in the architectural and optimization principles they establish rather than in direct application-specific conclusions about energy consumption or resource scheduling.

This architectural arrangement introduces an important distinction in the interpretation of optimization-assisted intelligent systems. The optimizer does not necessarily determine the final operational decision directly. Instead, it can alter the information space from which the subsequent learning mechanism constructs its decision boundary. This position gives the optimization stage an indirect but potentially substantial influence over the final analytical outcome. Such a configuration requires careful coordination between the criterion used to evaluate candidate variable subsets and the behavior of the downstream analytical model. A subset that appears compact from a dimensional perspective may not be useful if it removes variables required to discriminate relevant operating states, whereas a comparatively larger subset may preserve unnecessary information and increase computational requirements. Consequently, the quality of an optimization-assisted selection mechanism should be interpreted according to its interaction with the subsequent learner rather than according to the reduction ratio alone. This observation provides a useful basis for evaluating intelligent architectures in which multiple computational components influence one another sequentially.

A related branch of the literature explicitly combines machine-learning procedures with metaheuristic optimization for simultaneous feature-selection and classification objectives. Work on botnet attack detection, for example, has examined machine learning together with metaheuristic optimization algorithms as a coordinated framework in which the search procedure determines informative feature combinations and the learning mechanism evaluates the discriminatory usefulness of those combinations [11]. This formulation differs from conventional workflows in which feature reduction and classification are selected independently. The optimizer and learner instead become mutually dependent: the search process requires an evaluation mechanism capable of indicating

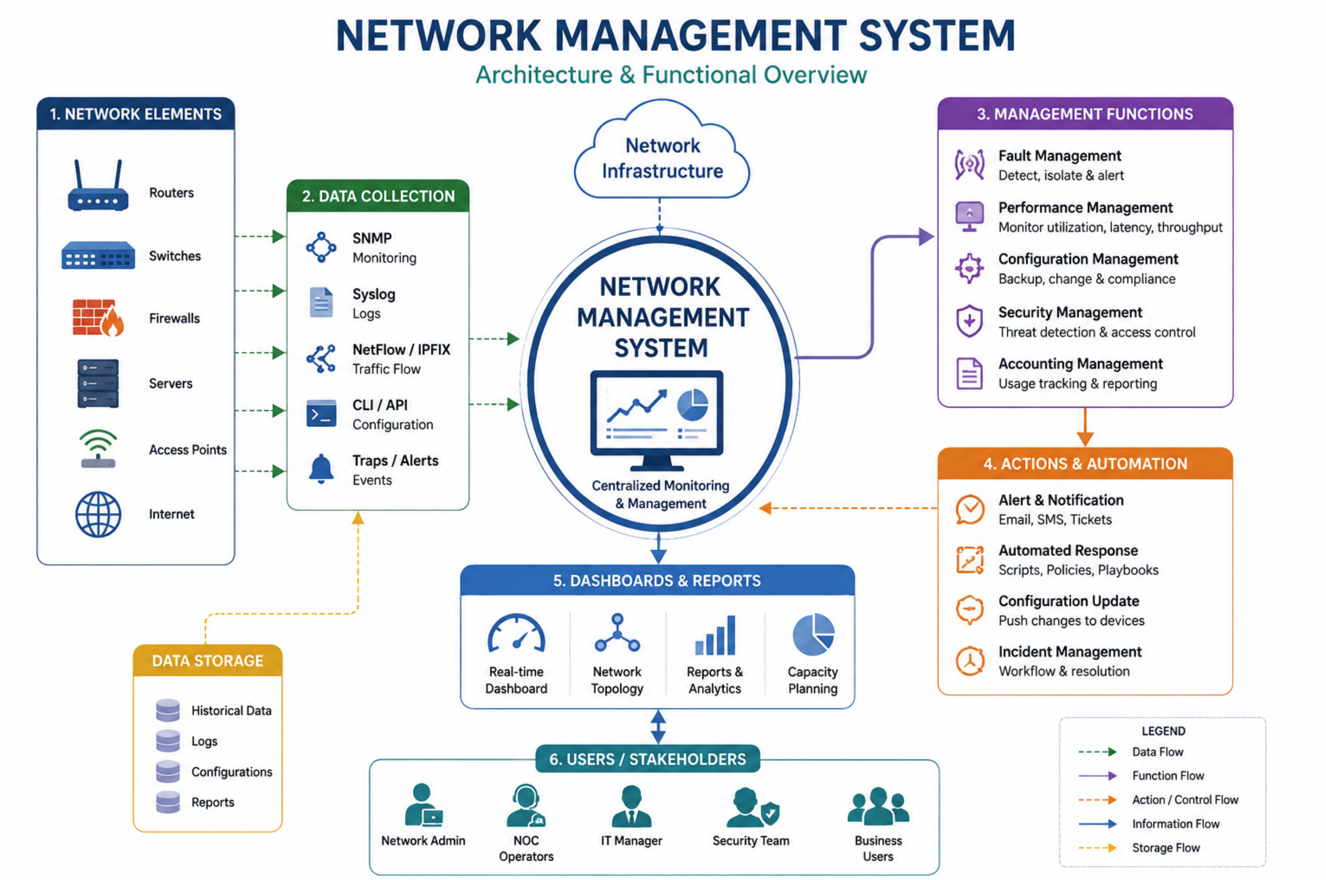

Figure 2. Conceptual architecture of a Network Management System integrating network elements, data collection, centralized management functions, automated actions, dashboards, data storage, and stakeholder interaction.

whether a candidate subset is useful, while the learning model depends on the search process to define the representation on which its subsequent classification behavior is based. The methodology therefore creates a feedback relationship between data representation and model performance.

This interaction has important implications for the design of intelligent computational frameworks. When machine learning is embedded inside a metaheuristic search process, the objective function becomes one of the most consequential methodological decisions. An optimization procedure can only search for solutions according to the information supplied by its fitness formulation. If the objective considers only predictive performance, the resulting search may favor comparatively large candidate subsets whenever they provide even a marginal improvement in classification quality. Conversely, an objective that emphasizes dimensional reduction excessively may produce highly compact representations that omit relevant information. Multi-criteria formulations can address this tension, but they introduce additional decisions concerning weighting, normalization, dominance relations, or trade-off selection. Accordingly, the effectiveness of a hybrid machine-learning–metaheuristic framework depends as much on the formulation of the search objective as on the identity of the optimizer itself. This methodological issue is particularly important when different optimization algorithms are compared because an apparent difference in algorithmic performance may actually result from differences in fitness construction, stopping criteria, or learner configuration.

The literature has also progressed toward hybrid search mechanisms that modify the optimization process itself rather than simply combining one optimizer with one predictive model. A feature-selection framework for network intrusion detection employed a hybrid metaheuristic dynamic optimization strategy, illustrating a design in which feature-space exploration is performed through a dynamically coordinated search mechanism rather than through a single fixed search procedure [12]. The significance of this approach lies in the treatment of the search process as a configurable computational system. Instead of assuming that one search behavior is equally effective throughout the entire optimization trajectory, dynamic hybridization provides a basis for changing the relative contribution of different search behaviors as candidate solutions evolve.

This dynamic-search perspective raises a separate methodological issue concerning how exploration and refinement should be distributed across the optimization process. Search mechanisms that produce strong diversification during the early iterations may be beneficial when the candidate space contains many distant promising regions, whereas stronger local refinement may become more useful once high-quality regions have been identified. A dynamically coordinated metaheuristic framework can therefore be interpreted as an attempt to control the search trajectory rather than merely accelerate convergence. This is distinct from simple parameter adjustment because the underlying search behavior itself may be altered according to the state of the optimization process. For intelligent computational applications, such mechanisms create an additional design dimension: optimization performance depends not only on the selected algorithm but also on when and how different search operators are activated.

This observation is particularly relevant to high-dimensional analytical problems because the topology of a search space can change substantially as variables are added or removed. In a binary feature-selection problem, every additional candidate feature doubles the number of possible subsets, causing the theoretical solution space to expand exponentially. An optimizer operating in such an environment must therefore balance sufficient diversification to avoid becoming confined to a limited region with adequate refinement to distinguish among similar high-quality subsets. Dynamic hybridization provides one mechanism for managing this balance. However, it also increases implementation complexity because the framework must specify how search components interact, how transitions are controlled, and whether additional mechanisms improve the final solution sufficiently to justify their computational overhead.

A further line of research has investigated explicit hybridization between established population-based optimization algorithms. An IoT intrusion-detection framework combines the Whale Optimization Algorithm (WOA) with the Grey Wolf Optimizer (GWO) for feature selection and parameter optimization [13]. Unlike a comparison of independent optimizers, this framework combines different search behaviors within one procedure. WOA contributes a spiral search mechanism, while GWO uses population guidance based on a social hierarchy. For intelligent energy-management research, the transferable principle is the coordination of complementary exploration and refinement mechanisms; the reported intrusion-detection performance should not be interpreted as evidence of energy-management performance.

Hybrid optimizer design creates several analytical questions that should be distinguished from conventional algorithm comparison. The first concerns the level at which the constituent methods interact. A sequential hybrid may allow one optimizer to produce an initial solution population that is subsequently refined by another, whereas an embedded hybrid may exchange candidate information continuously during the same optimization process. A parallel architecture can maintain multiple search processes simultaneously and subsequently combine their results. A hierarchical architecture may assign different optimization responsibilities to different layers. These configurations are not methodologically equivalent, even when they contain the same constituent optimizers. Consequently, describing a framework only according to the names of the included algorithms can conceal the mechanism through which the hybrid system actually searches the solution space.

A second issue concerns the computational cost introduced through hybridization. Combining multiple optimizers can increase the number of objective-function evaluations, population-management operations, model-training cycles, or information-exchange procedures required before termination. Therefore, superiority in a final predictive metric cannot automatically establish that a hybrid framework is operationally preferable. The computational budget required to obtain that performance must also be considered. This concern becomes particularly important for intelligent systems intended to support frequent model updating or near-real-time analytical decisions, because optimization latency can become part of the practical suitability of the framework. Accordingly, comparative evaluation of hybrid optimizers should distinguish improvements produced by genuinely more effective search behavior from improvements obtained primarily through a substantially larger computational budget.

