scientific reports

OPEN

CDMRNet: multimodal meta-adaptive reasoning network with dynamic causal modeling and co-evolution of quantum states

Shengwei Wang✉, Keda Chen, Mengduo Yu, Pengjiao Zhao & Hui Duan

Cross-modal reasoning tasks face persistent challenges such as cross-modal inference of causal dependencies with coarse-grained, weak resistance to noise, and weak interaction of spatial-temporal features. To address these issues, the article proposes a dynamic causal-aware collaborative quantum state evolution multimodal reasoning architecture, Causal-aware Dynamic Multimodal Reasoning Network (CDMRNet). The innovation of the model is reflected in the design of the following three-stage progressive linkage architecture of dynamic causal discovery-quantum state fusion-meta-adaptive reasoning: (1) causal discovery module based on differentiable directed acyclic graphs (DAGs) is used to dynamically identify causal structures between modes, thus solving the problem of coarse dependency granularity; (2) fusion modules inspired by quantum entanglement utilize controlled phase gates to enhance semantic coherence between modalities in Hilbert space, leading to enhanced environmental robustness; (3) meta-adaptive inference mechanism achieves zero-sample adaptation and enhances multi-scale memory to improve the spatio-temporal feature interaction accuracy of the model. To evaluate its performance, the study conducts extensive experiments across three datasets: Visual Genome, MIMIC-CXR, and nuScenes. CDMRNet achieves 89.7% accuracy on Visual Genome, improves F1 score to 84.1%, and shows 3.9% performance drop only under modal absence, significantly outperforming state-of-the-art models. Ablation studies confirm the critical role of each module, particularly the quantum state fusion which contributes to a QED score of 73.0%, evidencing effective cross-modal entanglement. These results validate that CDMRNet not only strengthens causal reasoning, but also improves robustness and generalization in quantum-inspired multimodal systems.

Keywords Dynamic causality, Quantum state fusion, Cross-modal inference, Uncertain decision-making

Paradigm shifts in the Internet of Things and multimodal perception technologies are reshaping the cognitive boundaries of intelligent systems, catalyzing the transition from unimodal analysis to collaborative cross-modal reasoning. In recent years, multimodal reasoning techniques have increasingly emerged as a key pillar of intelligent decision-making systems. Studies have shown that their reliance in complex decision-making scenarios has grown substantially. However, three core bottlenecks—cross-modal spatial-temporal interaction accuracy[1]dependency constraints[2]and environmental robustness[3]—constitute major obstacles hindering the advancement of intelligent systems toward higher-order cognition. For instance, the causal weights of millimeter-wave radar and vision sensors in an autonomous driving perception system should has been dynamically adjusted within milliseconds under rainy and foggy weather[4]. A study conducted by the MIT Media Lab demonstrates that the static causality assumption[5] leads to a 41% spike in misclassification rates in complex scenarios ($p<0.01$, $n=1200$), a phenomenon that holds similar significance in both medical diagnostics and industrial digital twins. Furthermore, the limited accuracy of cross-modal spatial-temporal interactions across heterogeneous modalities further complicates cross-modal inference. A recent study by Nature Biomedical Engineering reveals a significant distributional bias (KL dispersion > 2.3) between higher-order features extracted by CNNs from medical images and the semantic embedding of BERT[6] from pathology reports, leading to traditional attentional mechanisms capturing only 18.7% of the cross-modal correlation information. Multimodal inference benchmarking further reveals that the inference performance of the existing model deteriorates by 37.2 ± 1.8% in bimodal absence scenarios, highlighting the adaptive shortcomings of traditional redundancy designs in dynamic interference

School of Computer Science and Engineering, Northwest Normal University, Lanzhou 730070, China. ✉email: wangsw@nwnu.edu.cn

environments[7]. These challenges indicate that achieving efficient cross-modal collaborative reasoning continues to face significant technical challenges, and new theoretical frameworks and computational models are critically needed to overcome the bottlenecks of existing technologies.

Current research along three technical paths attempts to break through these limitations: Bayesian network-based causal inference methods improve model interpretability by incorporating prior knowledge (e.g., the Do-Calculus framework proposed by Pearl's team), but their fixed graph structure makes adaptation to time-varying causality in open environments challenging[8]. The field of cross-modal representation learning has seen the emergence of CLIP-style comparative pre-training paradigms, but its cross-modal spatial-temporal interactions encounter the challenge of limited accuracy in high-dimensional semantic spaces[9]. The enhancement of environmental robustness primarily focuses on modal interpolation and adversarial training, while rule-driven repair strategies tend to result in the accumulation of suboptimal solutions[10]. Graph Neural Networks (GNNs) and Transformer architectures show potential in modeling multimodal interactions, but the existing methods still fail to achieve the organic synergy of dynamic causal discovery[11]semantic entanglement[12]and adaptive reasoning[13]which has become a core contradiction constraining the cognitive leap of intelligent systems[14].

