EAI Endorsed Transactions
on Energy Web
Research Article EAI.EU

Intelligent Distribution-Network Fault Handling via Cross-Modal Semantic Fusion and Knowledge-Graph Reasoning

Fangzhou Liu [1,*], Yang Zhao[1], Jiqiao Chen[1] and Ting Luo[2]

[1]Lishui Power Supply Company, State Grid Zhejiang Electric Power Co., Ltd., Lishui, China
[2]State Grid Lishui Liandu District Power Supply Company, Lishui, China

Abstract

Power distribution network fault handling imposes stringent requirements on both response timeliness and decision reliability. The high penetration of distributed renewable generation, the widespread deployment of power-electronic devices, and increasing operating-condition fluctuations further intensify the complexity and uncertainty of fault mechanisms. To address the challenges of cross-validating multi-source heterogeneous information, structurally invoking fault-handling knowledge, and balancing latency and accuracy under resource constraints, this paper develops a cross-modal semantic fusion and knowledge-graph reasoning–based fault-handling system, termed CM-SKG. The system takes synchronous phasor measurement unit (PMU) electrical measurements and UAV dual-spectrum inspection images as its primary information sources, and establishes a closed-loop architecture that integrates semantic sensing, transmission, and reasoning. Specifically, task-oriented semantic representations and controllable compression are performed at the edge device; at the edge server, a learnable semantic alignment matrix and a bidirectional interactive attention mechanism are employed for cross-modal fusion, effectively mitigating granularity discrepancies across heterogeneous data and producing consistent fault semantic evidence; at the control center, a fault-handling knowledge graph is introduced for reasoning, mapping fused semantics to executable handling actions, while a tunable reasoning-depth mechanism enables smooth switching between fast response and deeper inference. Furthermore, we propose a semantic communication efficiency (SCE) metric that jointly accounts for cross-modal fusion quality, reasoning reliability, and end-to-end latency, and use it to drive the coordinated optimization of compression ratio, reasoning depth, and computational resources. Online policy learning is realized via a distributional soft actor–critic (DSAC) algorithm. Simulation results demonstrate that CM-SKG significantly improves fault-handling accuracy and decision stability while satisfying real-time constraints.

Keywords: Power distribution network fault handling; semantic communications; cross-modal fusion; knowledge graph reasoning; deep reinforcement learning

Received on 08 January 2026, accepted on 01 March 2026, published on 23 June 2026

doi: 10.4108/ew.13683

1. Introduction

Distribution network fault handling is a critical procedure for ensuring power-supply reliability. With the high penetration of renewable generation and the widespread deployment of power-electronic devices, faults exhibit rapid propagation, strong coupling, and frequent external disturbances, making it difficult for handling strategies based solely on a single type of measurement or a single model to simultaneously meet the requirements of fast response and high-reliability decision-making. Focusing on the fault location, isolation, and service restoration (FLISR) workflow, prior studies have evaluated the reliability of cyber-physical distribution systems from a cyber–physical coupled perspective and incorporated the

*Corresponding author. Email: FANGZHOULIU123456@163.com

FLISR process into reliability impact analysis, thereby providing a basis for engineering constraints and performance characterization of the handling chain [1]. Further algorithmic studies have integrated restoration procedures with operational constraints and reliability-index losses, reflecting an engineering-oriented design philosophy for practical deployment [2]. However, in real-world scenarios, information sources such as field inspections, online measurements, dispatching rules, and handling plans are highly heterogeneous. Without a unified semantic organization and mutual cross-validation mechanism, systems often face the dilemma that information is easy to acquire but hard to exploit, models are trainable but lack interpretability, and handling plans can be generated but are difficult to execute. Meanwhile, studies on the vulnerability of relative data alignment in phasor data concentrators (PDCs) show that time-synchronization attacks can significantly undermine measurement trustworthiness and propagate along the subsequent decision chain [3].

At the sensing and diagnosis level, multimodal information fusion is widely regarded as an effective way to improve fault identification accuracy and coverage. Related work has discussed how multimodal data integration can be applied to smart-grid fault diagnosis [4], and has introduced attention mechanisms into UAV-based power inspection for equipment detection to enhance recognition performance under complex backgrounds [5]. Explorations from the perspective of integrated sensing and communication (ISAC) also provide system-level insights for high-frequency data acquisition and low-latency interaction in distribution networks [6]. Nevertheless, most existing fusion methods adopt late fusion or simple feature concatenation. Because PMU signals are high-frequency but sparse time series, whereas images are low-frequency but dense spatial matrices, direct concatenation ignores the disparity of heterogeneous data manifolds. As a result, high-dimensional visual features may overwhelm low-dimensional electrical transient information, leading to manifold drift [7]. Therefore, for engineering-deployable fault handling, there is a pressing need for a deep fusion mechanism that enables bidirectional interactive cross-validation in a unified latent space and explicitly constrains cross-modal consistency.

In distribution-network edge environments, uplink transmission often relies on bandwidth-limited 4G/5G slices or power-line communication (PLC). Uploading raw high-resolution inspection images and high-rate PMU data can quickly congest the uplink and induce severe queuing delays or even packet loss. To address this issue, semantic communication proposes a paradigm that transmits meaning rather than bits, replacing complete raw data with task-relevant semantic representations. Hoa et al. [8] combine edge computing with semantic communication in UAV-enabled networks and significantly reduce transmission load while improving task efficiency by extracting high-level semantics. Liao et al. [9] propose an "electric semantic short-packet communication" perspective for power systems and emphasize the necessity of lowering load and energy consumption via semantics under millisecond-level measurements. At the theoretical and metric level, Getu et al. [10] systematically review the research landscape of semantic communications. Classical information theory provides a mathematical foundation for information measurement [11], while goal-oriented semantic metrics offer practical tools to connect semantic effectiveness with task utility [12, 13]. Edge AI research for 6G further that future mission-critical services will pursue both ultra-low latency and ubiquitous intelligence, and that edge-side cooperative computing and self-learning under resource constraints will be essential capabilities [14]. From the evolution of massive IoT, Guo et al. [15] provide a comprehensive survey of IoT toward 6G, and Khan et al. [16] discuss the 6G vision and architectural elements. Xiao et al. [17] focus on self-learning edge intelligence in 6G, whereas Zhou et al. [18] highlight the importance of pushing AI capabilities down to the edge.

When semantic information comes from heterogeneous modalities, cross-modal alignment and collaboration become crucial. Zhou et al. [19] propose a cross-modal collaborative communication framework to support information complementarity across modalities, and Chen et al. [20] further discuss cross-modal semantic communications for 6G. In addition, research on knowledge-augmented neural reasoning for multi-hop reasoning tasks shows that introducing external structured knowledge can improve reasoning reliability and generalization [21], providing inspiration for an explainable closed loop that connects evidence fusion, knowledge constraints, and action decisions in distribution-network handling. Knowledge graphs (KGs) therefore provide a suitable carrier for organizing topology, equipment, operating rules, and mechanism constraints. Chaccour et al. [22] advocate a "less data, more knowledge" philosophy for building next-generation semantic communication networks. Zhang et al. [23] propose a knowledge-graph-based explainable and generalized zero-shot semantic communication framework, and Song et al. [24] demonstrate robustness gains from knowledge-graph-enabled semantic communication in UAV object detection. In the power domain, Li et al. [25] construct a power-equipment knowledge graph for intelligent fault detection; Sun et al. [26] propose a dispatch knowledge graph for distributed energy management; and Pi et al. [27] present the construction and application of a knowledge graph for intelligent decision-making in power-grid fault handling. Meanwhile, graph neural network–based knowledge reasoning provides general support for complex relational modeling and inference [28].

Nevertheless, distribution-network fault handling is a typical online closed-loop decision process: the semantic compression strength, computing-resource allocation, and reasoning depth must be adjusted dynamically with operating conditions and link states, and decision policies should be learned online to maximize long-term cumulative rewards. At the algorithmic level, Soft Actor–Critic (SAC), proposed by Haarnoja et al. [29], is a representative maximum-entropy reinforcement learning method; it encourages exploration and improves training stability by adding an entropy regularizer to the expected return, while retaining off-policy learning and high sample efficiency. Its temperature parameter can adaptively balance exploration and exploitation, making it well suited for continuous control under dynamic uncertainty.