A more structurally elaborate design appears in research combining ensemble-based feature selection with optimization-driven deep learning for attack detection in cloud-computing environments [14]. This type of architecture introduces multiple analytical stages before the final decision is produced. Ensemble-based selection attempts to determine informative inputs using the combined behavior of several selection mechanisms, while optimization is subsequently associated with the configuration of the deep-learning component. The architecture therefore distributes computational intelligence across different levels of the pipeline rather than expecting one algorithm to perform all analytical responsibilities. Such a design represents a progression from single-stage hybridization toward modular intelligent systems in which separate mechanisms are assigned to information screening, model construction, and final analytical inference.

The ensemble component introduces a methodological principle that differs from the population-search concepts discussed previously. Whereas a metaheuristic searches candidate solutions through iterative population evolution, an ensemble-based selection mechanism can exploit agreement, complementarity, or aggregation among multiple selectors. The purpose is not necessarily to locate one candidate solution through a single search trajectory but to construct a more stable representation from the outputs of several analytical mechanisms. This distinction becomes important when evaluating multi-stage frameworks because an ensemble and a metaheuristic should not be interpreted as interchangeable computational tools. They solve different methodological problems. The ensemble addresses diversity among analytical judgments, while the optimizer addresses the search for an appropriate configuration according to an objective function.

The coexistence of these mechanisms within one framework also introduces the issue of error propagation across computational stages. In a multi-stage intelligent architecture, an early analytical decision influences every subsequent component. If the input-selection stage removes useful information, later optimization cannot recreate variables that are no longer available. Similarly, a high-quality feature representation does not guarantee strong final performance if the optimized learning architecture is poorly matched to the structure of the target problem. Evaluation of multi-stage systems should therefore examine the contribution of each stage separately rather than reporting only the final output of the complete pipeline. Ablation analysis is particularly useful in this context because it can determine whether each additional module produces a measurable contribution when the remaining architecture is held constant. Without such analysis, increasingly complex frameworks can appear effective while leaving uncertainty regarding which component is actually responsible for the observed improvement.

Taken collectively, the reviewed methodological evidence reveals a progression in the structural complexity of optimization-assisted intelligent systems. The earliest architectural form considered here places a metaheuristic directly before a classifier as an attribute-selection mechanism. A second pattern couples the learning algorithm with the optimization procedure so that candidate representations are evaluated according to downstream predictive behavior. Dynamic hybridization then modifies the internal search trajectory itself, while explicit multi-optimizer frameworks combine distinct population-based search mechanisms. Finally, ensemble-driven architectures distribute analytical responsibilities across several computational layers. These configurations form a useful taxonomy because they demonstrate that the phrase "metaheuristic-assisted intelligence" can describe substantially different system architectures. Meaningful comparison therefore requires identification of the exact role played by optimization rather than merely recording whether an optimizer was included.

This architectural distinction is particularly useful when transferring methodological lessons into intelligent energy-management research. The key design question is not simply which optimizer achieves the highest performance on a particular benchmark. A more informative question concerns where optimization is positioned within the complete management architecture and which computational object it modifies. Optimization may influence the representation of measured information, the internal configuration of a learning model, the structure of a hybrid analytical pipeline, or a higher-level decision mechanism. These placements generate different objective functions, different decision spaces, and different computational costs. Consequently, comparisons across studies should first determine whether the competing algorithms actually solve the same optimization problem. Comparing a method used for input selection with another used for model configuration, for example, would provide little methodological insight unless the different optimization roles were explicitly recognized.

Another important theme emerging from the literature concerns evaluation granularity. Intelligent computational frameworks often report one final performance indicator even though several stages contribute to the result. A more rigorous evaluation strategy should assign metrics to the stage they are intended to measure. A feature-selection component can be assessed according to dimensional reduction, subset stability, and its influence on downstream performance. An optimizer can be evaluated according to convergence behavior, computational budget, robustness across independent runs, and sensitivity to initialization. A predictive learner can be evaluated using task-specific predictive metrics. A complete operational framework may additionally require assessment of latency, scalability, memory requirements, or decision feasibility. Maintaining this separation prevents one strong final metric from concealing weaknesses elsewhere in the architecture.

The stochastic character of most metaheuristic approaches creates an additional requirement for experimental design. A result from a single optimization run provides limited information regarding algorithmic reliability because different random initializations can produce different candidate trajectories and final solutions. Comparative studies should therefore distinguish best-case performance from typical performance and should report the distribution of outcomes across repeated independent runs. Measures such as the mean, median, standard deviation, interquartile range, or confidence intervals can provide information regarding the stability of the optimizer, while appropriate statistical tests can determine whether observed differences between competing methods are sufficiently consistent to support comparative claims. This is particularly important when performance differences are numerically small, since a marginal improvement in a single run may fall within ordinary stochastic variability.

Reproducibility should also be considered independently from predictive quality. Hybrid intelligent frameworks can contain numerous implementation decisions, including population size, iteration limits, termination conditions, initialization procedures, algorithmic parameters, preprocessing choices, model configurations, feature bounds, and objective-function definitions. Incomplete reporting of these settings makes it difficult to determine whether an observed improvement originates from the proposed methodological contribution or from differences in experimental configuration. Computational-budget reporting is equally important when several optimizers are compared because equal iteration counts do not necessarily correspond to equal computational effort if the algorithms perform different numbers of objective evaluations per iteration. A reproducible evaluation should therefore expose enough implementation detail to allow the complete search and learning process to be reconstructed.

The reviewed literature further suggests that increasing architectural complexity should not automatically be interpreted as methodological advancement. Adding more optimizers, selection stages, ensemble components, or learning layers can increase representational capacity, but every additional component also introduces parameters, computational overhead, and potential sources of instability. A simpler framework may therefore be preferable when it achieves comparable performance with lower computational demand and clearer analytical behavior. This principle is particularly important for practically deployed intelligent systems because real-world implementation often imposes resource constraints that are absent from offline experimental evaluation. Model complexity should consequently be justified by demonstrable functional contribution rather than by architectural novelty alone.

An additional consideration concerns the relationship between benchmark-specific success and cross-domain methodological transfer. The five studies examined in this section primarily evaluate cybersecurity, network, industrial connected-device, and cloud-computing problems. Their results should therefore not be interpreted as direct empirical evidence that identical performance would be achieved in intelligent energy-management applications. Data distributions, objective functions, temporal dependencies, physical constraints, and operational consequences differ substantially between these domains. Nevertheless, the studies provide methodological evidence concerning the organization of hybrid intelligent systems, including how optimization can interact with learning, how multiple search mechanisms can be combined, how ensemble selection can precede deep models, and how architectural complexity can be distributed across several computational stages. Their value to an intelligent energy-management review is thus methodological rather than application-specific.

This distinction is essential for maintaining rigorous interpretation of the literature. Transferability should be evaluated at the level of computational principle rather than at the level of reported numerical performance. A feature-selection mechanism demonstrated in network data may motivate investigation of an analogous search structure for another high-dimensional dataset, but its original accuracy cannot be transferred to a different task. Similarly, the successful combination of two optimizers within one classification framework demonstrates the feasibility of hybrid search but does not establish that the same hybridization will be optimal under a different objective landscape. Recognizing this boundary prevents cross-domain evidence from being overstated while still allowing useful computational principles to inform subsequent research.

The literature can therefore be understood according to an emerging hierarchy of methodological integration. At the lowest level, optimization is attached to one isolated analytical task. More integrated systems couple optimization directly with learning. More advanced configurations coordinate several search behaviors or multiple optimizers. Multi-stage architectures then distribute distinct responsibilities among selection, optimization, and deep-learning components. This hierarchy provides a useful analytical framework for reviewing intelligent systems because it enables methodological complexity to be described according to functional organization rather than according to the number of algorithms listed in a model name.

For intelligent energy-management research, such an organization can support more meaningful evaluation of future computational architectures. The central issue is whether additional computational layers provide clearly identifiable value within the complete decision process. Each component should possess a defined function, its contribution should be independently measurable, and its computational requirements should be reported. This perspective also encourages researchers to avoid constructing hybrid models in which algorithms are combined primarily to create novelty without establishing a clear functional reason for each component.

Overall, the literature reviewed in this section demonstrates that contemporary intelligent computational research is progressively shifting toward deliberately structured hybrid architectures. The principal methodological development is not simply the growing use of metaheuristic algorithms but the increasing diversity of roles assigned to optimization within larger analytical systems. Feature-space modification, learner-dependent search, dynamic operator coordination, multi-optimizer hybridization, and ensemble-supported deep architectures represent distinct design strategies with different computational implications. Recognizing these distinctions provides a more precise foundation for assessing intelligent energy-management frameworks and prevents fundamentally different optimization tasks from being grouped together under a single broad methodological category. The subsequent Discussion can therefore concentrate on the higher-level implications arising from this literature without reproducing the individual methodological patterns examined here. To provide a structured methodological view without repeating the analytical arguments developed in the preceding literature synthesis, Table 1 organizes representative computational architectures according to the specific methodological function implemented in each study. The included works originate from several data-intensive application domains and are therefore not presented as direct empirical evidence for energy-management performance. Instead, the table isolates transferable computational designs that can inform intelligent energy-management research, including alternative strategies for feature-space construction, model calibration, multi-objective learning, hybrid optimization, comparative optimizer assessment, explainability, environmental modeling, and decision-oriented prediction. This organization distinguishes the functional contribution of each methodology and avoids treating fundamentally different computational tasks as equivalent merely because they employ machine learning or metaheuristic optimization.

Table 1 demonstrates that the methodological landscape represented by the reviewed studies cannot be characterized through a single dominant computational architecture. Instead, the selected references illustrate a gradual movement from relatively focused feature-selection procedures toward increasingly integrated systems in which data representation, model configuration, optimization, prediction, and decision support are distributed across several computational stages. The significance of this progression lies in the changing role assigned to optimization. In some methodologies, optimization acts directly on the available input variables; in others, it modifies the internal configuration of the predictive model; more elaborate frameworks distribute optimization across several components or combine it with deep representation learning, ensemble analysis, or explainability. This methodological diversity is particularly relevant to intelligent energy-management research because it shows that optimization should not automatically be regarded as a single interchangeable processing step. Its practical contribution depends on the computational object being optimized, the criterion used to measure candidate quality, and the position of the optimizer within the overall analytical pipeline.