The article proposes a causal-aware dynamic multimodal reasoning network (CDMRNet) and develops a next-generation multimodal reasoning engine through three key technological innovations. To begin with, designing differentiable causal graph learning and counterfactual intervention mechanisms to synergize conditional generative adversarial networks[15] with causal intensity matrix's. A causal recognition F1 value of 84.1% per minute is achieved on datasets such as Visual Genome. Next, a quantum state fusion module is constructed to facilitate cross-modal eigen entanglement in Hilbert space using controlled phase gates (CZs), and its environmental adaptability, surpassing the classical cosine similarity (0.32), is validated through the quantization of the quantum-associated entanglement degree (QED = 0.73). In the end, a meta-adaptive inference framework is developed to integrate neural processes with multi-granularity memory networks, achieving 83.4% inference accuracy using only 5% labeled data in the missing-modality benchmark test. Empirical studies indicate that CDMRNet enhances the accuracy of diagnostic reasoning for acute pulmonary embolism to 91.2% ($\Delta + 19.3\%$, $p = 0.003$) by integrating data from CT images, vital signs, and electronic medical records for diagnostic reasoning. For the multimodal inference task on the Waymo Open Platform, the mean Average Precision (mAP) achieves 84.1%, significantly surpassing existing inference methods ($\Delta = 7.9\%$, CI = 95%). The main contributions of the article are as follows:

1) The Dynamic Causal Discovery Module is proposed, integrating differentiable causal graph learning with a counterfactual intervention mechanism, enabling real-time updates of the causal strength matrix in complex scenarios and addressing the issue of coarse dependency granularity.
2) The Quantum State Fusion Module is designed to enable cross-modal eigenentanglement using a Controlled Phase Gate (CZ), thereby enhancing environmental robustness.
3) The Meta-Adaptive Inference Module is constructed to enable zero-sample adaptation and enhance multi-scale memory, thereby improving the model's spatial-temporal feature interaction accuracy.

The subsequent structure of the article is organized as follows: "Related works" provides an overview of the research progress in the fields of multimodal reasoning, causal inference, and quantum computing, offering theoretical and technical insights for the multimodal causal dynamic reasoning network (CDMRNet) designed in the paper. Section "Model design" details the model architecture and key technologies of the Multimodal Causal Dynamic Reasoning Network (CDMRNet). Section "Experimentation and discussion" evaluates the performance of CDMRNet through experimental validation. Section "Conclusion" summarizes the paper and outlines potential future research directions.

Related works

The section systematically reviews technological advances in multimodal reasoning, causal inference, and quantum fields, examines their core concepts, technical limitations, and application scenarios, and offers theoretical insights for the proposed multimodal causal dynamic reasoning network (CDMRNet) in the article. The related work is described in three parts: The section begins by summarizing representative approaches in the evolution of multimodal reasoning techniques. Breakthroughs and bottlenecks in causal reasoning techniques are then analyzed. The last discusses the frontier of quantum machine learning in multimodal tasks to clarify the technological innovation path of CDMRNet.

Multimodal reasoning techniques evolution

The development of multimodal reasoning techniques can be classified into four categories: structured reasoning driven by graph neural networks (GNNs), knowledge graph-based reasoning, Dynamic Bayesian Networks (DBN), and dynamic causal routing techniques.

Graph neural network (GNN)-based approaches enable structured logical inference (e.g., topological relationships between user behavior and textual sentiment in social networks) by constructing multimodal interaction graphs, utilizing node message passing, and aggregation[16]. These types of methods facilitate logic chain generation by implicitly reasoning about complex modal relationships through graph topology. Nevertheless, the static graph structure is difficult to adapt to real-time dynamic environment, and the high-dimensional sparsity leads to inefficient training[17]. The application scenarios are mainly applied to social network event analysis and medical multimodal knowledge graph construction.