Furthermore, Duan et al. [30] introduce distributional value modeling in Distributional SAC (DSAC), which models the return distribution rather than only the expected Q-value, thereby mitigating value-estimation errors and overestimation bias and improving robustness and convergence stability under stochastic channels and fluctuating compute budgets. Duan et al. [31] also provide a more systematic discussion of the characteristics and application boundaries of such methods.

To address these challenges, this paper proposes a cross-modal semantic fusion and knowledge-graph reasoning–based fault-handling system (CM-SKG). The system establishes a closed-loop architecture that integrates semantic sensing, transmission, and reasoning: task-oriented semantic generation and controllable compression are performed at the device side; cross-modal interactive cross-validation is realized at the edge via a learnable matrix and attention mechanisms; and at the control center, knowledge-graph reasoning maps fused semantics to handling actions, while online coordination is driven by a unified effectiveness metric.

2. System Architecture

This paper addresses the dual requirements of low-latency response and highly reliable decision-making for distribution-network fault handling in scenarios with high penetration of renewable energy. As shown in Figure 1, we construct a hierarchical collaborative framework that combines CM-SKG. The framework follows an integrated paradigm of semantic sensing, transmission, and reasoning: task-oriented semantic representations are used as the information carrier, multi-modal raw observations are mapped into compact semantic embeddings at the device side and transmitted under a budget-controlled scheme, cross-modal semantic alignment and fusion are performed at the edge, and controllable-depth reasoning is conducted at the master station based on a distribution-network knowledge graph. In this way, the credibility and executability of fault-handling schemes are enhanced under end-to-end latency constraints.

The set of acquired modalities is defined as $M = \{\text{pmu},\text{img}\}$, corresponding to electrical measurements from phasor measurement units (PMUs) and visible/infrared images collected by unmanned aerial vehicles (UAVs), respectively. To characterize online regulation under dynamic operating conditions, we adopt a discrete decision-cycle model by partitioning continuous time into T equal-length decision slots, each with duration Δt. The slot index $t \in \{1,\ldots,T\}$ corresponds to the time interval $[(t-1)\Delta t, t\Delta t)$ Within each slot, the system completes a closed-loop update consisting of device-side semantic feature extraction and modality budgeted compression, semantic uploading and edge-side cross-modal fusion, and master-station controllable-depth knowledge reasoning (κ_t), thereby enabling temporally coordinated decision-making for distribution-network fault handling.

The overall architecture comprises four layers: the device layer, sensing layer, edge layer, and master-station layer. The device layer consists of key electrical equipment and switching components in the distribution network—including distributed photovoltaic inverters, energy-storage converters, distribution transformers, feeders, and switches—which serve as the physical objects where faults occur and control actions are executed. The sensing layer includes PMUs and inspection UAVs, which synchronously collect nodal electrical quantities and visual/thermal image information, respectively, and perform task-oriented semantic feature extraction and modality budgeted compression at the device side to generate lightweight semantic packets for uplink transmission. The edge layer is formed by on-site edge servers, which receive the multi-modal semantic packets and conduct semantic alignment and cross-modal attention-based fusion to produce an integrated fault semantic representation, and then forward it to the master station to support global reasoning and action planning. The master-station layer integrates a distribution-network knowledge graph and a reasoning engine; it performs controllable-depth graph reasoning conditioned on the edge-fused semantics to output fault-handling schemes such as fault isolation, load transfer, and maintenance strategies, and sends the resulting control commands back to the field for execution.

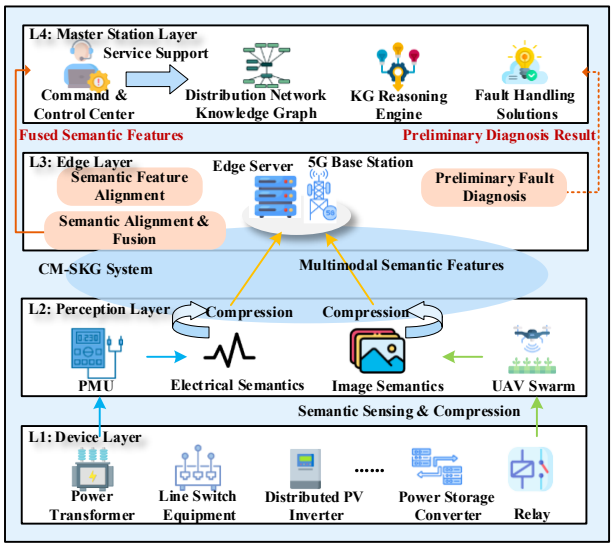

Figure 1. CM-SKG hierarchical collaborative fault-handling framework based on cross-modal semantic fusion and knowledge-graph reasoning

2.1. Semantic Feature Extraction and Budgeted Compression

Figure 2 provides an overview of the proposed CM-SKG processing pipeline. Specifically, raw PMU measurements and UAV images are first converted into task-oriented semantic representations and compressed under a controllable budget at the device side; the resulting semantic packets are then aligned and fused at the edge; finally, the

fused semantics are fed into knowledge-graph reasoning to produce fault-handling decisions at the control center. In this subsection, we focus on the device-side semantic extraction and budget-controlled compression modules.

PMUs serve as the core devices for situational awareness in distribution networks. Within slot t, a PMU synchronously samples electrical quantities at a sampling frequency f_s. The raw waveform matrix collected in this slot is defined as $X_{\mathrm{pmu},t} \in \mathbb{R}^{9\times N_s}$:

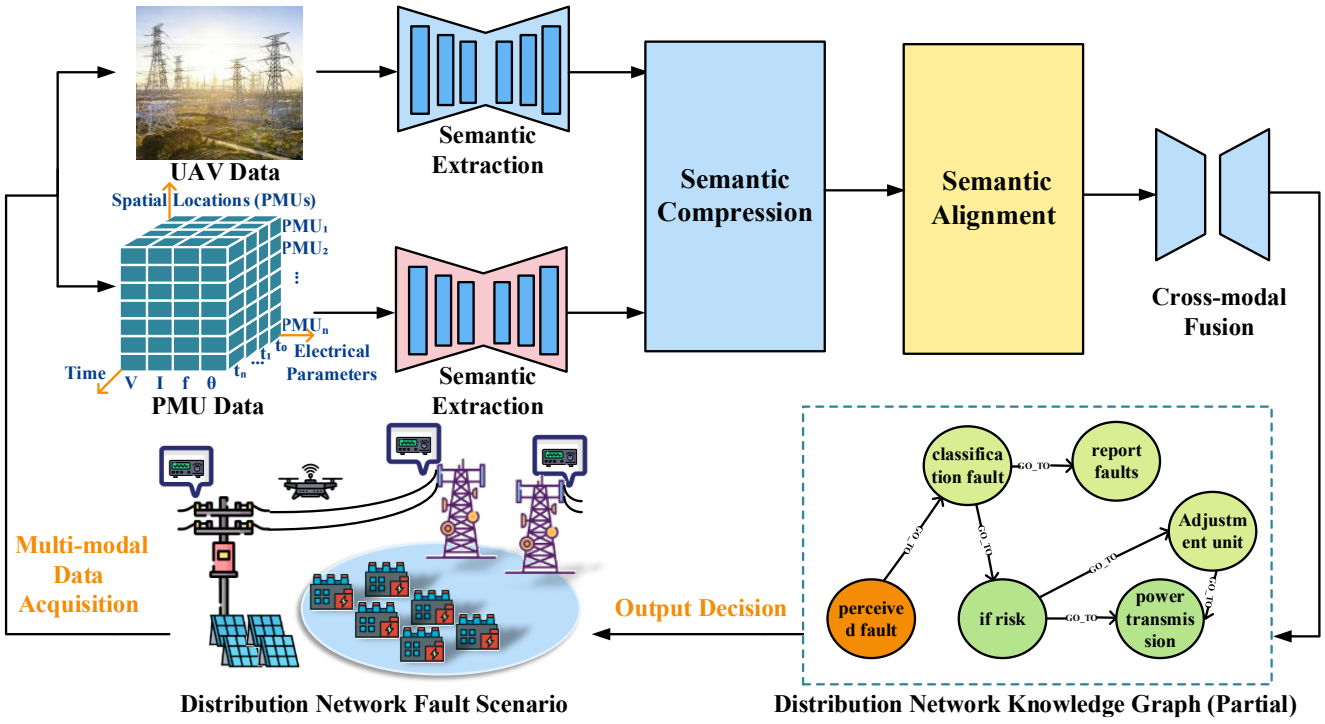

Figure 2. End-to-end CM-SKG workflow

$$X_{\mathrm{pmu},t} = \left[V_{abc}(\tau), I_{abc}(\tau), \theta_{abc}(\tau)\right]^{\top} \tag{1}$$