The hybrid feature-selection framework represented in the first methodological category illustrates an important architecture in which information reduction occurs before the principal predictive or classification stage. This organization establishes a clear separation between determining which information should be retained and determining how the retained information should subsequently be interpreted. Such separation can be valuable in data-rich intelligent systems because raw datasets may contain variables with substantially different predictive contributions. Rather than forcing the predictive model to process the complete input space, a preliminary selection mechanism attempts to construct a more concentrated representation. Methodologically, this approach can reduce the effective dimensionality of the learning problem while allowing the subsequent learner to operate on a deliberately selected subset. The important principle is therefore not simply the removal of variables, but the architectural decision to place an information-screening mechanism before the main learner. For intelligent energy-management applications involving measurements from multiple meters, sensors, operational devices, environmental sources, and temporal records, such an architecture provides a general mechanism for separating data-space refinement from subsequent predic-

Table 1. Representative computational methodologies relevant to intelligent energy-management research

Study focus	Data/application type	Methodological approach	Main methodological objective
Hybrid feature-selection framework [15]	Distributed computing network data	Hybrid machine-learning architecture incorporating feature selection before intrusion classification	Constructs a reduced and discriminative input representation before model training, illustrating how data-space refinement can be positioned as an integrated stage of an intelligent analytical pipeline.
Search-driven variable selection [16]	UCI benchmark classification datasets	Grey Wolf Optimizer with a self-repulsion strategy for feature selection	Uses population-based search to reduce redundant variables while maintaining classification performance, illustrating optimizer-controlled input selection.
Sensitivity-informed modeling [17]	Ransomware-related behavioral data	Feature selection combined with sensitivity analysis and an optimized hybrid predictive model	Evaluates variable influence alongside optimized model construction so that model development incorporates information regarding the relative contribution of candidate inputs rather than relying only on final predictive output.
Comparative multi-objective optimization [18]	Transit network design and frequency setting under station congestion	Objective-space decomposition evolutionary algorithm, assessed against alternative multi-objective methods	Balances passenger and operator costs and compares competing search procedures within a common transit-planning problem.
Binary hybrid search [19]	Kronodroid Android malware features	Multi-wrapper feature selection combining Harris Hawks, Grey Wolf, and Sine Cosine optimizers	Combines complementary search strategies to select binary feature subsets, with Random Forest used to evaluate candidate representations.
Cross-model feature-selection assessment [20]	Seismic observations of debris flows	Feature-reduction strategies evaluated with Random Forest, XGBoost, and LSTM models	Examines how feature importance, input redundancy, and learner choice affect debris-flow classification and early-warning performance.
Swarm-assisted neural modeling [21]	Clinical diagnostic datasets	Swarm-intelligence metaheuristic optimization integrated with neural-network models	Synthesizes the role of swarm-based search in neural-model development and demonstrates how population-level optimization can be used to support configuration of nonlinear predictive architectures.
Multi-objective network fine-tuning [22]	Clinical diabetes data	Multi-objective metaheuristic-assisted fine-tuning of a deep network	Treats deep-model configuration as a multi-objective search problem, allowing competing optimization requirements to be considered simultaneously instead of reducing model selection to a single performance criterion.
Deep-model input optimization [23]	Cotton leaf image data	Metaheuristic feature-selection algorithms coupled with deep learning	Reduces the information supplied to the deep-learning stage through search-based feature selection, providing a framework in which input optimization precedes high-capacity representation learning.

tion or control.

The Black Hole Algorithm-based methodology represents a more explicit search-oriented interpretation of feature selection. In this formulation, variable selection is treated as an optimization problem rather than as a deterministic ranking process. Every candidate subset can be regarded as a possible solution within a very large combinatorial search space, and the optimizer is responsible for navigating that space toward configurations associated with favorable analytical performance. The methodological implication is substantial because the search procedure must determine not merely whether an individual variable appears useful in isolation, but whether a combination of variables collectively provides an effective representation. Variables that appear weak individually may become important when interacting with other inputs, whereas several individually informative variables may contain overlapping information. Search-driven feature selection therefore provides a mechanism for considering combinations rather than isolated relevance scores. For intelligent energy-management datasets containing correlated meteorological, temporal, behavioral, generation, and consumption variables, the same methodological principle can support the identification of compact input configurations without requiring exhaustive examination of every possible subset.

The sensitivity-informed hybrid methodology introduces a different analytical dimension because it places explicit attention on the influence of variables in addition to their inclusion or exclusion. Sensitivity analysis can provide information regarding how changes in particular inputs affect model behavior or predicted outcomes, thereby extending the methodological objective beyond simple subset construction. When combined with feature selection and optimized predictive modeling, this creates a pipeline in which the data are not only reduced but also examined according to their relative influence on the analytical process. Such a design is especially valuable when the objective includes understanding why a model responds strongly to certain measurements. Within

Table 1. Representative computational methodologies relevant to intelligent energy-management research (continued)

Study focus	Data/application type	Methodological approach	Main methodological objective
GWO-guided deep framework [24]	IoT and medical-IoT network traffic	Grey Wolf feature selection followed by CNN feature extraction and LightGBM classification	Separates variable selection, learned feature extraction, and classification into distinct stages of an intrusion-detection framework.
Model-level metaheuristic calibration [25]	Parkinson's disease clinical variables	Metaheuristic optimization of machine-learning models	Uses optimization at the model-configuration level rather than exclusively at the input-selection level, demonstrating a distinct role for search algorithms in refining predictive-model settings.
Multi-component forecasting architecture [26]	Cryptocurrency time-series data	Dual-branch Transformer combining self-attention and auto-correlation mechanisms	Combines complementary temporal representations to model dependencies across short and long forecasting horizons.
Explainable hybrid prediction [27]	Chronic kidney disease clinical data	Binary breadth-first feature selection, Extra Trees classification, and SHAP explanations	Combines compact input selection with ensemble prediction and feature-level explanations of the resulting classifications.
Connected-device hybrid prediction [28]	IoT-based kidney and heart disease measurements	TriKernel support-vector prediction with Levy Gazelle hyperparameter optimization	Connects patient-data acquisition, optimization, and supervised prediction within an IoT-oriented healthcare framework.
Metaheuristic-enhanced environmental modeling [29]	Treated-water quality measurements	Extreme learning machines optimized with PSO, GA, biogeography-based optimization, and a BBO–PSO hybrid	Compares optimizer-assisted regression models for continuously varying water-quality parameters.
Hybrid credit-risk feature optimization [30]	Peer-to-peer lending records	Hybrid modified Grey Wolf Optimizer with an optimized decision-tree model	Coordinates feature-space search with a decision-oriented predictive model, illustrating how optimization can be distributed between representation construction and the subsequent learner.
Physical-source localization [31]	Microseismic arrival-time measurements	Hybrid Genetic Algorithm–Particle Swarm Optimization for source location	Searches a physical coordinate space using adaptive swarm parameters and evolutionary operations rather than a learned feature representation.
Hybrid susceptibility mapping [32]	Flood-related geospatial variables	CatBoost with Whale and Zebra optimization algorithms for hyperparameter tuning	Uses swarm-based parameter search to improve spatial susceptibility models and support flood-risk mapping.

intelligent energy management, this distinction could become important when variables represent operationally meaningful quantities such as external conditions, historical demand, equipment states, renewable generation, or market signals. An optimization framework that identifies a strong predictive configuration provides one type of knowledge, whereas sensitivity analysis contributes another by revealing how individual inputs influence the resulting model response. The methodological value lies in maintaining these two analytical objectives as complementary rather than treating predictive performance as the only criterion of interest.

The comparative multi-objective methodology represented in the table examines transit network design and frequency setting under station congestion. Its objective-space decomposition evolutionary algorithm balances passenger and operator costs and is evaluated against alternative search methods. This example demonstrates why algorithmic comparison requires a shared decision problem rather than disconnected performance figures. Differences in objectives, computational budgets, constraints, or stopping criteria can confound comparisons. For intelligent energy management, the relevant lesson is to assess competing optimizers under equivalent operational requirements and to distinguish trade-offs among stakeholders rather than assume that one algorithm is universally preferable.

The multi-wrapper hybrid feature-selection methodology addresses the special characteristics of feature-selection problems in which the decision associated with each candidate variable is inherently discrete. A feature is generally either retained or excluded, meaning that the optimizer must operate within a binary representation even when its original search behavior was developed for continuous variables. This requires an explicit mechanism through which search movements are translated into binary selection decisions. Hybridizing multiple population-based strategies introduces an additional layer because the resulting method attempts to combine distinct exploration characteristics within the same binary search environment. The methodological interest lies in the construction of a search process suited to combinatorial decisions rather than the simple transfer of a continuous optimizer without adaptation. This consideration is directly relevant to intelligent energy-management modeling when the objective is to determine whether specific sensors, measurements, lagged observations, contextual variables, or derived features should participate in the predictive framework. The underlying decision is categorical, and the optimizer must therefore represent that structure appropriately.

The cross-model feature-selection methodology provides an important means of examining the interaction between data representation and learner choice. A feature subset cannot always be considered independently of the predictive algorithm that subsequently processes it. Different learning models possess different capacities for handling nonlinear relationships, redundant inputs, interactions, and complex decision boundaries. Consequently, a subset that performs particularly well with one learner may not necessarily provide equivalent behavior with another. Comparing multiple machine-learning algorithms together with feature-reduction strategies enables the evaluation to distinguish between improvements associated with better information representation and those associated with the intrinsic characteristics of the learner. This methodological separation is highly relevant to intelligent energy management because comparative studies frequently report a winning model without establishing whether its superiority results from model architecture, preprocessing, feature selection, or optimization. Cross-model assessment provides a more rigorous framework for identifying where the improvement actually originates.

Swarm-assisted neural modeling illustrates a broader integration between population-based optimization and nonlinear learning systems. Neural models contain numerous configurable elements whose selection can substantially influence their behavior. Swarm-based optimization offers a mechanism for searching this configuration space without relying solely on manual trial-and-error procedures. The important methodological contribution is the conversion of neural-model design into an explicit search problem. Candidate configurations can be represented as solutions, evaluated according to a defined criterion, and iteratively refined. This process can be used to explore settings that would otherwise require repeated manual experimentation. For intelligent energy-management modeling, where neural networks may be employed to represent nonlinear temporal or operational relationships, such optimization can provide a systematic mechanism for identifying model configurations suited to a particular dataset. The methodology therefore differs from feature selection because the optimization target is located inside the predictive architecture rather than exclusively within the input representation.