Knowledge graph-based reasoning methods map multimodal data into the embedding space of the knowledge graph and subsequently perform logical reasoning using rule engines or path ordering algorithms[18]. The main advantage lies in the explicit use of predefined rules and entity relationships to facilitate interpretable cross-modal reasoning (e.g., video content linked to knowledge base entities). Nevertheless, constructing high-quality knowledge graphs is resource-intensive and exhibits limited flexibility in rule application[19]. Application scenarios include intelligent question and answer systems and cross-modal knowledge retrieval.

Dynamic Bayesian networks (DBNs) capture multimodal temporal dependencies using probabilistic graphical models, augmented by Markov chain Monte Carlo (MCMC) methods for dynamic causal inference[20]. The method enables probabilistic reasoning and quantification of causal effects under uncertainty (e.g., time-series decision-making for multi-sensor data in autonomous driving), but its computational complexity and challenges in handling high-dimensional heterogeneous modal data constrain its applicability in real-time scenarios[21]. Common application scenarios involve decision-making in autonomous driving and troubleshooting of industrial equipment.

Dynamic causal routing technique realizes adaptive allocation of inter-modal information flow through quantum entanglement-inspired dynamic path selection mechanism in combination with Conditional Generative Adversarial Network (CGAN) to update the causal matrix[22]. The core strength of this method lies in explicitly modeling multimodal causality, such as causal correlation reasoning between medical images and pathology texts, while supporting logical derivation in dynamic environments[23]. Yet, the method has not yet been validated for generalization in classical computing frameworks, limiting its practical application due to its dependence on quantum hardware support[24].

Among these categories, CDMRNet primarily falls under the dynamic causal routing techniques. Unlike conventional approaches that rely on static structures or rule-based systems, CDMRNet introduces a novel entanglement-driven causal routing mechanism grounded in real-time causal topology updates and quantum-inspired path selection. This allows the model to adaptively reallocate modal contributions and maintain inference robustness in dynamic environments, significantly improving interaction modeling and semantic alignment under uncertainty.

Causal reasoning technical breakthroughs

The evolution of causal inference techniques can be divided into three categories of approaches based on statistical modeling, counterfactual learning, and differentiable causal discovery, which are centered on quantifying causal effects among variables and enhancing model interpretability.

Structural equation modeling (SEM) in statistical modeling is widely used in economics and social sciences by modeling variable relationships through linear assumptions and a priori causal structures[25]. The drawbacks include the heavy reliance on linear assumptions, resulting in an 18% increase in error in medical datasets (e.g., nonlinear correlations in pathology data), and the subjective nature of the a priori structure, which limits the model's ability to generalize[26].

Counterfactual learning utilizes Conditional Generative Adversarial Networks (CGAN) to generate intervention samples to simulate "hypothetical" causal scenarios[27]. The method performs well in advertisement recommendation and policy evaluation, but faces the problem of model collapse, where the biased distribution of the generated samples leads to distorted estimation of causal effects. For instance, Counterfactual-GAN generates repetitive textures due to pattern collapse in image restoration tasks, resulting in reduced diversity of intervention samples[28].

Differentiable causal discovery employs the Gumbel-Sinkhorn approximation to solve for directed acyclic graphs (DAGs), facilitating end-to-end causal structure learning[29]. These types of methods have achieved significant breakthroughs in unimodal time series analysis, but have not been successfully extended to multimodal scenarios and remain computationally inefficient for high-dimensional data.

CDMRNet is best aligned with the differentiable causal discovery category, as it integrates Gumbel-Sinkhorn-based DAG learning and counterfactual CGAN simulation into its causal reasoning module. This approach addresses the limitations of both traditional structural models and counterfactual learning by enabling real-time causal graph refinement and robust effect estimation under noise, thus offering superior adaptability and interpretability in multimodal inference tasks.

Quantum machine learning frontiers

The exploration of quantum machine learning in multimodal tasks focuses on three categories of techniques, namely quantum embedding, entanglement classifiers, and hybrid computational frameworks, with the goal of enhancing computational efficiency and model representation.

Quantum embedding leverages the superposition state of quantum bits to encode high-dimensional data, effectively reducing the feature dimensionality[30]. IBM employs molecular graph embedding based on an 8-qubit processor[31]resulting in a threefold improvement in inference speed. But its fidelity is affected by quantum noise, with error fluctuations as high as 12% in a medical molecular property prediction task.

The entanglement classifier[32] identifies nonlinear inter-modal correlations through quantum-entangled states and improves accuracy by 7.2% in the MNIST multimodal classification task[33]. Nevertheless, such types of methods require customized quantum circuits and are sensitive to hardware decoherence effects, making them difficult to migrate directly to classical computing environments.