Where the 9 channels cover three-phase voltages, three-phase currents, and three-phase phase angles, $\tau \in \{1,2,\ldots,N_s\}$, and $N_s = f_s \Delta t$, $\cdot \Delta t$ denotes the number of time-domain samples within one decision period. To realize the "understand-then-transmit" paradigm of semantic communications, the system constructs a deep semantic encoder $f_e^{\mathrm{pmu}}(\cdot)$ using a 1D-CNN and a self-attention mechanism, mapping massive samples to a compact high-dimensional physical semantic vector $Z_{\mathrm{pmu},t} \in \mathbb{R}^{d_{pmu}}$:

$$Z_{\mathrm{pmu},t} = f_e^{\mathrm{pmu}}\left(X_{\mathrm{pmu},t}\right) \tag{2}$$

For UAV-based visual sensing, the inspection system synchronously acquires operating-state observations of distribution equipment using visible-light and infrared dual-spectrum sensors. Let the raw image data collected in slot t be $I_{\mathrm{img},t}$. A deep convolutional neural network $f_e^{\mathrm{img}}(\cdot)$ is used to map the high-dimensional pixel space into a compact visual semantic feature $Z_{\mathrm{img},t} \in \mathbb{R}^{d_{img}}$:

$$Z_{\mathrm{img},t} = f_e^{\mathrm{img}}\left(I_{\mathrm{img},t}\right) \tag{3}$$

The visual semantic encoding process aims to decouple task-relevant semantic anchors that are strongly correlated with fault diagnosis from complex backgrounds, thereby suppressing non-critical information.

To reduce communication load while preserving the usability of fault-discriminative semantics, we introduce a task-oriented semantic compression mechanism after device-side multimodal semantic encoding. Let the system-level semantic compression ratio be $x_t \in [x_{\min}, 1]$, which is treated as a continuous and adjustable variable in modeling and optimization. Considering that the two modalities $m \in \{\mathrm{pmu}, \mathrm{img}\}$ differ in semantic representation dimensionality and encoding overhead, we define the full-semantic bit size $A_m(\mathrm{bit})$ for each modality and construct the modality-wise effective semantic bit budget as

$$B_{m,t}^{(\mathrm{sem})} = x_t A_m, m \in \{\mathrm{pmu}, \mathrm{img}\} \tag{4}$$

We evaluate the overall semantic importance of the current task at the sample/slot granularity and adaptively determine the shared compression ratio x_t based on this importance. The device side first performs global

aggregation over the semantic representations of both modalities to obtain summary vectors $u_{m,t}=\mathrm{Pool}\left(Z_{m,t}\right)$, and then assesses the importance of the cross-modal joint semantics in the fused space. The overall semantic-importance indicator $I_t\in[0,1]$ is defined as

$$I_t=\sigma\left(a^{\cdot}\ \tanh\left(W[u_{pmu,t}\|u_{img,t}]+b\right)\right),I_t\in[0,1] \quad (5)$$

Here, $[\cdot\|\cdot]$ denotes vector concatenation, W,b,a are parameters of a lightweight importance estimator, and $\sigma(\cdot)$ is the sigmoid function. This definition scores the task relevance directly on the joint semantics of the two modalities, and is therefore consistent with the "shared compression ratio" design: a larger I_t indicates that fault discrimination in the current slot relies more on high-fidelity cross-modal semantic information, whereas a smaller I_t suggests higher semantic redundancy or lower discriminative contribution.

Accordingly, we employ a monotonic and differentiable continuous mapping to convert the overall semantic importance into the compression level for slot t:

$$x_t=\phi\left(I_t\right)=x_{\min}+\left(1-x_{\min}\right)\cdot\sigma\left(\kappa\left(I_t-\tau\right)\right) \quad (6)$$

Here, $\kappa>0$ is the slope coefficient and $\tau\in(0,1)$ is the importance threshold, ensuring $x_t\in[x_{\min},1]$ and monotonic increase with I_t. This importance-based compression strategy aligns with the priority characteristics of distribution network fault handling. The indicator I_t effectively acts as a priority filter: data segments containing critical fault signatures (high relevance) are assigned a larger transmission budget to preserve diagnostic details, whereas steady-state monitoring data (low relevance) are heavily compressed to conserve bandwidth resources. Given x_t, the compressor performs task-oriented compression coding on the device-side semantic representations, denoted by the compression operator $Z_{m,t}^{(\mathrm{comp})}=C_m\left(Z_{m,t};x_t\right)$. The resulting $Z_{m,t}^{(\mathrm{comp})}$ is a low-bit-rate semantic code controlled by the shared compression ratio, and its uplink bit size is uniformly characterized by $B_{m,t}^{(\mathrm{sem})}$. Finally, the device uploads $Z_{\mathrm{pmu},t}^{(\mathrm{comp})}$ and $Z_{\mathrm{img},t}^{(\mathrm{comp})}$ to the edge server, providing compact yet more discriminative semantic inputs for subsequent cross-modal fusion and knowledge reasoning.

2.2. Edge Cross-Modal Fusion

To deeply exploit the complementary information between electrical semantics and visual semantics, we design the bidirectional cross-attention mechanism illustrated in Figure 3 Specifically, let the compressed PMU semantic feature and the compressed UAV image semantic feature received at the edge in slot t be $Z_{pmu,t}^{(comp)}\in\mathrm{R}^{d_{pmu}}$ and $Z_{img,t}^{(comp)}\in\mathrm{R}^{d_{img}}$, respectively. Since the two modalities differ significantly in dimensionality and distribution characteristics, we first introduce modality-specific alignment mappings to project them into a shared semantic space of dimension d_{aln}, thereby establishing a consistent semantic coordinate system:

$$\bar{Z}_{m,t}=\phi\left(W_mZ_{m,t}^{(comp)}+b_m\right),m\in\{pmu,img\} \quad (7)$$

Here, $W_{pmu}\in\mathrm{R}^{d_{aln}\times d_{pmu}}$ and $W_{img}\in\mathrm{R}^{d_{aln}\times d_{img}}$ are learnable alignment matrices, and $\phi(\cdot)$ is a nonlinear activation function. This alignment step makes cross-modal features comparable and interactive in the unified semantic space, laying the foundation for subsequent cross-attention computation.

Building on semantic alignment, we adopt bidirectional multi-head cross-attention to enable deep semantic interaction and mutual verification between modalities. Along the PMU →image path, $\bar{Z}_{pmu,t}$ serves as the query, while $\bar{Z}_{img,t}$ serves as the key and value, retrieving visual evidence correlated with electrical abnormal patterns and yielding an image-enhanced PMU representation:

$$Z_{pmu\to img,t}=MHA\left(\bar{Z}_{pmu,t},\bar{Z}_{img,t},\bar{Z}_{img,t}\right) \quad (8)$$

Symmetrically, along the image →PMU path, $\bar{Z}_{img,t}$ acts as the query and matches electrical patterns consistent with visual anomalies, producing an electrically enhanced image representation:

$$Z_{img\to pmu,t}=MHA\left(\bar{Z}_{img,t},\bar{Z}_{pmu,t},\bar{Z}_{pmu,t}\right) \quad (9)$$

Here, $\mathrm{MHA}(\cdot)$ denotes the multi-head cross-attention module with H parallel heads. The outputs of all heads are concatenated and projected back to the shared semantic dimension:

$$MHA(Q,K,V)=[head_1\oplus head_2\oplus\cdots\oplus head_H]W_O \quad (10)$$

In the above, $\oplus$ denotes concatenation and head_h is the attention output of the h-th head given (Q,K,V), $W_O\in\mathrm{R}^{d_{aln}\times d_{aln}}$. This multi-head structure captures cross-modal correlations and complementary evidence in multiple subspaces in parallel, improving alignment capability and fusion robustness under complex fault conditions. Specifically, this bidirectional mechanism implicitly addresses the granularity discrepancy between heterogeneous data. By using sparse PMU temporal samples as queries to retrieve dense visual spatial features (and vice versa), the model aligns electrical transients with corresponding visual regions in the latent space. This overcomes the structural misalignment caused by their differing physical definitions without requiring strict temporal-spatial synchronization.