The multi-objective deep-network fine-tuning methodology extends this concept by rejecting the assumption that model configuration should necessarily be optimized according to one criterion only. Intelligent computational systems frequently involve competing objectives. A configuration associated with strong predictive performance may require greater computational complexity, while a simpler model may offer lower computational demand at the cost of a modest reduction in predictive quality. Multi-objective optimization provides a formal mechanism for representing such competing requirements rather than combining them implicitly or ignoring all but one. The resulting search process can identify a set of trade-off solutions rather than a single universally optimal configuration. This methodological perspective has particular relevance for intelligent energy management because practical deployments may require simultaneous consideration of accuracy, execution time, memory requirements, energy consumed by computation, robustness, or other operational criteria. Multi-objective fine-tuning therefore provides a conceptual basis for designing predictive intelligence according to practical system requirements rather than according to one isolated metric.

The methodology involving metaheuristic feature selection before deep learning illustrates a sequential organization in which the information supplied to a high-capacity model is optimized before representation learning begins. Deep-learning models are often capable of automatically extracting complex representations from data, but this does not imply that every possible input variable should necessarily be supplied without prior screening. A preliminary optimization stage can restrict the information space, after which the deep architecture concentrates on learning higher-level relationships within the selected representation. The methodological significance lies in the explicit division between variable-level search and hierarchical representation learning. Rather than expecting the deep model to perform both functions implicitly, the architecture assigns them to separate computational components. In intelligent energy-management applications, this type of separation may be particularly useful when structured measurements from many sources are supplied to a deep predictive architecture and only a subset provides meaningful information for the target task.

The GWO–CNN–LightGBM framework adds another perspective by linking feature-space optimization with a specialized deep-learning architecture. Its methodological structure demonstrates how a study can deliberately separate the optimization of the input representation from the nonlinear processing performed by the deep-learning component. The optimizer addresses the question of which variables should be retained, while the subsequent CNN feature extractor and LightGBM classifier addresses how the retained information should be transformed into the final prediction or classification output. This separation can facilitate more precise analysis of the computational responsibilities assigned to each stage. It also demonstrates why hybrid architectures should not be described merely by listing their constituent algorithms. The position and function of each component determine the actual methodology. For intelligent energy-management research, this provides a useful design principle: when several computational methods are combined, each should have a clearly defined role rather than being introduced simply to increase architectural complexity.

The model-level metaheuristic calibration methodology moves the optimization target away from the input variables entirely. Instead of determining which information enters the model, the search process is used to improve the configuration of the machine-learning models themselves. This distinction is important because feature selection and model optimization operate on different decision spaces. The former searches combinations of input variables, whereas the latter searches parameter or configuration alternatives governing how the learner processes those inputs. Treating these tasks separately allows researchers to identify whether predictive improvements are being generated by improved representation or improved model configuration. In intelligent energy-management systems, this distinction becomes particularly important when the same dataset is evaluated using several predictive models. A model should not necessarily be considered inferior simply because it was evaluated under an arbitrary default configuration while another received extensive optimization. Model-level metaheuristic calibration therefore contributes to fairer and more systematic predictive comparison.

The multi-component forecasting architecture included in the table assigns different temporal-modeling functions to two Transformer branches. Self-attention and auto-correlation mechanisms capture complementary dependencies over short and long forecasting horizons. This study illustrates modular prediction rather than metaheuristic feature selection. Its relevance to intelligent energy forecasting is therefore architectural: different processing branches can address different temporal structures. Such a design still requires component-level evaluation to establish whether each branch contributes enough predictive value to justify its computational cost.

The explainable hybrid prediction methodology introduces interpretability as an explicit architectural requirement rather than treating the predictive output as the final endpoint. Feature selection and ensemble prediction can generate powerful decision systems, but complex models can make it difficult to understand the basis on which particular outputs are produced. An explainable framework attempts to supplement predictive capability with mechanisms that provide insight into the factors influencing the model. This represents a methodological shift because model evaluation is no longer confined to determining whether predictions are correct; the analytical framework must also provide information that supports interpretation. Intelligent energy management can particularly benefit from this principle when automated recommendations affect building operation, industrial processes, distributed resources, or economically significant energy decisions. Operators may require an explanation of the conditions that contributed to a recommendation before accepting or acting upon it. Explainability therefore becomes a separate functional component of the intelligent system rather than simply a desirable descriptive property.

The connected-device hybrid methodology demonstrates how sensing infrastructure can be integrated with optimization and supervised learning within one analytical environment. The method combines an Internet of Things-enabled data-acquisition context with a TriKernel support-vector predictor whose hyperparameters are tuned using Levy Gazelle optimization. Its methodological significance is that the analytical model is positioned within a connected sensing ecosystem rather than treated as an isolated offline experiment. Data acquisition, communication, optimization, and inference therefore belong to the same operational chain. For intelligent energy management, this type of architecture closely reflects the conditions encountered in practical smart-energy environments, where measurements originate from distributed devices and must travel through communication infrastructure before analytical decisions can be generated. The quality of the final model consequently depends not only on the learning procedure but also on the availability and reliability of the sensing layer that supplies its inputs.

The metaheuristic-enhanced extreme learning machine methodology illustrates the application of optimization-assisted prediction to continuously varying environmental measurements. This methodology is useful for understanding how a relatively efficient learning architecture can be combined with a search mechanism when the target phenomenon exhibits nonlinear variation. The optimizer can contribute to identifying an improved configuration, while the learning model provides the mapping from observed variables to the continuous target quantity. Methodologically, this demonstrates that metaheuristic-assisted learning is not restricted to categorical classification problems. The same computational principle can be applied to regression and continuous estimation. This distinction is highly relevant to intelligent energy management because many core variables, including load, generation, temperature-related demand, storage-related quantities, and resource availability, are continuous rather than categorical. Optimization-assisted regression therefore represents a different methodological pathway from intrusion classification or disease detection.

The hybrid modified Grey Wolf Optimizer methodology combined with an optimized decision tree provides another example of distributed optimization, but with a structurally interpretable learner. The feature-selection component controls the input representation, while the decision-tree stage produces the subsequent risk-oriented prediction. This type of architecture is methodologically interesting because it balances search-based dimensional reduction with a model whose decision structure can generally be inspected more readily than that of many deep architectures. The resulting framework demonstrates that hybrid optimization does not inherently require a highly complex downstream learner. Intelligent computational systems can instead combine advanced search with comparatively transparent predictive mechanisms when interpretability, implementation simplicity, or explicit decision structure is valuable. For intelligent energy management, such a design can be relevant when operators require understandable decision rules in addition to predictive capability.

The physical-source localization methodology represents a different role for metaheuristic optimization because the candidate solutions are source coordinates rather than feature subsets or model weights. A hybrid GA–PSO procedure uses microseismic observations to estimate source location, with adaptive swarm parameters and evolutionary operations controlling the search. Unlike a deep-learning pipeline, this approach performs optimization directly over a physical solution space. For intelligent energy management, the transferable idea is that metaheuristics can support physical-state estimation and operational search tasks, not only preprocessing or predictive-model construction.

The final susceptibility-mapping methodology combines CatBoost with Whale and Zebra optimization algorithms to construct spatially oriented predictive models. Its relevance lies in the integration of a structured predictive learner with optimization for a problem characterized by heterogeneous geographic information. Spatial prediction differs from conventional independent-sample classification because nearby observations may share environmental characteristics and because the resulting output must often be interpreted over a geographic surface. The use of hybrid models in this setting illustrates how metaheuristic optimization can support structured prediction beyond purely tabular or temporal datasets. Intelligent energy-management systems increasingly incorporate geographically distributed resources, network components, renewable-generation sites, and regional demand information. Although the original application differs, the methodological principle demonstrates the adaptability of optimization-assisted learning to spatially structured decision environments.

When the methodologies in Table 1 are considered collectively, a clear functional classification emerges. Some methods optimize the information entering the predictive system; others optimize the model that interprets that information; another group combines multiple computational mechanisms to distribute responsibilities across a larger architecture; and still others apply optimization directly to a physical or decision-oriented search space. These categories should not be collapsed into one general notion of "metaheuristic optimization" because the mathematical representation of the candidate solution, the objective function, and the interpretation of the optimized output differ substantially. A binary feature-selection vector represents a fundamentally different decision object from a set of model parameters or a candidate physical location. Recognizing these distinctions is essential when transferring methodologies into intelligent energy-management applications.

The reviewed methodologies also reveal an important progression in how computational intelligence is organized. Early or comparatively simple architectures can be represented as a sequence in which data are first processed and then passed to a learner. More integrated approaches create feedback between the optimizer and the learner, meaning that candidate solutions are evaluated according to model behavior. Multi-component systems extend this interaction by assigning separate responsibilities to several algorithms, while explainable and sensing-integrated frameworks incorporate additional requirements beyond prediction itself. This progression indicates that the concept of an intelligent system is becoming increasingly architectural. Performance is determined by the coordination of computational functions rather than by the identity of one isolated algorithm.

From the perspective of intelligent energy management, the methodological evidence contained in the table suggests that future computational architectures can be designed according to the specific location of uncertainty or complexity within the management process. If the primary problem is excessive or redundant sensing information, optimization may be most useful at the feature-selection stage. If the input representation is sufficiently informative but model configuration is difficult, search can be directed toward hyperparameter or architectural calibration. When several objectives must be balanced, multi-objective optimization provides an alternative to a single aggregated fitness criterion. If decision sequences involve distinct analytical responsibilities, modular hybrid frameworks may be more appropriate. When operator trust is important, an explainability layer can be integrated after prediction. If measurements originate from distributed infrastructure, sensing and communication must become explicit elements of the architecture rather than invisible assumptions.

This functional interpretation also prevents unnecessary algorithmic complexity. A common risk in hybrid computational research is the assumption that adding another optimizer or predictive component will inherently improve the system. The methodologies summarized in the table instead suggest that every additional component should be justified according to a unique analytical function. A feature-selection module should solve an information-dimensionality problem; a model optimizer should solve a configuration problem; an attention mechanism should address selective representation; an explainability component should address interpretive transparency; and a physical search mechanism should solve a decision-space problem. When two components perform essentially the same role without a clear reason for their coexistence, architectural complexity can increase without providing a correspondingly distinct methodological contribution.