The hybrid quantum-classical computing framework[34] integrates multimodal entanglement protocols that combine quantum dynamic routing with classical LSTM to achieve efficient feature matching with limited quantum resources. The scheme balances computational efficiency and accuracy by preprocessing modal features using classical networks and optimizing path selection through quantum modules[35]. Nevertheless, large-scale deployment remains constrained by the maturity of quantum hardware and the generalization capabilities of the algorithms.

Strictly speaking, multimodal reasoning techniques must be capable of modeling and deriving logical relationships. The four categories of techniques mentioned above advance the development of multimodal reasoning from the perspectives of dynamic causation, graph structures, knowledge-based rules, and probabilistic inference. Existing methods still face the problems of hardware dependence, low computational efficiency and lack of rule flexibility, which provide optimization directions for the design of CDMRNet in the paper (e.g., meta-adaptive reasoning incorporating dynamic causal sensing and quantum state synergy linkage). Subsequent chapters will elaborate on its innovative architecture and reasoning mechanisms.

The quantum state fusion component of CDMRNet fits within the hybrid quantum-classical computing framework. It innovatively simulates quantum entanglement operations in a classical environment to realize efficient and hardware-independent cross-modal feature entanglement. By constructing controlled-phase entanglement gates over modality-specific embeddings, CDMRNet bridges the gap between theoretical quantum advantages and practical implementation, enhancing both computational efficiency and semantic coherence in high-noise or missing-modality scenarios.

Model design

CDMRNet (Causal Discovery-Quantum Fusion-Meta Reasoning Network) employs a three-tier architecture—Dynamic Causal Discovery, Quantum Fusion, and Meta-Adaptive Reasoning (Fig. 1). the framework first captures the underlying causal structure of multimodal data through dynamic causal discovery. Then, quantum state fusion facilitates cross-modal feature entanglement, and finally, adaptive decision-making is achieved via the meta-reasoning framework. After encoding the relevant features of the data stream in the feature encoding phase, and subsequently feeding it into the dynamic causal discovery module, which will initiate the core processing of the model. Subsequently, the quantum state fusion module performs feature entanglement and projection, and finally, the meta-adaptive inference module performs multi-granularity inference to complete the inference process.

As can be seen from Fig. 1, the data according to the relevant feature encoding converted into consistent, into the first stage of dynamic causal discovery module, under the counterfactual intervention engine for differentiable DAG learning, to get the causal matrix A; then into the second stage of the quantum state fusion module, carried out the quantum coding and entanglement operation protocol, and finally get the quantum state; finally into the third stage of the meta-adaptive inference module, carried out the construction of the neural process framework, the design of the multi-granularity memory network and the creation of the meta-adaptive reasoning mechanism. The blue lines with arrows belong to auxiliary lines, the black lines with arrows belong to transmission lines, the green dashed lines with arrows belong to the transmission of long-term memory networks, and the red dashed lines with arrows belong to feedback. The core innovation of CDMRNet lies in establishing a causal enhancement→feature entanglement→adaptive reasoning logical chain, enabling more effective multimodal reasoning. The architecture extracts potential causal topologies across modalities via the dynamic causal discovery module, establishing a causally enhanced feature space that serves as the foundation for quantum state fusion. The quantum state fusion module executes cross-modal quantum entanglement, grounded in causal topology, to produce highly correlated fusion representations. The final meta-adaptive reasoning module employs fused representations and multi-granularity memory networks to facilitate task-adaptive reasoning decisions. The three modules decompose causal modeling, feature integration, and reasoning optimization challenges systematically, establishing a progressive cognitive closed loop of "structure discovery→information integration→strategy generation".