To adaptively adjust the contributions of the two interaction paths according to fault types and evidence salience, we introduce dynamic weights $\alpha_{pmu,t}$ and $\alpha_{img,t}$. We first globally pool the two cross-attention outputs to obtain summary vectors, then feed them into a lightweight

gating network to produce unnormalized scores $\tilde{\alpha}_{pmu,t}$ and $\tilde{\alpha}_{img,t}$, which are normalized by Softmax:

$$\alpha_{pmu,t} = \frac{\exp(\tilde{\alpha}_{pmu,t})}{\exp(\tilde{\alpha}_{pmu,t}) + \exp(\tilde{\alpha}_{img,t})} \tag{11}$$

$$\alpha_{img,t} = \frac{\exp(\tilde{\alpha}_{img,t})}{\exp(\tilde{\alpha}_{pmu,t}) + \exp(\tilde{\alpha}_{img,t})} \tag{12}$$

The final fused semantic representation is obtained by weighted integration:

$$Z_t^{(fus)} = \alpha_{pmu,t} Z_{pmu\to img,t} + \alpha_{img,t} Z_{img\to pmu,t} \tag{13}$$

Furthermore, the dynamic weights α_{pmu} and α_{img} endow the system with robustness against environmental interference. For instance, if extreme weather (e.g., heavy fog) degrades visual clarity, the gating network detects the reduced confidence in visual semantics. Consequently, it automatically down-weights α_{img} and up-weights α_{pmu}, shifting reliance to reliable electrical measurements. This salience-based adaptive weighting ensures system resilience even when one modality is compromised.

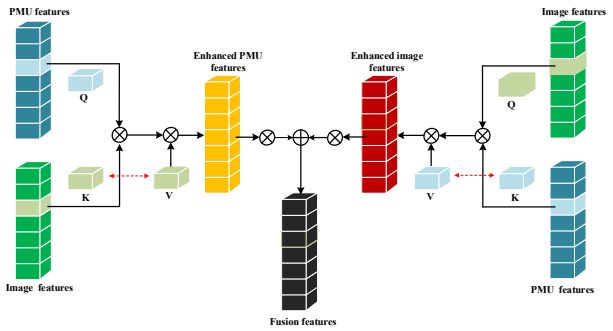

Figure 3. Cross-modal semantic fusion via bidirectional multi-head cross-attention

To quantify cross-modal semantic fusion performance, we define a fusion-quality metric $Q_t^{(fus)}$ to measure the semantic deviation between the fused representation and a reference representation. Let $Z_t^{(fus)}$ denote the reference fused semantics under ideal conditions. The fusion quality in slot t is defined as

$$Q_t^{(fus)} = \left(1 + \sigma_f \left\| Z_t^{(fus)} - Z_t^{(ref)} \right\|_2\right)^{-1} \tag{14}$$

Where σ_f controls the sensitivity to semantic deviation and $\|\cdot\|_2$ is the Euclidean norm. The metric lies in $(0,1]$: higher consistency with the reference yields a smaller distance and thus a larger $Q_t^{(fus)}$.

2.3. Controllable KG Reasoning

The fused semantic representation $Z_t^{(\mathrm{fus})}$ output by the edge layer aggregates electrical anomaly evidence and visual anomaly evidence. we construct a fault-handling knowledge graph and introduce a reasoning-depth control factor κ_t, enabling a graded switch from fast matching to deep inference and thereby forming an adaptive reasoning mechanism. The fault-handling knowledge graph is represented as:

$$KG = (E, R, T), \quad T \subseteq E \times R \times E \tag{15}$$

Where E is the entity set, including key entities such as distribution-network equipment, fault types, operating states, and handling actions; R is the relation set, characterizing topological connections, causal relations, and handling-constraint relations among entities; and T is the triple set, organizing network-topology facts, fault-mechanism rules, and empirical handling knowledge to provide explainable support for the reasoning process. To improve the discriminability and scale stability of the query vector, we adopt a composite mechanism comprising affine mapping, nonlinear activation, and normalization:

$$q_t = \mathrm{Norm}\left(\phi\left(W_q Z_t^{(fus)} + b_q\right)\right) \tag{16}$$

Where W_q and b_q are trainable parameters, $\phi(\cdot)$ is a nonlinear activation function, and $\mathrm{Norm}(\cdot)$ denotes ℓ_2 normalization, which suppresses the impact of feature-magnitude fluctuations across operating conditions on matching scores. Let A_{KG} denote the candidate action set. For any candidate action $a \in A_{KG}$, the reasoning module outputs a matching/feasibility score between the query q_t and action a under knowledge-graph constraints, supporting subsequent action selection.

To capture the differences in computational cost and reasoning capability under different reasoning depths, we construct a lightweight reasoning branch and a deep reasoning branch, which output candidate-action scores $S^{(L)}(a|\ q_t)$ and $S^{(D)}(a|\ q_t)$, respectively. The lightweight branch focuses on fast matching based on local neighborhoods and limited-step relation expansion to reduce reasoning latency, whereas the deep branch emphasizes multi-hop path inference and compositional relational reasoning in the embedding space to enhance discriminability under complex fault scenarios. The two branches are gated and fused via the reasoning-depth control factor $\kappa_t \in [0,1]$. To improve interpretability and numerical stability, we write the fusion as a temperature-smoothed interpolation:

$$\begin{aligned} S(a|\ q_t) = (1-\kappa_t) S^{(L)}(a|\ q_t) \\ + \kappa_t S^{(D)}(a|\ q_t) - \mu \kappa_t (1-\kappa_t) \end{aligned} \tag{17}$$

Where $\mu \geq 0$ is a regularization coefficient used to discourage frequent extreme switching of κ_t during training. Physically, this tunable depth mechanism adapts to the response requirements of distinct fault scenarios. For evident faults like short circuits, the agent tends to lower κ_t to leverage the lightweight branch for rapid isolation (ms-level). Conversely, for ambiguous scenarios such as high-impedance grounding, the system increases κ_t to

activate deep multi-hop reasoning, ensuring decision reliability at the cost of marginally increased latency.

Furthermore, we normalize the candidate-action scores into a handling decision distribution. Considering score-scale variations across scenarios and the influence of rule constraints on action executability, we introduce a temperature coefficient $\tau_s > 0$ and a soft feasibility weight $\omega(a) \in (0,1]$, yielding:

$$p(a|\ q_t) = \frac{\exp\left(\tau_s S(a|\ q_t)\right)\omega(a)}{\sum_{a' \in A_{KG}} \exp\left(\tau_s S(a'|\ q_t)\right)\omega(a')} \tag{18}$$

2.4. Latency Modeling

In the proposed CM-SKG framework, end-to-end latency is a key indicator for evaluating the system's fast-response capability. We build a full-process latency model that covers multimodal data acquisition, device-side semantic extraction and compression, wireless transmission, edge-side cross-modal fusion, and master-station knowledge reasoning. Considering that the PMU chain and the visual chain are executed in parallel at the device side, while the fusion and knowledge reasoning stages are performed sequentially at the edge and the control center, respectively, the total end-to-end latency in slot t is expressed as

$$\tau_t^{(total)} = \max\left(\tau_{pmu,t}^{(chain)}, \tau_{img,t}^{(chain)}\right) + \tau_t^{(fus)} + \tau_t^{(KG)} \tag{19}$$

Here, $\tau_{m,t}^{(\text{chain})}$ denotes the device-side processing-chain latency of modality $m \in \{pmu, img\}$, which consists of semantic extraction, semantic compression, and data transmission in series:

$$\tau_{m,t}^{(chain)} = \tau_{m,t}^{(ext)} + \tau_{m,t}^{(comp)} + \tau_{m,t}^{(trans)} \tag{20}$$

The semantic-extraction latency depends on the feature dimension, the per-dimension computational complexity, and the allocated device-side compute resource, and is modeled as

$$\tau_{m,t}^{(ext)} = \frac{d_m^{(ft)} C_m^{(ext)}}{f_{m,t}^{(CPU)}} \tag{21}$$

Where $d_m^{(ft)}$ is the semantic feature dimension of modality m, $C_m^{(ext)}$ is the computation required for extracting one feature dimension, and $f_{m,t}^{(CPU)}$ is the CPU frequency allocated to modality m.