Another insight derived from the combined methodologies concerns the importance of stage-specific evaluation. Because hybrid architectures distribute intelligence across several components, evaluating only the final prediction may obscure the contribution of intermediate stages. An optimizer used for feature selection should be evaluated according to the quality and compactness of the selected representation. A model-calibration procedure should be examined according to the improvement obtained over an equivalent non-optimized configuration. A hybrid search should demonstrate that the combination provides value beyond its individual components. An explainability mechanism should be assessed according to whether it produces meaningful and consistent interpretations. A sensing-integrated system should consider whether the complete acquisition-to-inference chain can operate under realistic conditions. This stage-specific perspective is essential for determining whether a complex architecture provides genuine functional improvement.

Computational efficiency forms another methodological consideration that becomes increasingly important as architectures become more elaborate. Metaheuristic algorithms can require repeated evaluation of candidate solutions, and when each candidate involves training or validating a learning model, the total computational burden can become substantial. A multi-stage architecture containing feature selection, model optimization, and deep learning can therefore incur a much larger computational cost than a conventional learner. In energy-management environments requiring frequent updates or operational decisions within restricted time intervals, this cost can become an important practical constraint. Consequently, methodological assessment should consider not only whether optimization improves predictive behavior but also whether the additional search effort remains compatible with the intended deployment environment.

The methodologies also highlight the difference between offline model development and online intelligent operation. Feature selection and extensive hyperparameter optimization can often be performed offline because they are associated primarily with model construction. In contrast, physical search, adaptive decision-making, or some forms of operational optimization may need to execute repeatedly while the system is active. This distinction affects the acceptable computational budget. An algorithm that requires substantial processing may remain practical during occasional model retraining but become unsuitable when invoked for every real-time management decision. Intelligent energy-management design should therefore position computationally expensive procedures at stages where their cost can be tolerated and reserve faster decision mechanisms for time-critical operation.

Finally, the methodological references summarized in Table 1 demonstrate that the broader contribution of metaheuristic optimization is its capacity to provide a flexible search layer between complex decision spaces and intelligent predictive systems. Its role changes according to the problem being addressed: it can identify variables, configure learning models, coordinate hybrid components, balance competing objectives, support physical-state estimation, or participate within connected sensing architectures. The methodological value of metaheuristics should therefore be evaluated according to this functional role rather than according to algorithmic novelty alone. For intelligent energy management, such a perspective provides a more rigorous foundation for selecting computational methods because the optimizer can be chosen according to the structure of the energy-management problem, the nature of its decision variables, the operational time scale, and the computational constraints of the deployment environment.

Accordingly, the methodology matrix should be interpreted as evidence of several distinct computational design patterns rather than as a ranking of algorithms. The selected references collectively illustrate how intelligent systems can be constructed through different combinations of information selection, predictive learning, model calibration, hybrid search, explainable decision support, connected sensing, and physical-space optimization. These patterns provide methodological building blocks that can be adapted to energy-management problems when their functional relevance is clearly established. Importantly, the transfer of a methodology from another application domain should preserve the computational principle while allowing the objective function, constraints, input representation, and evaluation criteria to be reformulated according to the physical and operational characteristics of the target energy system.

2.1 Real-Time Energy State Awareness and Digital Representation

An increasingly important direction in intelligent energy management concerns the construction of continuously updated representations of the current operating state of an energy system. Predictive modeling alone provides information regarding expected future behavior, but an intelligent management platform must first possess an accurate representation of what is occurring within the physical environment at the present moment. This requirement has encouraged the development of architectures in which measurements collected from meters, sensors, distributed devices, and control systems are continuously integrated into a unified digital representation of the managed environment. Such representations can describe consumption levels, generation availability, storage conditions, equipment status, network operating variables, and other state-dependent quantities that influence subsequent management decisions.

The methodological importance of real-time state awareness arises from the difference between historical analysis and operational intelligence. Historical datasets are useful for model development and long-term pattern discovery, whereas operational management requires information reflecting the actual state of the system at the instant a decision is generated. Delays, missing measurements, synchronization problems, and differences in sensing frequency can therefore directly influence decision quality. Intelligent energy-management architectures must consequently incorporate mechanisms capable of organizing asynchronous measurements, distinguishing outdated information from current observations, and maintaining a consistent representation of the energy environment.

Digital representations can further provide an intermediate analytical layer between physical energy assets and computational decision mechanisms. Rather than allowing every predictive or optimization algorithm to interact directly with heterogeneous sensors, a unified state model can supply standardized information to multiple computational components. This architecture can simplify integration because forecasting models, control procedures, scheduling algorithms, and anomaly-detection mechanisms operate on a shared representation of the physical environment. It also permits the management system to separate the acquisition layer from the analytical layer, reducing direct dependence between individual sensors and specific algorithms.

The usefulness of such a digital representation becomes especially significant in environments containing numerous distributed components. A microgrid, for example, may simultaneously contain local generators, renewable resources, energy-storage systems, flexible loads, metering devices, and power-electronic interfaces. An industrial facility may contain several production units whose energy requirements differ according to operating state, while a smart building may contain multiple zones with distinct thermal and electrical characteristics. In these environments, the management problem cannot be represented adequately through one aggregated consumption measurement. Instead, intelligent operation depends on understanding the relationships among several interacting subsystems.

A continuously updated energy-state representation also provides a foundation for detecting inconsistencies between expected and observed behavior. When the management platform maintains an estimate of how the system should behave under particular conditions, deviations can be identified as operational abnormalities requiring additional examination. Importantly, this form of state awareness should be distinguished from the predictive and anomaly-related mechanisms themselves. The digital representation provides the organized operational context upon which those mechanisms operate rather than serving as the final analytical decision.

Another methodological issue involves temporal resolution. Different energy-management decisions operate at different time scales. Some control actions may require second-level or minute-level information, while other planning decisions may rely on hourly, daily, or seasonal patterns. A real-time representation must therefore support multiple temporal resolutions without unnecessarily processing every measurement at the highest available frequency. This creates a data-management problem in which information must be aggregated or retained according to its operational relevance. High-frequency electrical measurements may be necessary for immediate equipment control, whereas long-term demand planning may require only aggregated consumption trends.

The concept of a digital representation can be extended further toward digital-twin-oriented energy management, in which the computational representation does not merely reproduce measured states but also models the behavior of the underlying physical system. In such an architecture, the digital counterpart can be used to evaluate potential actions before they are implemented physically. Candidate scheduling decisions, control modifications, or resource allocations can be tested computationally to estimate their possible consequences. This creates an additional decision-support layer between optimization and physical execution because the management platform can evaluate proposed actions within a simulated representation before applying them to actual energy assets.

Such an arrangement is particularly valuable when operational mistakes carry significant economic or technical consequences. Direct experimentation on real energy infrastructure may be impractical when an inappropriate control action could increase demand peaks, violate equipment constraints, or reduce system reliability. A sufficiently representative digital environment can provide a safer mechanism for examining alternative actions. However, its usefulness depends heavily on how accurately the digital representation reflects changes in the physical system. A model that becomes outdated can create false confidence by evaluating decisions according to conditions that no longer correspond to reality.

The synchronization between physical and digital states therefore becomes a central methodological requirement. Intelligent energy-management architectures incorporating digital representations must specify how frequently the computational model is updated, how conflicting measurements are resolved, and how uncertainty in the observed state is represented. A system should not necessarily interpret every measured value as perfectly accurate. Sensor noise, calibration error, communication interruption, and incomplete observations can introduce uncertainty that should be reflected in the digital representation rather than ignored.

Overall, real-time state awareness introduces an important layer of intelligence that is conceptually different from prediction, feature selection, or optimization. It establishes the operational context within which these methods function. Intelligent energy management therefore benefits from architectures capable of maintaining a coherent and continuously updated description of the managed environment, thereby ensuring that subsequent computational decisions correspond as closely as possible to the actual condition of the physical energy system.

2.2 Human-Aware and Preference-Constrained Energy Management

Intelligent energy management cannot be evaluated exclusively according to technical energy reduction because many managed environments directly serve human users whose preferences, behavioral patterns, and operational requirements influence the feasibility of automated decisions. Residential buildings, commercial spaces, educational facilities, and workplaces all contain energy-consuming activities that exist to provide services rather than simply to consume electricity. Heating, cooling, lighting, ventilation, appliance operation, and electric-vehicle charging are therefore connected to comfort, convenience, productivity, mobility, and individual behavioral expectations. An energy-management strategy that minimizes consumption while substantially reducing these services may be computationally efficient but practically unacceptable.

Human-aware management introduces user-related constraints directly into the decision process rather than treating human behavior as an unpredictable disturbance. This distinction is important because conventional optimization may formulate energy consumption or electricity cost as the principal objective while representing user requirements through fixed limits. More advanced management can instead account for preferences dynamically, allowing acceptable operating ranges to change according to occupancy, activity, time of day, or expressed user priorities. The resulting energy-management problem becomes a negotiation between resource efficiency and service quality rather than a one-dimensional minimization task.

Residential environments provide a particularly clear example. Flexible appliances can often be shifted away from expensive or high-demand periods, but the acceptable degree of shifting varies considerably according to the function of the appliance and household routine. Some devices can be delayed for several hours with minimal inconvenience, whereas others are closely connected to immediate household requirements. Treating all controllable loads as equally flexible can therefore produce schedules that appear optimal mathematically but are unlikely to be accepted by occupants. Intelligent scheduling should consequently distinguish technical flexibility from behavioral flexibility.

Preference representation creates its own methodological challenges. Explicit preferences can be obtained directly when users specify desired temperatures, appliance deadlines, charging requirements, or acceptable cost limits. However, continuously requesting manual input can reduce usability and increase interaction burden. Intelligent systems can instead attempt to infer recurring preferences from observed behavior, although this introduces uncertainty because historical behavior does not always represent current intention. A robust management architecture can therefore combine explicit constraints with learned behavioral tendencies while retaining mechanisms through which users can override automated decisions when necessary.