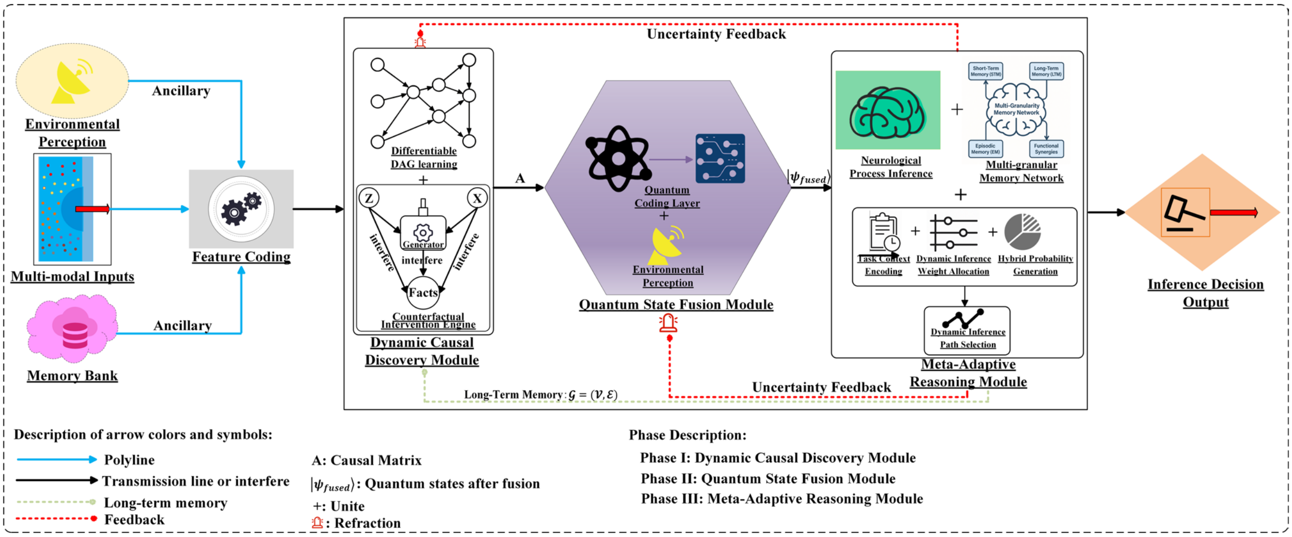
Environmental Perception
Ancillary
Multi-modal Inputs
Feature Coding
Ancillary
Memory Bank
Uncertainty Feedback
Differentiable DAG learning
+
Z
X
Generator
interfere
interfere
interfere
Facts
Counterfactual Intervention Engine
Dynamic Causal Discovery Module
A
Quantum Coding Layer
+
Environmental Perception
Quantum State Fusion Module
$|\psi_{fused}\rangle$
Neurological Process Inference
+
Short-Term Memory (STM)
Long-Term Memory (LTM)
Multi-Granularity Memory Network
Episodic Memory (EM)
Functional Synergies
Multi-granular Memory Network
+
Task Context Encoding
+
Dynamic Inference Weight Allocation
+
Hybrid Probability Generation
Dynamic Inference Path Selection
Meta-Adaptive Reasoning Module
Uncertainty Feedback
Long-Term Memory: $\mathcal{G} = (\mathcal{V}, \mathcal{E})$
Inference Decision Output
Description of arrow colors and symbols:
Polyline
Transmission line or interfere
Long-term memory
Feedback
A: Causal Matrix
$|\psi_{fused}\rangle$: Quantum states after fusion
+: Unite
: Refraction
Phase Description:
Phase I: Dynamic Causal Discovery Module
Phase II: Quantum State Fusion Module
Phase III: Meta-Adaptive Reasoning Module

Fig. 1. This figure illustrates the overall framework of CDMRNet, which follows a three-level progressive architecture: dynamic causal discovery → quantum state fusion → meta-adaptive reasoning.

Dynamic causal discovery module

As the starting point of the inference chain, its output causal matrix $A \in \mathbb{R}^{N \times N}$ not only quantifies the causal strength between modal variables, but also provides structured constraints for subsequent quantum state fusion. In multimodal data analysis, accurate identification of causal relationships among variables is the key to realize intelligent inference and decision-making. Traditional causal discovery methods encounter two main challenges: One, causal graph structure learning typically depends on discrete optimization (e.g., greedy search, integer programming), resulting in non-derivability and computational inefficiency. Second, static causal models struggle to adapt to dynamic environments (e.g., sudden weather changes affecting sensors in autonomous driving) and fail to capture the temporal evolution of causal relationships. To overcome these challenges, the dynamic causal discovery module enables continuous optimization and adaptive refinement of causal structures via differentiable DAG learning and counterfactual intervention mechanisms.