We use a system-level semantic compression ratio $x_t \in [x_{min}, 1]$ to control the uplink semantic budget. The compression process includes importance evaluation and bit allocation/quantization, whose computational overhead typically increases as compression becomes stronger. To capture this monotonic relation, the semantic-compression latency is modeled as

$$\tau_{m,t}^{(comp)} = \frac{d_m^{(ft)} C_m^{(comp)} \psi(x_t)}{f_{m,t}^{(CPU)}} \tag{22}$$

Where $C_m^{(comp)}$ is the baseline computation required for compressing one feature dimension, and $\psi(x_t) = \left(\frac{1}{x_t}\right)^{\alpha}$ with $\alpha > 0$ characterizes the additional complexity growth induced by the compression strength as x_t varies.

The transmission stage is mainly determined by the uploaded semantic payload and the available uplink rate. Since our focus is on the decision loop of semantic fusion and knowledge reasoning, and to avoid introducing overly detailed physical-layer modeling, we approximate link throughput using an average achievable rate. The transmission latency of modality $m$is defined as

$$\tau_{m,t}^{(trans)} = \frac{L_{m,t}}{R_{m,t}} \tag{23}$$

The uploaded semantic payload is determined by the compression ratio and satisfies $L_{m,t} = x_t A_m$, $m \in \{pmu, img\}$, where A_m is the full-semantic bit size of modality m. The transmission rate is approximated by a capacity-form expression:

$$R_{m,t} = B_m \log_2\left(1 + \gamma_{m,t}\right) \tag{24}$$

Where B_m is the uplink bandwidth allocated to modality m, and $\gamma_{m,t}$ is the effective signal-to-noise ratio (SNR). This simplified model captures the principal coupling among compression ratio, uplink payload, and transmission latency while preventing excessive communication-model complexity.

Cross-modal fusion is executed on the edge server, whose computational cost depends on the fusion network scale and allocated compute resources. For simplicity, let the equivalent computation of the fusion stage be C_{fus}, and let the allocated CPU frequency be $f_{fus,t}^{(CPU)}$. The fusion latency is then:

$$\tau_t^{(fus)} = \frac{C_{fus}}{f_{fus,t}^{(CPU)}} \tag{25}$$

At the control center, knowledge-graph reasoning uses the reasoning-depth control factor $\kappa_t \in [0,1]$ to smoothly interpolate between lightweight reasoning and deep reasoning. The KG reasoning latency is modeled as

$$\tau_t^{(KG)} = \tau_{min} + (\tau_{max} - \tau_{min})\kappa_t \tag{26}$$

where τ_{min} and τ_{max} correspond to the baseline latencies of lightweight reasoning and deep reasoning, respectively.

3. System Architecture

To comprehensively evaluate the decision reliability and response timeliness of the CM-SKG system under resource constraints, we develop a unified optimization framework that jointly considers fused semantic quality and end-to-end latency. Improving semantic quality typically incurs higher computational and communication overhead, which

in turn increases processing and transmission delays; conversely, pursuing ultra-low latency alone may sacrifice semantic fidelity and the reliability of fault-handling decisions. Therefore, it is necessary to dynamically balance semantic quality and response latency under stringent real-time constraints.

3.1. Definition of Semantic Communication Efficacy

The overall semantic quality of the system is jointly affected by cross-modal fusion performance and the reliability of knowledge-graph reasoning. To capture their complementarity and synergy along the decision chain, we define the overall semantic quality in slot t as

$$Q_t = \omega_f Q_t^{(fus)} + \omega_k Q_t^{(kg)} + \lambda Q_t^{(fus)} Q_t^{(kg)} \tag{27}$$

Where $Q_t^{(fus)}$ is the cross-modal fusion-quality score and $Q_t^{(kg)}$ is the KG reasoning-quality score; $\omega_f \geq 0$ and $\omega_k \geq 0$ are weighting coefficients satisfying $\omega_f + \omega_k = 1$, used to adjust the relative importance of fusion quality and reasoning quality in the overall evaluation; and $\lambda \geq 0$ is an interaction coefficient that characterizes potential nonlinear synergy between fusion quality and reasoning quality. If the synergy term is ignored, one may set $\lambda \geq 0$, which reduces the definition to a linear weighted form.

In distribution-network fault handling, the system must satisfy both a minimum semantic-quality requirement and a maximum end-to-end latency constraint. To uniformly quantify the overall benefit brought by semantic-quality gains and latency improvements, we propose Semantic Communication Efficacy (SCE) as an instantaneous performance metric for slot t:

$$\varpi_t = \xi_q (Q_t - Q_{min}) + \xi_d (\tau_{max} - \tau_t^{(total)}) \tag{28}$$

Where $\xi_q \in (0,1)$ and $\xi_d \in (0,1)$ are the weighting coefficients for the semantic-quality term and the latency term, respectively, satisfying $\xi_q + \xi_d = 1$, and can be configured according to operating conditions; Q_{min} is the minimum semantic-quality threshold, and τ_{max} is the maximum tolerable end-to-end latency. This metric provides a unified quantification of semantic fidelity and system responsiveness: a larger ϖ_t indicates better overall performance under the given constraints.

3.2. Joint Optimization Problem Formulation

We model the system's adaptive regulation under dynamic operating conditions as a Markov decision process (MDP) and employ deep reinforcement learning to learn a policy π that maximizes the expected long-term cumulative discounted efficacy. Under resource and quality-of-service constraints, the joint optimization problem is formulated as

$$\text{(P1):}\ \max_{\pi} E_{\pi}\left[\sum_{t=0}^{\infty} \gamma^t \varpi_t\right] \tag{29}$$

$$\text{s.t. } x_{min} \leq x_t \leq 1, \forall t \tag{29a}$$

$$0 \leq \kappa_t \leq 1, \forall t \tag{29b}$$

$$f_{m,min}^{(CPU)} \leq f_{m,t}^{(CPU)} \leq f_{m,avail}^{(CPU)}, m \in \{pmu, img\}, \forall t \tag{29c}$$

$$f_{fus,min}^{(CPU)} \leq f_{fus,t}^{(CPU)} \leq f_{fus,avail}^{(CPU)}, \forall t \tag{29d}$$

$$Q_t \geq Q_{min}, \forall t \tag{29e}$$

$$\tau_t^{(total)} \leq \tau_{max}, \forall t \tag{29f}$$

Here, $\gamma \in (0,1)$ is the discount factor. In each slot t, the system adjusts the semantic compression ratio x_t and the reasoning-depth control factor κ_t, and allocates device-side and edge-side computing resources to jointly optimize overall semantic quality and end-to-end latency. Specifically, $x_t \in [x_{min}, 1]$ characterizes the uplink semantic budget strength, while $\kappa_t \in [0,1]$ enables a smooth transition between lightweight reasoning and deep reasoning, thereby dynamically balancing reasoning accuracy and computational cost.

3.3. DSAC Cooperative Optimization

To ensure the observability and trainability of the state, this paper selects low-dimensional statistics that reflect link quality, computing resources, and historical performance feedback to form the state vector:

$$\begin{aligned} s_t = [&\gamma_{\text{pmu},t}, \gamma_{\text{img},t}, f_{\text{pmu},t}^{(\text{CPU})}, f_{\text{img},t}^{(\text{CPU})}, \\ &f_{\text{fus},t}^{(\text{CPU})}, Q_{t-1}^{(\text{fus})}, Q_{t-1}^{(\text{kg})}, \tau_{t-1}^{(\text{total})}] \end{aligned} \tag{30}$$

At each time slot t, the system outputs a joint decision over the key controllable parameters:

$$a_t = \left[x_t, \kappa_t, f_{pmu,t}^{(CPU)}, f_{img,t}^{(CPU)}, f_{fus,t}^{(CPU)}\right] \tag{31}$$

Here, $x_t \in [x_{min}, 1]$ is the semantic compression ratio, and $\kappa_t \in [0,1]$ is the reasoning-depth control factor for knowledge-graph reasoning. The fusion weights $\alpha_{pmu,t}$ and $\alpha_{img,t}$ are adaptively generated by the internal Softmax gating in the fusion module, and thus are not explicitly included as external control actions.