The treatment of comfort is especially important in building energy management. Thermal comfort is influenced by more than temperature alone and can vary across individuals and operating environments. Similarly, lighting preferences, ventilation expectations, and equipment-use patterns may differ across occupants. A management platform that optimizes a building according to one standardized comfort target may consequently produce uneven satisfaction across users. Human-aware systems require flexible representations of acceptable operating conditions rather than assuming a universally preferred state.

Human involvement is also relevant to the level of automation assigned to the intelligent system. Fully autonomous energy management can execute decisions without human approval, whereas decision-support systems may present recommendations that require confirmation before implementation. Between these extremes, semi-autonomous architectures can execute routine decisions automatically while escalating unusual or high-impact actions to an operator. The appropriate level of autonomy depends on the managed environment, the consequences of incorrect decisions, and the degree of trust placed in the computational model.

This consideration becomes particularly significant in industrial energy management. Energy optimization may recommend modifying equipment schedules, process timing, or production sequences, but these actions can affect manufacturing targets, safety procedures, maintenance plans, or product quality. Energy efficiency therefore cannot be optimized independently from operational priorities. Human operators possess contextual knowledge that may not be represented completely within the analytical model, making operator interaction an important component of practical deployment. Intelligent systems should consequently be capable of supporting expert decisions rather than assuming that computational optimization can replace all operational judgment.

User acceptance represents another important factor. An intelligent system may provide technically strong recommendations yet fail in practice if users repeatedly override its decisions. Frequent override behavior can indicate that the model does not adequately represent actual preferences or operating constraints. Such interaction data can itself become useful information for refining future management policies. A human-aware system can therefore interpret user interventions not simply as failures to follow recommendations but as feedback describing deficiencies in the existing decision model.

The incorporation of preferences also creates fairness considerations when energy resources are shared among multiple users. In multi-occupant buildings or community-level energy systems, one management decision can benefit some participants while imposing inconvenience or cost on others. A purely aggregated objective may conceal these differences. Intelligent management should therefore consider how benefits and burdens are distributed when resources or flexibility are shared. This issue becomes increasingly important in collective energy environments where storage, generation, or flexible demand is coordinated across several participants.

Another methodological requirement concerns transparency of automated decisions. Users are more likely to understand and evaluate a recommendation when the system provides a meaningful explanation of why an action has been proposed. For example, a recommendation to delay a controllable load can be accompanied by information indicating that the action avoids a high-demand interval or uses expected local renewable generation. Such explanations can improve the interpretability of automated management without requiring users to understand the underlying computational algorithm.

Human-aware intelligent energy management therefore broadens the concept of performance beyond energy and cost indicators. A practically effective system must also consider acceptability, intervention frequency, service preservation, and alignment with user priorities. These requirements should not be treated as secondary considerations introduced only after optimization. Instead, they can form explicit elements of the decision model so that the resulting management strategy remains technically effective while respecting the purpose for which energy is ultimately being consumed.

2.3 Multi-Energy Coordination and Sector-Coupled Intelligent Management

A further development in intelligent energy management concerns the transition from electricity-centered management toward coordinated operation of multiple energy carriers and interconnected sectors. Many real energy environments do not operate exclusively through electricity. Buildings can simultaneously use electrical energy, thermal energy, gas-based heating, stored thermal capacity, and locally generated renewable energy. Industrial facilities may consume electricity, heat, steam, compressed air, and fuel within the same production process. Transportation electrification introduces another connection by linking mobility demand directly with electrical infrastructure. Intelligent energy management therefore increasingly requires coordination across energy forms whose operational behaviors and storage possibilities differ substantially.

Multi-energy coordination changes the structure of the management problem because one energy service may be supplied through more than one pathway. Heating, for example, can be provided through electrical heat pumps, resistive heating, district thermal networks, or fuel-based systems. The management platform can therefore select not only when energy should be used but also which energy carrier or conversion pathway should satisfy a particular service requirement. This flexibility can create substantial opportunities for optimization because decisions can account for relative availability, cost, efficiency, and operational constraints across several energy sources.

Energy conversion technologies introduce coupling between previously separate infrastructures. A heat pump converts electrical input into thermal output, while combined heat and power systems produce multiple useful energy outputs from a shared process. Electrolysis can convert electricity into hydrogen, and stored hydrogen can subsequently support industrial, transportation, or power-generation applications. Battery systems shift electricity temporally, whereas thermal storage can shift heating or cooling demand. These conversion and storage mechanisms create a network of possible energy pathways, making management increasingly dependent on coordination across both time and energy form.

The presence of multiple storage technologies further increases decision complexity. Electrical batteries and thermal-storage systems possess different efficiency characteristics, operating constraints, degradation mechanisms, and response times. A management platform can determine which resource should absorb surplus energy according to expected future requirements. Storing electricity directly may be preferable when future electrical demand is anticipated, whereas converting surplus electricity into thermal energy may provide greater value when heating or cooling requirements are expected later. Intelligent coordination therefore requires forecasts covering more than one demand stream and optimization capable of representing interactions among different storage resources.

Sector coupling becomes particularly significant with the expansion of electric mobility. Electric vehicles introduce substantial flexible demand but can also function as distributed storage resources when appropriate technologies and operational agreements are available. Their availability is constrained by mobility requirements, meaning that an energy-management system must coordinate charging decisions with expected departure times and required driving range. At larger scales, widespread electric-vehicle charging can influence distribution-network loading and create new demand peaks if unmanaged. Coordinated management can instead distribute charging across time or align it with periods of favorable generation availability.

Industrial sector coupling presents a different challenge because energy consumption is closely tied to production processes. Certain energy-intensive operations may possess scheduling flexibility, while others must follow strict process sequences. Waste heat from one process may potentially support another thermal requirement, creating opportunities for internal energy recovery. Intelligent energy management in such environments therefore requires an understanding of material flow and production dependencies in addition to energy flow. Optimizing energy independently from production can generate infeasible solutions because the timing and quantity of energy demand are determined partly by operational requirements.

The increasing integration of renewable generation strengthens the importance of multi-energy flexibility. Variable renewable output may exceed immediate electrical demand during certain periods while falling below demand during others. Sector-coupled systems provide additional pathways through which surplus generation can be utilized rather than curtailed. Electricity can be directed toward storage, thermal loads, mobility charging, or conversion processes when conditions permit. The availability of several flexible consumption pathways can therefore increase the capacity of the overall system to accommodate variable renewable production.

However, multi-energy management also introduces substantial modeling complexity. Each energy carrier possesses different physical properties, efficiency relationships, conversion constraints, and temporal dynamics. Electrical systems respond almost instantaneously to changes in generation and load, whereas thermal systems often possess considerable inertia. Gas and hydrogen infrastructures exhibit different storage and transport characteristics. A unified management model must therefore represent these differences without simplifying them to the point that resulting decisions become physically unrealistic.

The optimization horizon is another important consideration. Short-term decisions may prioritize immediate balancing and equipment constraints, while longer-term planning can consider storage trajectories, anticipated production requirements, weather variation, or mobility schedules. Coordinating multiple energy carriers across several time horizons increases the dimensionality of the decision problem substantially. Hierarchical management architectures can address this challenge by separating long-term planning from short-term control while ensuring that decisions at different levels remain consistent.

Interoperability becomes particularly important in multi-energy environments because individual subsystems may be managed by different controllers, software platforms, or organizations. An electrical management platform may not automatically exchange information with thermal or mobility systems. Intelligent coordination therefore requires standardized representations through which information concerning demand, availability, constraints, and planned operation can be shared. Without such interoperability, theoretically flexible resources may remain operationally isolated.

At community or district scale, sector-coupled management can further involve multiple ownership structures. Electrical storage may belong to one participant, renewable generation to another, and thermal infrastructure to a shared operator. The management problem therefore includes institutional and economic boundaries in addition to technical constraints. Decision mechanisms must specify how shared benefits, operational responsibilities, and resource access are coordinated. This creates a transition from purely engineering optimization toward socio-technical energy management.

Multi-energy coordination consequently represents a substantial expansion of the intelligent energy-management problem. Instead of optimizing one energy stream, the management system must identify interactions among several forms of energy, conversion technologies, storage devices, and end-use sectors. The resulting flexibility can improve utilization of available resources and support higher renewable-energy integration, but it requires models capable of representing physical coupling accurately and optimization procedures capable of navigating a substantially enlarged decision space. This direction therefore provides a distinct pathway for intelligent energy-management research in which computational intelligence is applied not only to increasingly complex algorithms but also to increasingly interconnected physical energy infrastructures.

3. DISCUSSION

The evidence synthesized throughout this review indicates that the development of intelligent energy management is progressing toward computational systems whose effectiveness depends increasingly on the quality of coordination between analytical intelligence and the physical characteristics of the managed energy environment. The central research challenge is therefore no longer restricted to identifying models capable of producing favorable predictive results. A more consequential issue concerns whether the analytical mechanisms incorporated within an intelligent energy-management system remain physically meaningful, operationally feasible, and sufficiently reliable when their outputs are translated into real actions. This distinction is particularly important because energy-management decisions directly influence physical infrastructure whose behavior is governed by equipment limitations, network constraints, storage dynamics, thermal processes, operational schedules, and safety requirements. A computational model can therefore produce a mathematically attractive recommendation that remains unsuitable for practical implementation if the physical characteristics of the underlying energy system are not represented adequately.

One of the most important directions emerging from this observation concerns the incorporation of physical knowledge into intelligent computational models. Many machine-learning architectures are primarily data-driven and attempt to infer relationships directly from observations without explicitly representing the physical principles governing the target system. Although such flexibility can be advantageous when analytical relationships are complex, exclusively data-driven models may generate predictions that appear statistically plausible while violating fundamental operational relationships. Intelligent energy-management research can address this limitation by developing computational structures in which physical constraints, conservation relationships, equipment characteristics, or known system dynamics are incorporated directly into the learning or decision process. Such integration can reduce the possibility of generating analytically strong but physically unrealistic outputs.