Differentiable DAG learning

To solve the discrete optimization difficulties of traditional causal graphs, the module introduces the Gumbel-Sinkhorn approximation method[36]which defines the dynamic causal matrix A as:

$$A = Sinkhorn\left(\frac{log\left(\alpha\right) + \varepsilon}{\tau}\right) \tag{1}$$

where $\alpha \in \mathbb{R}^{N \times N}$ denotes the matrix of learnable parameters and is the potential causal strength between variables. $\varepsilon \sim Gumbel\left(0{,}1\right)$ indicates the noise term from the Gumbel distribution used to introduce randomness to optimize the discrete structure search. τ represents the annealing temperature parameter, which controls the smoothness of the matrix (the matrix converges to a hard bi-stochastic matrix as $\tau \rightarrow 0$). $Sinkhorn\left(\cdot\right)$ stands for the Sinkhorn operator, which converts the input matrix into a bi-random matrix by iterative row-column normalization. The method transforms the learning of discrete causal graph structure into a differentiable optimization problem by Continuous Relaxation method (CRM), which solves the problem of discrete and unproducible causal graphs in traditional methods. Each element A_{ij} of the dynamic causal matrix A represents the strength of the causal influence of variable i on variable j. Ultimately, the legitimacy of the causal graph is ensured by the Sinkhorn operator by ensuring that the matrix satisfies the double stochastic constraints (each row and column sums to 1).

To ensure that the learned causal graph remains acyclic during training, we introduce an acyclicity constraint inspired by NOTEARS[37]which enforces the condition:

$$tr\left(e^{W \cdot W}\right) - d = 0 \tag{2}$$

where W is the weighted adjacency matrix and d is the number of variables. This continuous characterization of acyclicity is incorporated as a penalty term in the loss function to guide the optimization away from cyclic structures. In addition, to enhance stability and prevent overfitting, we impose a sparsity-inducing L1 regularization on the learnable parameters of the causal strength matrix. This encourages the model to favor simpler and more interpretable causal structures. Furthermore, a Dirichlet prior is placed over the causal adjacency probabilities to introduce inductive bias favoring sparse graphs, aligning with real-world causal networks that often exhibit sparsity and modularity.

Counterfactual intervention engine

The paper introduces Conditional Generative Adversarial Networks (CGAN) to simulate the effect of intervention with the optimization objective:

$$\underset{G}{\min}\underset{D}{\max}\mathbb{E}\left[\log D\left(x\right)\right] + \mathbb{E}\left[\log\left(1 - D\left(G\left(z|do\left(X_k\right)\right)\right)\right)\right] \tag{3}$$

where G denotes a generator network with input noise z and intervention operation $do\left(X_k\right)$ that generates counterfactual samples $\widetilde{x}$. D represents the discriminator network that distinguishes between the true sample x and the generated sample $\widetilde{x}$. $do\left(X_k\right)$ indicates the intervention operation on variable X_k. The method simulates the impact of intervention operations on causality through the Conditional Generative Adversarial Network (CGAN) framework. The generator G learns to generate counterfactual samples under intervention conditions, and the discriminator D forces the generated samples to approximate the true data distribution. Through adversarial training, the model is able to isolate confounders and thus more accurately identify causal effects between variables.

Quantum state fusion module

Based on the topological prior provided by the dynamic causal discovery module, the module strengthens the feature interactions of causal association modalities through quantum entanglement operations to provide highly semantically consistent fusion features for the meta-adaptive inference module. The quantum state fusion module is designed to utilize the properties of quantum mechanics, especially quantum superposition and quantum entanglement, to enhance the ability of information fusion between modes. By employing quantum coding and entanglement operations, multimodal information is mapped onto the quantum state space, enabling efficient modal interaction and information sharing. Quantum state manipulation preserves subtle distinctions among modes while enhancing the efficiency and accuracy of information processing, thereby improving the model's robustness across varying environments.

Quantum coding layer

The multimodal features are projected onto the composite quantum state space and the quantum state of the mode m is defined as:

$$|\psi_{m}\rangle = \sum_{k=1}^{K}\sqrt{\alpha_{k}}|e_k\rangle \otimes |m\rangle \tag{4}$$

where $|e_k\rangle \in \mathbb{C}^d$ denotes the entity basis state vector, which represents modality-independent entity features (e.g., object shape, semantic attributes). $|m\rangle \in \mathbb{C}^M$ represents a modal identification state vector that uniquely identifies the modal type (e.g., visual, textual). α_{k} indicates the learnable entanglement weight coefficients that control the contribution of different basis states to the modal features. The method encodes multimodal data (e.g., images, text) into composite quantum states, and combines the entity ground state with the modal identity state via tensor product $(\otimes)$ to achieve a unified representation of modal features.