The instantaneous reward uses SCE as the main term and introduces a constraint-violation penalty:

$$r_t = \varpi_t - \rho \phi_t \tag{32}$$

Where $\rho > 0$ is the penalty coefficient. ϕ_t is a constraint-violation indicator: if the overall semantic quality in slot t is lower than the threshold Q_{min} or the end-to-end total latency exceeds the upper bound $\tau_{\max}$, then $\phi_t = 1$; otherwise, $\phi_t = 0$.

Under the maximum-entropy reinforcement learning framework, we define the soft return random variable as

$$Z(s_t, a_t) = \sum_{i=t}^{\infty} \gamma^{i-t} \left(r_i - \nu \log \pi_{\phi}(a_i | s_i)\right) \tag{33}$$

Where $\gamma \in (0,1)$ is the discount factor, $\nu > 0$ is the entropy regularization coefficient, and $\pi_{\phi}(a | s)$ is the actor (policy)

network parameterized by ϕ . The corresponding distributional Bellman equation is

$$Z(s_t,a_t) \overset{D}{=} r_t + \gamma\left(Z(s_{t+1},a_{t+1}) - \nu \log \pi_\phi(a_{t+1}|\ s_{t+1})\right) \quad (34)$$

Where $a_{t+1} \sim \pi_\phi(\cdot|\ s_{t+1})$ denotes that the action at the next moment a_t follows the conditional distribution output by the policy network π_ϕ .

For implementation, we approximate the return distribution with a Gaussian:

$$Z_\theta(s,a) \sim \mathrm{N}\left(\mu_\theta(s,a), \sigma_\theta^2(s,a)\right) \quad (35)$$

Where $\mu_\theta(\cdot)$ and $\sigma_\theta(\cdot)$ are output by the critic network, and θ denotes the critic parameters. We sample a mini-batch of transitions $(s_t,a_t,r_t,s_{t+1}) \sim \mathrm{B}$ from the replay buffer B . Using the target critic network parameters $\bar{\theta}$, we construct the target sample:

$$y_t = r_t + \gamma\left(\mu_{\bar{\theta}}(s_{t+1},a_{t+1}) - \nu \log \pi_\phi(a_{t+1}|\ s_{t+1})\right) \quad (36)$$

Then, the critic parameters are updated by minimizing the negative log-likelihood (NLL) of the target under the predicted Gaussian distribution:

$$\mathrm{L}_{critic}(\theta) = \mathrm{E}_{(s_t,a_t,r_t,s_{t+1})\sim \mathrm{B}}\left[\frac{\left(y_t - \mu_\theta(s_t,a_t)\right)^2}{2\sigma_\theta^2(s_t,a_t)} + \frac{1}{2}\log \sigma_\theta^2(s_t,a_t)\right] \quad (37)$$

The target network adopts a soft-update rule:

$$\bar{\theta} \leftarrow \tau\theta + (1-\tau)\bar{\theta} \quad (38)$$

Using the critic mean $\mu_\theta(s,a)$ as the value approximation, the actor parameters ϕ are updated by minimizing:

$$\mathrm{L}_{actor}(\phi) = \mathrm{E}_{s_t\sim \mathrm{B},\, a_t\sim\pi_\phi(\cdot|\ s_t)}[\nu \log \pi_\phi(a_t|\ s_t) - \mu_\theta(s_t,a_t)] \quad (39)$$

4. Simulation Results and Analysis

To comprehensively validate the proposed CM-SKG system, we built a simulation environment based on a modified IEEE 33-bus distribution network integrated with high-penetration distributed photovoltaics (PV) and UAV inspection paths.

We constructed a synchronous multi-modal dataset. Electrical data includes voltage/current phasors sampled at 4 kHz with Gaussian white noise (SNR: 20-40dB). Visual data consists of paired visible-light and infrared images resized to 224×224 , covering fault types such as short-circuits, high-impedance grounding, insulator breakage, and overheating.

The edge-cloud link is modeled as a Rayleigh fading channel with bandwidth dynamically fluctuating between 5 MHz and 50 MHz. The device-side CPU frequency is 1.2 GHz, and the edge server is 3.0 GHz.

We compare CM-SKG with: (1) Algorithms: Soft Actor-Critic (SAC), Multi-Agent PPO (MAPPO); (2) Modalities: PMU-Only, Image-Only, Traditional Bit-level Transmission; (3) Strategies: Fixed Reasoning (High Quality), Fixed Compression (Low Latency).

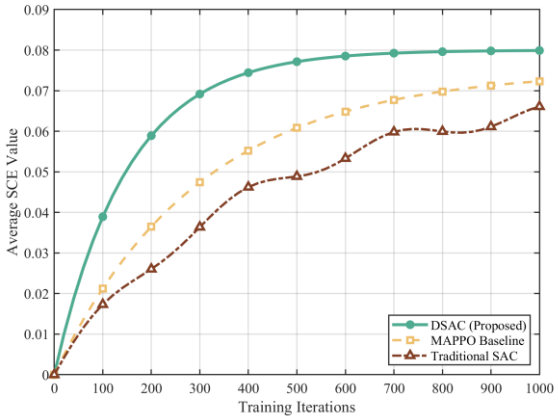

Figure 4. Convergence performance comparison of different reinforcement learning algorithms

Figure 4 illustrates the training trajectory of the proposed DSAC-based strategy compared to baseline RL algorithms. It is observed that DSAC achieves the highest average Semantic Communication Efficacy (SCE) and exhibits significantly tighter variance (higher stability) than SAC and MAPPO. The superiority of DSAC stems from its distributional value modeling capability. In distribution networks, channel fading and fault occurrences introduce high stochasticity. Traditional RL methods (like SAC) estimate the expectation of Q-values, which often leads to overestimation bias and risky actions in volatile environments. In contrast, by learning the full distribution of returns, DSAC allows the agent to capture the intrinsic uncertainty and adopt a risk-aware policy, ensuring stable convergence even under highly dynamic channel conditions.

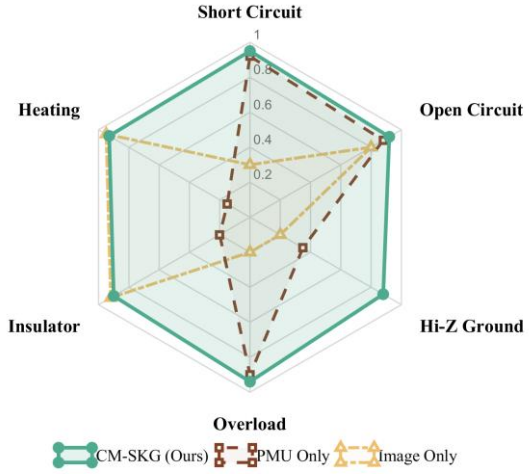

Figure 5. Fault diagnosis accuracy comparison across different fault types using radar chart

Figure 5 evaluates the fault diagnosis capability across heterogeneous fault scenarios. The single-modal approaches exhibit distinct perceptual blind spots: "PMU Only" excels at electrical transient faults but fails to detect physical defects like Insulator breakage; conversely, "Image Only" is sensitive to visual anomalies like Heating but struggles with internal state issues like Overload. CM-SKG forms a comprehensive hexagonal coverage map. This is attributed to the bidirectional cross-attention mechanism, which aligns heterogeneous feature manifolds in the latent space. Visual semantics provide context to disambiguate electrical signals, while electrical features corroborate visual anomalies, yielding a synergistic effect that surpasses the sum of individual modalities.

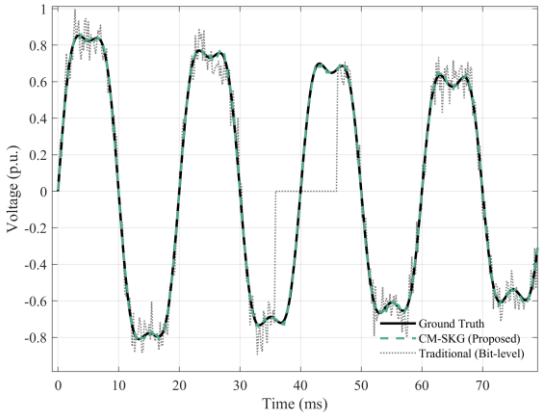

Figure 6. Time-domain voltage waveform reconstruction performance under low-SNR conditions

Figure 6 visualizes the signal reconstruction performance. Under low-SNR conditions, the Traditional method suffers from severe distortion due to packet loss and bit errors, leading to zero-filling artifacts. However, CM-SKG reconstructs a smooth waveform that closely matches the Ground Truth. This demonstrates the noise-resilience of semantic communication. Unlike traditional coding that treats noise as information, the semantic decoder learns the intrinsic topological manifold of valid power signals. It filters out stochastic noise and reconstructs the signal based on semantic meaning, effectively acting as a domain-knowledge-driven denoiser.