The use of physical knowledge also provides a pathway for improving the reliability of intelligent models when available datasets do not adequately represent every possible operating condition. Energy systems can experience rare configurations that are difficult to capture comprehensively during model training. Extreme weather conditions, unusual consumption events, abnormal resource availability, or unexpected equipment behavior may produce states that are poorly represented within historical observations. A purely data-driven model may therefore be forced to extrapolate beyond its learned experience. Physical constraints can provide an additional source of structure in such situations by restricting predictions to ranges that remain consistent with known system behavior. The relationship between physical modeling and artificial intelligence should consequently be considered complementary rather than competitive. Physical models provide mechanistic structure, while data-driven methods provide adaptability to relationships that may be difficult to describe analytically.

A second major issue concerns uncertainty. Intelligent energy-management systems frequently rely on forecasts that cannot be assumed to be perfectly accurate. Future consumption, renewable-energy production, electricity-market conditions, occupancy, environmental variables, and equipment availability may all contain substantial uncertainty. Nevertheless, many computational frameworks represent predicted values as deterministic quantities and subsequently optimize management decisions as though the forecasts were exact. This treatment can create fragile decisions because the selected operating strategy may perform well only when the predicted future state occurs precisely as expected.

Uncertainty-aware energy management requires a different decision philosophy. Rather than searching exclusively for the best action under one predicted scenario, the management framework should consider the range of plausible future conditions that may occur. A decision that provides slightly lower nominal performance but remains effective across several plausible scenarios may offer greater operational value than a highly optimized strategy that becomes unsuitable following a small forecasting error. The distinction between nominal optimality and robust operational suitability is therefore fundamental. Intelligent systems should increasingly communicate not only what they predict but also how confident they are in those predictions and how sensitive the recommended action is to forecast error.

This requirement introduces probabilistic reasoning into the management process. Prediction intervals, scenario distributions, confidence measures, or other representations of uncertainty can allow the decision layer to distinguish high-confidence operating conditions from circumstances in which the available information remains ambiguous. Management policies could subsequently adopt different levels of conservatism according to uncertainty. When predictions are highly reliable, the system may exploit available flexibility more aggressively. When uncertainty increases, the management mechanism may preserve additional operational margins and avoid decisions that depend heavily on one narrowly predicted outcome. Such behavior would enable intelligent energy management to respond not only to expected conditions but also to the quality of the information on which those expectations are based.

A related research issue concerns the verification of automated decisions before physical implementation. As intelligent energy-management systems become increasingly autonomous, the consequences of computational errors become more significant. A prediction error may remain an analytical problem if it is used only for reporting, but it can become an operational problem when the predicted quantity determines an automated control action. The relationship between artificial intelligence and control therefore creates a need for explicit safety mechanisms capable of examining proposed decisions before execution.

One promising architectural principle is the separation between decision generation and decision authorization. An intelligent model can propose an action, while an independent verification layer evaluates whether that action satisfies predefined physical and operational constraints. Such a layer does not need to reproduce the complete optimization process. Its purpose is instead to determine whether the proposed action remains within an acceptable operating envelope. This architecture provides a safeguard against unexpected model behavior and prevents optimization outputs from being implemented directly when they violate essential restrictions.

Safety verification becomes increasingly important as decision frequency increases. A management system generating one recommendation per day provides considerable opportunity for manual inspection, whereas a controller operating at short intervals may produce thousands of decisions during continuous operation. Human examination of every action becomes unrealistic under such conditions. Automated verification can therefore serve as a computational intermediary between artificial intelligence and physical execution. The presence of such a mechanism can also support the gradual introduction of greater autonomy because organizations can permit automated decision generation while retaining independent control over the boundary within which those decisions may operate.

Another issue requiring greater attention is the distinction between laboratory optimization and operational value. Academic studies frequently evaluate intelligent models using fixed datasets and predefined experimental partitions. Such evaluation is essential for methodological comparison, but real energy-management systems operate continuously and interact with changing physical environments. The operational value of a computational framework therefore depends on characteristics that may not be fully visible within conventional offline testing. A model must remain useful after deployment, interact reliably with physical infrastructure, maintain acceptable behavior during unusual conditions, and support decisions at the speed required by the target application.

This difference suggests that intelligent energy-management research should increasingly consider deployment-oriented evaluation environments. Hardware-in-the-loop testing, controlled operational pilots, realistic simulations, and staged implementation can provide intermediate validation levels between purely offline experiments and full deployment. Such approaches allow researchers to examine how models behave when exposed to communication delays, imperfect measurements, actuator limitations, and operational interruptions that may be absent from static datasets. The objective is not to replace conventional model evaluation but to complement it with evidence concerning practical behavior.

The transition from academic evaluation toward deployment also requires consideration of model lifecycle management. An intelligent energy-management model should not be viewed as a finished artifact immediately after training. Its performance may change after deployment because the environment evolves, sensing configurations change, equipment is replaced, and operational objectives are modified. Lifecycle management therefore includes procedures for detecting model degradation, deciding when updates are necessary, validating newly trained versions, and controlling the transition between model generations.

This introduces an important governance challenge. Updating an operational model cannot always be treated as a routine software modification because a newly trained model may produce different decisions even when supplied with similar inputs. Organizations deploying intelligent energy management therefore require mechanisms for version control, validation, rollback, and audit. A new model should be evaluated before replacing the currently deployed version, and previous configurations should remain recoverable if unexpected behavior occurs. Such lifecycle governance becomes especially important in large infrastructures where several analytical models may operate simultaneously across different equipment groups or locations.

The environmental cost of computational intelligence itself represents another issue that deserves greater consideration within energy-management research. Intelligent energy management is commonly motivated partly by the objective of improving sustainability, yet increasingly complex learning architectures and extensive optimization processes can require substantial computational resources. Repeated model training, large population-based optimization procedures, extensive parameter searches, and high-frequency inference all consume electricity. The energy consumed by the computational solution should therefore become part of the broader assessment when the purpose of the system is to reduce energy use or environmental impact.

This issue does not imply that complex models should be avoided. Instead, computational cost should be evaluated relative to the energy benefit produced by the management system. A computationally expensive optimization process may be entirely justified if it generates substantial long-term energy savings and is executed infrequently. Conversely, marginal predictive improvement may provide limited practical value if achieving that improvement requires dramatically greater computational resources. Energy-aware artificial intelligence therefore introduces an additional dimension of methodological evaluation in which the computational footprint of the analytical system is considered alongside its operational benefit.

The location at which computation is performed also becomes relevant. Cloud computing can provide substantial processing capacity but requires communication between the managed energy environment and remote computational infrastructure. Local computation can reduce dependence on continuous connectivity and support faster responses but may operate under restricted hardware resources. Distributed architectures can divide responsibilities between local and remote processing according to time sensitivity and computational complexity. The appropriate allocation of computation therefore depends on the characteristics of the management task rather than on the assumption that one infrastructure model is universally preferable.

Energy-management intelligence must also operate within economic and institutional structures rather than solely within technical systems. Electricity tariffs, market participation rules, regulatory constraints, contractual arrangements, and ownership structures can determine which actions are practically permissible even when they are technically feasible. An intelligent controller cannot optimize energy exchange independently from the rules governing that exchange. Similarly, flexible demand may possess technical capability that cannot be utilized fully because contractual arrangements do not permit dynamic operation. The integration of institutional constraints into intelligent management therefore represents an important step toward models that reflect the actual decision environment.

Policy changes can further modify the objective of an intelligent energy-management system. Financial incentives for renewable generation, dynamic electricity tariffs, demand-response programs, carbon pricing, and network-service markets can alter the economic value of particular actions. An intelligent management framework intended for long-term operation should consequently avoid assuming that present economic relationships will remain unchanged indefinitely. Flexible objective formulations can enable management systems to adapt when pricing structures or regulatory requirements evolve.

This institutional perspective becomes particularly significant when intelligent energy management operates across organizational boundaries. The party owning a distributed resource may differ from the party managing the network, purchasing the electricity, or benefiting from demand flexibility. An optimization that minimizes system-wide cost may not minimize the cost experienced by every participant. Practical implementation therefore requires clarity concerning whose objective is being optimized and how benefits are distributed among stakeholders. This question cannot be resolved purely through improved predictive accuracy because it concerns the governance of shared energy resources.

Another important research direction concerns the appropriate level of abstraction at which intelligent management decisions should be generated. Highly detailed models can represent individual devices and operating states but may become computationally difficult to scale. More aggregated models can reduce complexity but may conceal constraints that determine whether a proposed action is physically feasible. Intelligent energy management therefore requires a balance between analytical detail and tractable decision-making. The appropriate level of representation depends on the decision being supported. Strategic planning may tolerate aggregated representations, whereas equipment-level control may require detailed models of individual resources.

Hierarchical decision architectures provide one potential solution to this problem. Higher-level components can determine broad operational targets, while lower-level controllers translate those targets into resource-specific actions. Such architectures allow different computational methods to operate at different temporal and spatial scales without requiring one centralized model to represent every decision variable simultaneously. The challenge lies in maintaining consistency between levels. A high-level target should remain achievable given lower-level physical constraints, while local decisions should contribute collectively to the broader management objective.

The discussion also reveals that intelligent energy management should be assessed according to decision quality rather than predictive performance alone. Prediction is valuable because it informs action, but the ultimate purpose of an energy-management system is to improve operation. A model with slightly lower forecasting accuracy could theoretically produce better operational outcomes if its errors occur in regions that have little influence on management decisions. Conversely, a highly accurate predictor may still lead to poor decisions if relatively small errors occur at critical operating points. Evaluation frameworks should therefore increasingly examine the downstream effect of predictive errors rather than interpreting all errors as equally important.

Decision-oriented evaluation would connect analytical metrics with operational consequences. Instead of asking only whether a forecast deviates from the observed value, researchers can examine whether the deviation changes the selected schedule, resource allocation, or control action. This approach provides a closer relationship between machine-learning assessment and energy-system performance. It may also alter model-selection conclusions because the model achieving the lowest statistical error is not necessarily the model producing the most beneficial management decisions.

The broader implication of these observations is that intelligent energy management is moving toward a stage in which methodological maturity depends less on introducing additional algorithms and more on integrating computational intelligence responsibly with physical systems. Advances in prediction and optimization have established substantial analytical capability. The next challenge is ensuring that this capability produces decisions that remain meaningful under uncertainty, physically feasible, verifiable before execution, maintainable across operational lifecycles, economically appropriate, and environmentally justified. Addressing these requirements would shift intelligent energy management from experimentally successful computational modeling toward dependable operational intelligence.