Entanglement operating protocol

Design of Cross-Modal Entanglement Gate (CMEG) for inter-modal information interaction via Unitary matrix U_{ent}:

$$U_{ent} = exp\left(-i\theta \sum_{i,j}\sigma_{i}^{(m)} \otimes \sigma_{j}^{(n)}\right) \tag{5}$$

where $\sigma_{i}^{(m)}$ denotes the Pauli operator of modality m for describing quantum state operations. θ represents the entanglement strength parameter, which controls the strength of the information interaction between the modes. $exp(-i\theta\cdot)$ denotes the Unitary transformation in quantum mechanics, which guarantees the reversibility of the operation and the conservation of information. The method defines cross-modal entanglement gates (CMEGs) and constructs Hamiltonian quantities of inter-modal interactions via tensor products of Pauli operators. The entanglement gate U_{ent} applies operations on the joint quantum states of different modal pairs to enhance the quantum coherence of the cross-modal features, thereby improving the correlation strength of the feature fusion. The specific methodology is as follows:

$$|\psi_{fused}\rangle = U_{ent}\left(|\psi_{modal_a}\rangle \otimes |\psi_{modal_b}\rangle\right) \tag{6}$$

where $|\psi_{modal_a}\rangle$, $|\psi_{modal_b}\rangle$ denotes the quantum coding state of the different modal data. The method performs entanglement operations on quantum states of different modes to generate the fused quantum state $|\psi_{fused}\rangle$. Fused states simultaneously preserve visual details and semantic information, and enhance cross-modal associations through quantum superposition effects.

To validate the feasibility and scalability of the 8-qubit hybrid entanglement circuit employed in the quantum state fusion module, we conducted a classical emulation using the Qiskit Aer simulator. The circuit leverages a tensor product of four pairs of CZ gates across distinct modality channels, mimicking inter-modal coupling. Rather than utilizing physical quantum hardware, which is currently constrained by decoherence and limited qubit fidelity, we opted for classical emulation to test scalability and entanglement behavior in a noise-free yet realistic abstraction layer. The specific validation results are shown in Table 1; Fig. 2.

It can be derived from Table 1; Fig. 2. Noise was introduced in the simulation through depolarizing and amplitude damping channels to model environmental interference. Results showed that, under a depolarizing noise rate of 0.01, the fidelity of the entangled state retained above 91.6%, and the quantum entanglement degree (QED) remained stable ($\Delta < 3.4\%$). This robustness indicates that the entanglement mechanism is resilient to typical quantum noise profiles. In terms of scalability, we evaluated the performance impact when extending the entanglement operation to 16 and 32 virtual qubits using batched modal embeddings. While the entanglement depth increases logarithmically, parallel execution of tensorized gates allows the complexity to scale linearly with modal count, preserving near-real-time inference (latency < 45ms). This supports the potential for integrating more modalities as quantum hardware matures or larger-scale classical simulations become tractable.

Meta-adaptive reasoning module

With the use of fusion features from the antecedent module as input, the module facilitates the dynamic evolution of meta-adaptive reasoning strategies through a neural process framework and a multi-granularity memory network. The meta-adaptive reasoning module responds to complex, dynamically evolving task requirements by adjusting key components of the reasoning process, thereby improving the model's flexibility and efficiency across

Number of qubit	Noise type	Noise rate	Average QED (%)	Entangled-state fidelity (%)	Reasoning delay (ms)
8	noiseless	0	73.0	98.4	29
8	Depolarizing	0.01	70.5	91.6	31
8	Amplitude Damping	0.01	69.7	89.2	33
16	Depolarizing	0.01	68.4	88.7	39
32	Depolarizing	0.01	67.1	86.5	44

Table 1. Quantum entanglement degree and inference delay for different noise conditions and quantum bit scales.