Figure 7 quantifies the relationship between semantic compression rate x_t and reconstruction quality (PSNR). A critical phenomenon, the cross-modal compensation zone, is identified. In this ultra-low bandwidth regime, single-modal performance degrades linearly.CM-SKG achieves a significant 5.8 dB gain at the optimal point ($x_t \approx 0.4$). This confirms the information compensation effect: when electrical features are heavily compressed, the aligned visual semantics serve as anchors to guide the reconstruction of electrical waveforms. This mechanism ensures that high-fidelity evidence can be recovered even when the transmission budget is extremely tight.

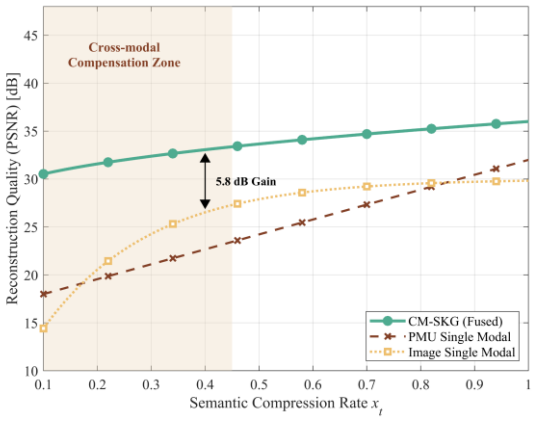

Figure 7 . Impact of semantic compression rate on reconstruction quality (PSNR) and cross-modal compensation effect

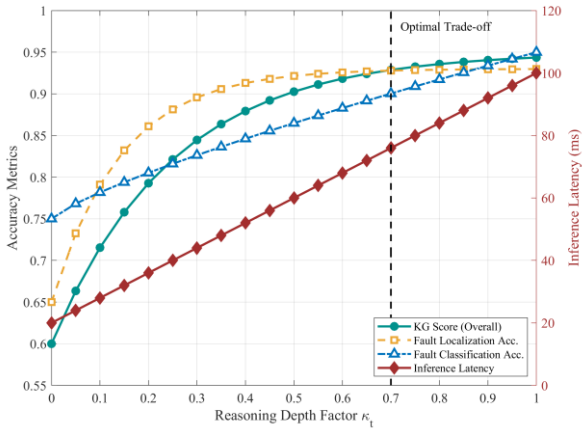

Figure 8. Trade-off analysis between fault classification accuracy and inference latency with varying reasoning depth factors

Figure 8 reveals the non-linear coupling between reasoning depth κ_t and system performance. As κ_t increases from shallow matching to deep reasoning, accuracy improves and saturates, while latency grows linearly. The intersection of the curves highlights an "Optimal Trade-off" region ($\kappa_t \approx 0.65$). The proposed adaptive strategy allows the system to oscillate dynamically around this point. For critical/complex faults, it increases κ_t to ensure reliability; for routine checks, it

reduces κ_t to save time, thereby maximizing the SCE metric.

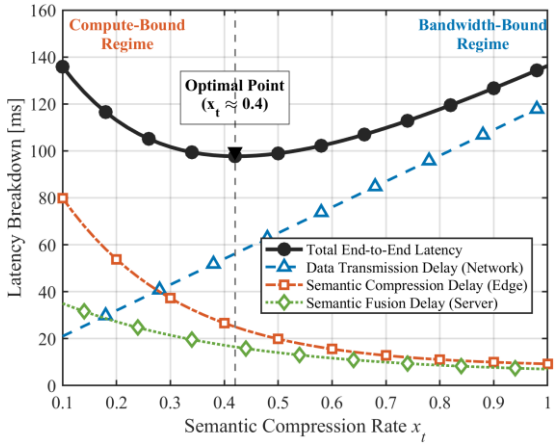

Figure 9. End-to-end latency breakdown and minimization with respect to semantic compression rate

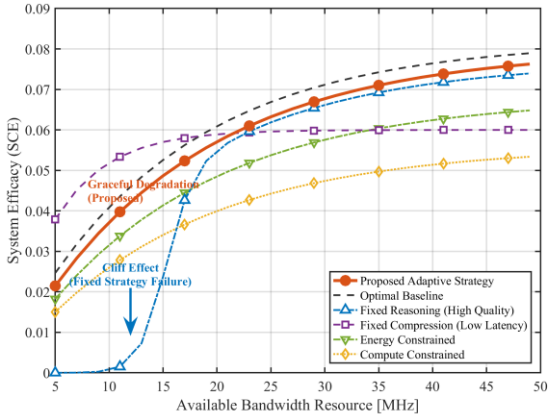

Figure 10. Robustness analysis: System efficacy variation against available bandwidth resource

Figure 9 decomposes the end-to-end latency. The total latency curve exhibits a convex U-shape, indicating two distinct operating regimes: the "Compute-Bound Regime" (left, dominated by compression delay) and the "Bandwidth-Bound Regime" (right, dominated by transmission delay). The proposed algorithm successfully locates the global minimum (Optimal Point $x_t \approx 0.4$). This validates the system's ability to balance computing and communication resources. In edge environments, blindly increasing compression is counterproductive due to computing overhead; CM-SKG finds the mathematically optimal operating point that satisfies real-time FLISR constraints.

Finally, Figure 10 provides a profound insight into the system's robustness mechanism by contrasting the performance trajectories of different strategies under severe bandwidth fluctuations. It is observed that the "Fixed Reasoning" strategy, which prioritizes high-quality transmission, suffers a catastrophic cliff effect when the available bandwidth drops below a critical threshold (approx. 12 MHz). This failure arises from its rigid demand for high-dimensional semantic transmission, which triggers infinite queuing delays and task packet loss under resource constraints. In sharp contrast, the proposed adaptive strategy exhibits superior "Graceful Degradation" characteristics: even under extreme conditions of severe bandwidth scarcity (<10 MHz), the DSAC agent proactively increases the semantic compression ratio and dynamically adjusts the KG reasoning depth, leveraging prior knowledge from the knowledge graph to effectively compensate for the loss of perceptual details. This "soft" adjustment mechanism ensures that the distribution network maintains basic observability and controllability even during severe communication blockage or congestion, fully validating the system's powerful resilience against uncertainty risks in practical engineering.

5. Conclusion

To address the challenges of heterogeneous data fusion and the trade-off between latency and accuracy in distribution network fault handling with high renewable penetration, this paper proposes a semantic communication-enabled Cross-Modal Semantic fusion and Knowledge Graph reasoning (CM-SKG) system. By establishing a semantic closed-loop architecture encompassing device-side compression, edge-side fusion, and master-station reasoning, the proposed system introduces learnable semantic alignment matrices and a bidirectional cross-attention mechanism. This design effectively overcomes the manifold shift and granularity inconsistency between PMU time-series data and UAV visual data, achieving deep mutual verification of heterogeneous semantics in a latent space. Furthermore, a dynamic semantic compression mechanism based on an exponential decay model and a Semantic Communication Efficacy (SCE) metric system—integrating fusion quality, reasoning reliability, and end-to-end latency—are constructed to transform traditional bit-oriented transmission optimization into decision-oriented semantic optimization.

Simulation results demonstrate that the CM-SKG system exhibits significant robustness and adaptability in handling complex fault scenarios. Compared with single-modal or traditional shallow fusion methods, CM-SKG significantly enhances semantic reconstruction quality under strong compression conditions through the cross-modal information compensation effect. Moreover, the DSAC-based joint optimization strategy successfully realizes the online adaptive scheduling of compression rates, reasoning depths, and computational resources, ensuring that the system consistently operates at the

optimal Pareto frontier of the latency-accuracy trade-off under dynamic channel and computational constraints. This study validates the immense potential of semantic communication and knowledge reasoning technologies in the intelligent upgrading of distribution networks, providing a feasible solution to the engineering dilemma where massive data exists but effective information is scarce.