4. CONCLUSION

Intelligent energy management is evolving from a collection of isolated analytical techniques into a decision-oriented discipline in which computational intelligence must operate as an accountable component of real energy infrastructure. The principal significance of this evolution lies in the transition from demonstrating that intelligent algorithms can analyze energy-related information toward establishing that their outputs can support dependable operational decisions.

The review indicates that the next stage of progress should be defined by the quality of the connection between computational reasoning and physical energy operation. Intelligent systems will provide their greatest value when analytical capability is accompanied by mechanisms that preserve feasibility, quantify uncertainty, verify automated actions, and support controlled operation throughout the lifetime of the deployed model. Under this perspective, methodological advancement is measured not simply by increasing algorithmic sophistication but by improving the trustworthiness of the complete path from information acquisition to physical decision execution.

The future of intelligent energy management therefore depends on developing computational frameworks that can function as persistent operational components rather than temporary experimental models. Such systems must remain capable of supporting changing energy environments while retaining clear boundaries between analytical recommendation and physical authorization. They must also demonstrate that the resources consumed by computational intelligence are proportionate to the operational benefits that intelligence provides.

Ultimately, the transition toward genuinely intelligent energy systems will occur when artificial intelligence is no longer treated as an external analytical addition to conventional energy infrastructure but as a carefully governed decision layer operating in coordination with physical, economic, and institutional processes. Achieving this integration would establish intelligent energy management as a practical foundation for the long-term operation of increasingly complex energy environments and would move the field beyond algorithm-centered experimentation toward reliable computationally supported energy decision-making.

REFERENCES

[1] F. Amirghafouri, A. A. Neghabi, H. Shakeri, and Y. E. Sola, "Nature-Inspired Meta-Heuristic Algorithms for Resource Allocation in the Internet of Things," *International Journal of Communication Systems*, vol. 38, no. 5, p. e6141, 2025.

[2] T. Etem, M. R. Baker, and S. Buyrukoğlu, "Hybrid approach for traffic flow detection using metaheuristics optimization and machine learning algorithms," *Journal of the Chinese Institute of Engineers*, pp. 1–21, 2025.

[3] H. Ouyang, X. Lin, S. Li, L. Gao, and E. H. Houssein, "Recent metaheuristic algorithms for multi-objective feature selection: review, applications, open issues and challenges," *Cluster Computing*, vol. 28, no. 7, p. 467, 2025.

[4] H. Von Linde and O. Riedel, "A methodology for evaluating feature selection and clustering methods with project-specific requirements," *International Journal of Production Research*, vol. 63, no. 5, pp. 1692–1706, 2025.

[5] M. Aly and A. S. Alotaibi, “Hybrid Butterfly-Grey Wolf Optimization (HB-GWO): A novel metaheuristic approach for feature selection in high-dimensional data,” *Statistics, Optimization & Information Computing*, vol. 13, no. 6, pp. 2575–2600, 2025.

[6] Zhang XinHao, “A novel hybrid DOA-PSO-enhanced LSSVM model for monthly runoff forecasting in the upper Heihe river basin,” *Scientific Reports*, vol. 15, Art. no. 35869, 2025.

[7] Q. Wen and Y. Liu, “Feature engineering and selection for prosumer electricity consumption and production forecasting: A comprehensive framework,” *Applied Energy*, vol. 381, Art. no. 125176, 2025.

[8] B. R. Nayana, R. Subha, R. Radhakrishnan, and P. Geethanjali, “Scalable bearing fault diagnosis using metaheuristic feature selection and machine learning for diverse operating conditions,” *Systems Science & Control Engineering*, vol. 13, no. 1, p. 2469606, 2025.

[9] J. P. Roselyn, “Golden eagle optimization approach for feature selection and XGBoost algorithm-based intrusion detection system in wireless mesh networks of smart grid.” *Journal of the Chinese Institute of Engineers*, vol. 48, no. 8, p. 1312, 2025.

[10] A. Dhiya Eddine, G. Abdelkader, and B. Mourad, “Enhancing IIoT security: a hybrid AI framework for intrusion detection powered by Gen-X AI,” *International Journal of Computers and Applications*, vol. 47, no. 9, pp. 749–765, 2025.

[11] M. Maazalahi and S. Hosseini, “Machine learning and metaheuristic optimization algorithms for feature selection and botnet attack detection,” *Knowledge and Information Systems*, vol. 67, no. 4, pp. 3549–3597, 2025.

[12] X. Gong, Y. Yang, Y. Zhang, N. Li, Y. Guan, and R. Jiang, “Feature selection method for network intrusion based on hybrid meta-heuristic dynamic optimization algorithm,” *Computers & Security*, vol. 156, p. 104512, 2025.

[13] Luo Shan, “(IoT) network intrusion detection system using optimization algorithms,” *Scientific Reports*, vol. 15, Art. no. 21706, 2025.

[14] S. Nagarajan, D. F. S. Jayapalan, and P. P. L. Devaraj, “Ensemble-based feature selection and optimization-driven deep learning for attack detection in cloud computing,” *Knowledge-Based Systems*, vol. 325, p. 113984, 2025.

[15] Y. Pourardebil Khah, M. Hosseini Shirvani, and H. Motameni, “A hybrid machine learning approach for feature selection in designing intrusion detection systems (IDS) model for distributed computing networks,” *The Journal of Supercomputing*, vol. 81, no. 1, p. 254, 2025.

[16] Y. Wang, Y. Yin, H. Zhao, J. Liu, C. Xu, and W. Dong, “Grey wolf optimizer with self-repulsion strategy for feature selection,” *Scientific Reports*, vol. 15, Art. no. 12807, 2025.

[17] K. Zhang, Y. Wang, U. A. Bhatti, Y. Zhou, and M. Jin, “Enhanced ransomware attacks detection using feature selection, sensitivity analysis, and optimized hybrid model,” *Journal of Big Data*, vol. 12, no. 1, p. 245, 2025.

[18] M. Liang, M. Xu, and S. Wang, “A novel multi-objective evolutionary algorithm for transit network design and frequency-setting problem considering passengers’ choice behaviors under station congestion,” *Transportation Research Part B: Methodological*, vol. 197, Art. no. 103238, 2025.

[19] S. Sharma, Prachi, R. Chhikara, and K. Khanna, “A novel feature selection technique: Detection and classification of Android malware,” *Egyptian Informatics Journal*, vol. 29, Art. no. 100618, 2025.

[20] Q. Zhou, H. Tang, C. Hibert, M. Chmiel, F. Walter, M. Dietze, and J. M. Turowski, “Enhancing debris flow warning via machine learning feature reduction and model selection,” *Journal of Geophysical Research: Earth Surface*, vol. 130, no. 4, Art. no. e2024JF008094, 2025.

[21] K. Veeranjaneyulu, M. Lakshmi, and S. Janakiraman, “Swarm Intelligent Metaheuristic Optimization Algorithms-Based Artificial Neural Network Models for Breast Cancer Diagnosis: Emerging Trends, Challenges and Future Research Directions,” *Archives of Computational Methods in Engineering*, vol. 32, no. 1, pp. 381–398, 2025.

[22] P. Dhal, B. Pradhan, U. Fiore, S. A. J. Francis, and D. S. Roy, “A clinical diabetes prediction based support system based on the multi-objective metaheuristic inspired fine tuning deep network,” *Information Fusion*, vol. 122, p. 103188, 2025.

[23] N. Ullah, F. Martínez-Álvarez, I. De Falco, and G. Sannino, “Optimizing deep learning for cotton leaf disease detection using meta-heuristic feature selection algorithms,” in *2025 20th Conference on Computer Science and Intelligence Systems (FedCSIS)*, IEEE, pp. 605–614, 2025.

[24] T. Abid, A. Ahmim, F. Maazouzi, D. Chefrour, I. Ullah, M. Ahmim, and R. Almukhlifi, “A novel IoT threat detection using GWO feature selection and CNN-enhanced LightGBM,” *Journal of Cloud Computing*, vol. 14, Art. no. 72, 2025.

[25] A. A. Soladoye, D. B. Olawade, A. O. Esan, N. Aderinto, B. A. Omodunbi, I. A. Adeyanju, and S. Boussios, “Enhancing Parkinson’s disease prediction using metaheuristic optimized machine learning models,” *Personalized Medicine*, vol. 22, no. 4, pp. 223–234, 2025.

[26] R. Younas, H. M. Raza Ur Rehman, and G. S. Choi, “Crypto foretell: A novel hybrid attention-correlation based forecasting approach for cryptocurrency,” *Journal of Big Data*, vol. 12, Art. no. 229, 2025.

[27] A. M. Elshewey, E. Selem, and A. H. Abed, "Improved CKD classification based on explainable artificial intelligence with extra trees and BBFS," *Scientific Reports*, vol. 15, Art. no. 17861, 2025.

[28] K. Dhanushkodi, P. Vinayagasundaram, V. Anbalagan, S. Subbaraj, and R. Sethuraman, "TriKSV-LG: A robust approach to disease prediction in healthcare systems using AI and Levy Gazelle optimization," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 28, no. 11, pp. 1783–1799, 2025.

[29] S. I. Abba, Q. B. Pham, A. Malik, R. Costache, M. S. Gaya, J. Abdullahi, S. Mati, A. G. Usman, and G. Saini, "Optimization of extreme learning machine with metaheuristic algorithms for modelling water quality parameters of Tamburawa Water Treatment Plant in Nigeria," *Water Resources Management*, vol. 39, no. 3, pp. 1377–1401, 2025.

[30] M. Sam'an, M. Mat Deris, and Farikhin, "Feature selection in peer-to-peer lending based on hybrid modified grey wolf optimization with optimized decision tree for credit risk assessment," *International Journal of Management Science and Engineering Management*, vol. 20, no. 1, pp. 141–158, 2025.

[31] Y. Han, F. Zeng, L. Fu, and F. Zheng, "GA-PSO algorithm for microseismic source location," *Applied Sciences*, vol. 15, no. 4, Art. no. 1841, 2025.

[32] S. V. Razavi-Termeh, A. Sadeghi-Niaraki, S. I. Abba, J. Hussain, and S.-M. Choi, "Flood-prone area mapping using a synergistic approach with swarm intelligence and gradient boosting algorithms," *Scientific Reports*, vol. 15, Art. no. 27924, 2025.