Future work will focus on two aspects: First, extending the current architecture to a broader collaborative scenario spanning cloud, edge, and device layers, and exploring semantic coding/decoding mechanisms based on Large Language Models (LLMs) to further enhance system generalization; second, investigating hardware acceleration techniques for task-oriented semantic communication to reduce energy consumption and computational latency at the edge, thereby promoting the practical engineering deployment of this technology.

Acknowledgements

This work was supported by the project "Research on Intelligent Decision-Making Technology for Distribution Network Operation and Inspection Based on Integrated Semantic Sensing, Communication, and Computation" (Grant No. 5211LS250009).

References

[1] R. He, H. Liang, J. Wu, H. Xie, and M. Shahidehpour, "Reliability assessment of cyber-physical distribution system using multi-dimensional information network model," IEEE Trans. Smart Grid, vol. 14, no. 6, pp. 4683–4692, Nov. 2023.

[2] A. R. Aji, K. M. Banjarnahor, M. S. Hadi, and N. Hariyanto, "Development of FLISR algorithm with additional consideration of reliability index, losses, and load forecasting: Case study in Central Java distribution, Indonesia," in Proc. 6th Int. Conf. Power Eng. Renew. Energy (ICPERE), Bandung, Indonesia, 2024, pp. 1–5.

[3] B. Moussa, A. Al-Barakati, M. Kassouf, M. Debbabi, and C. Assi, "Exploiting the vulnerability of relative data alignment in phasor data concentrators to time synchronization attacks," IEEE Trans. Smart Grid, vol. 11, no. 3, pp. 2541–2551, May 2020.

[4] Y. Zhao and Y. He, "Research on application of multimodal data integration in smart grid fault diagnosis," in Proc. 6th Int. Conf. Electr., Electron. Inf. Commun. Eng. (EEICE), 2025.

[5] J. Zhu, C. Fan, S. Li, K. Nie, J. Li, and Z. He, "SEDNet: Substation equipment detection network with an attention mechanism for UAV automatic power inspection," IEEE Trans. Instrum. Meas., vol. 74, Art. no. 5036114, 2025.

[6] R. Wang, H. Zhang, S. Zhang, H. Liao, and Z. Zhou, "High-frequency acquisition and low-latency interaction method for distribution network data based on sensing-communication collaboration," Trans. China Electrotech. Soc., published online, 2025. (in Chinese)

[7] D. Sarathkumar et al., "A technical review on classification of various faults in smart grid systems," IOP Conf. Ser.: Mater. Sci. Eng., vol. 1055, no. 1, 2021.

[8] N. T. Hoa, C. T. T. Hai, H. L. Hung, N. C. Luong, and D. Niyato, "Joint edge computing and semantic communication in UAV-enabled networks," IEEE Commun. Lett., vol. 29, no. 1, pp. 80–84, Jan. 2025.

[9] H. Liao, W. Che, Z. Zhou, X. Wang, A. Ali, and M. Guizani, "Electric semantic short packet communication: A green ISAC perspective," in Proc. IEEE Int. Conf. Commun. (ICC), Montreal, QC, Canada, 2025, pp. 674–679.

[10] T. M. Getu, G. Kaddoum, and M. Bennis, "Semantic communication: A survey on research landscape, challenges, and future directions," Proc. IEEE, vol. 112, no. 11, pp. 1649–1685, Nov. 2024.

[11] C. E. Shannon and W. Weaver, The Mathematical Theory of Communication. Urbana, IL, USA: Univ. Illinois Press, 1949.

[12] T. M. Getu, G. Kaddoum, and M. Bennis, "Making sense of meaning: Semantic metrics for goal-oriented communication," IEEE Access, vol. 11, pp. 45456–45492, 2023.

[13] T. M. Getu, G. Kaddoum, and M. Bennis, "Tutorial-cum-survey on semantic and goal-oriented communication: Research landscape, challenges, and future directions," Apr. 2023, arXiv:2301.12786.

[14] K. B. Letaief, Y. Shi, J. Lu, and J. Lu, "Edge artificial intelligence for 6G: Vision, enabling technologies, and applications," IEEE J. Sel. Areas Commun., vol. 40, no. 1, pp. 5–36, Jan. 2022.

[15] F. Guo, F. R. Yu, H. Zhang, X. Li, H. Ji, and V. C. M. Leung, "Enabling massive IoT toward 6G: A comprehensive survey," IEEE Internet Things J., vol. 8, no. 15, pp. 11891–11915, Aug. 2021.

[16] L. U. Khan, I. Yaqoob, M. Imran, Z. Han, and C. S. Hong, "6G wireless systems: A vision, architectural elements, and future directions," IEEE Access, vol. 8, pp. 147029–147044, 2020.

[17] Y. Xiao, G. Shi, Y. Li, W. Saad, and H. V. Poor, "Toward self-learning edge intelligence in 6G," IEEE Commun. Mag., vol. 58, no. 12, pp. 34–40, Dec. 2020.

[18] Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, "Edge intelligence: Paving the last mile of artificial intelligence with edge computing," Proc. IEEE, vol. 107, no. 8, pp. 1738–1762, Aug. 2019.

[19] L. Zhou, D. Wu, J. Chen, and X. Wei, "Cross-modal collaborative communications," IEEE Wireless Commun., vol. 27, no. 2, pp. 112–117, Apr. 2020.

[20] M. Chen, M. Liu, W. Wang, H. Dou, and L. Wang, "Cross-modal semantic communications in 6G," in Proc. IEEE/CIC Int. Conf. Commun. China (ICCC), Dalian, China, 2023, pp. 1–6.

[21] S. Zhang et al., "Knowledge-augmented neural reasoning for multi-hop QA," in Proc. ACL, 2021, pp. 1234–1244.

[22] C. Chaccour, W. Saad, M. Debbah, Z. Han, and H. V. Poor, "Less data, more knowledge: Building next generation semantic communication networks," Nov. 2022, arXiv:2211.14343.

[23] Z. Zhang, L. Wang, W. Wu, F. Zhou, and Q. Wu, "Knowledge graph-based explainable and generalized zero-shot semantic communications," IEEE Wireless Commun. Lett., vol. 14, no. 9, pp. 2977–2981, Sep. 2025.

[24] X. Song, F. Zhou, R. Ding, Z. Qu, Y. Li, Q. Wu, and N. Al-Dhahir, "UAV cognitive semantic communications enabled by knowledge graph for robust object detection," IEEE Trans. Commun., vol. 73, no. 8, pp. 6052–6067, Aug. 2025.

[25] J. Li, Z. Wang, and F. Luo, "Power equipment knowledge graph for intelligent fault detection," IEEE Trans. Ind. Informat., vol. 17, no. 12, pp. 8452–8463, 2021.

[26] Z. Sun, H. Lin, and X. Wang, “Dispatch knowledge graph for distributed energy management,” IEEE Trans. Smart Grid, vol. 13, no. 1, pp. 388–400, 2022.
[27] J. Pi, S. Qi, Q. He et al., “Construction and application of knowledge graph for intelligent decision-making of power grid fault handling,” in Proc. IEEE PES Gen. Meeting, 2023, pp. 1–5.
[28] H. Yang et al., “Graph neural networks for knowledge reasoning: Foundations and advances,” IEEE Trans. Neural Netw. Learn. Syst., vol. 33, no. 10, pp. 5275–5293, 2022.
[29] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in Proc. ICML, 2018, pp. 1861–1870.
[30] J. Duan, Y. Guan, S. E. Li, Y. Ren, Q. Sun, and B. Cheng, “Distributional soft actor-critic: Off-policy reinforcement learning for addressing value estimation errors,” IEEE Trans. Neural Netw. Learn. Syst., vol. 33, no. 11, pp. 6584–6598, Nov. 2022.
[31] J. Duan, Y. Ren, F. Zhang, J. Li, S. E. Li, Y. Guan, and K. Li, “Encoding distributional soft actor-critic for autonomous driving in multi-lane scenarios [Research Frontier],” IEEE Comput. Intell. Mag., vol. 19, no. 2, pp. 96–112, 2024.