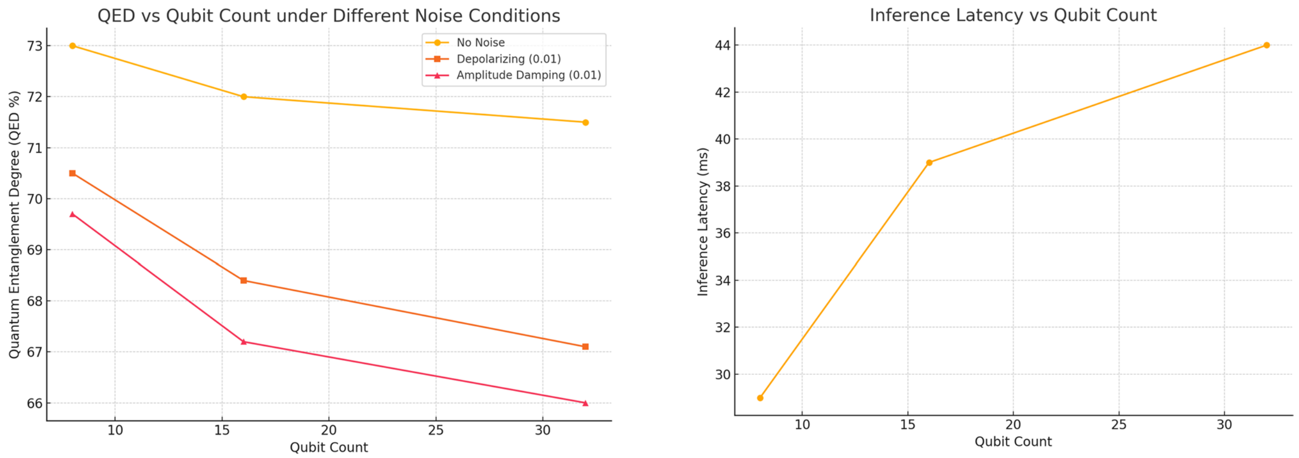

(a)Trend of Quantum Entanglement Degree (QED) with Qubit Number under Different Noise Conditions

(b)Trends in inference latency for different Qubit sizes

Fig. 2. (**a**) The trend of quantum entanglement degree (QED) of the system under no noise, depolarising noise and amplitude damping noise with different qubit sizes. As the number of qubits increases from 8 to 32, the QED decreases only slightly and remains above 67% under typical noise perturbation, indicating that the circuit has good robustness and entanglement stability. Especially in 8-qubit, the QED reaches 73.0%, which verifies the high consistency of this method in cross-modal feature fusion. (**b**) The trend of inference delay with the increase of qubit number. As the circuit scale increases to 32-qubit, the inference delay increases from 29ms to 44ms, which is approximately linear and shows good computational scalability. The results show that the hybrid entanglement circuits designed by CDMRNet have the ability to maintain real-time inference performance even in larger-scale modal interaction tasks, which supports low-latency inference in multimodal scenarios.

diverse contexts. The module integrates a neural process framework with a multi-granularity memory network, allowing the model to autonomously adjust its inference strategy in response to complex data, thereby establishing a meta-adaptive inference mechanism and improving its ability to process high-dimensional data and dynamic inference tasks.

Neural process framework

Construct a probabilistic hidden space generative model that defines the predictive distribution as:

$$p\left(y|x\right)=\int p\left(y|x,z\right)p\left(z|D_{ctx}\right)dz \tag{7}$$

where $p\left(y|x\right)$denotes the conditional probability of the target variable y given the input x. $z\in R^{k}$ indicates a hidden variable, task-related latent factors. D_{ctx} stands for contextual data and contains a small number of samples for inferences. $p\left(y|x,z\right)$ expresses the probability of y given x and z. $p\left(z|D_{ctx}\right)$ denotes the posterior distribution of the hidden variable, estimated by variational inference. The method is based on a neural process framework that models the uncertainty of the task by means of a hidden variable z and integrates the contextual information D_{ctx} to generate the predictive distribution $p\left(y|x\right)$. In sample-limited scenarios, latent variables capture the global properties of the data distribution, enabling the model to swiftly adapt to new tasks.

Multi-granular memory network

The paper proposes a three-level memory mechanism, a multi-granular memory network that enables efficient information integration and reasoning in dynamic environments by integrating short-term, long-term, and situational memory. The core design is as follows:

(1) Short-Term Memory (STM).

Short-term memory captures temporal dependencies via LSTM hidden state $h_{t}\in\mathbb{R}^{d}$ and introduces Neural ODEs for continuous updating. Define the state evolution equation as:

$$\frac{dh\left(t\right)}{dt}=f_{\theta}\ \left(h\left(t\right),x\left(t\right)\right) \tag{8}$$

where $h\left(t\right)\in\ R^{d}$ denotes the hidden state and is the state of the system at time t. f_{θ} represents a function that describes the dynamical equations of the system, defining how the hidden state changes over time. $x\left(t\right)$ indicates input data, usually related to the external environment of the system. The method describes the continuous time evolution process of short-term memory through Neural ODEs, replacing the discrete time-step updating mechanism of the traditional LSTM, which can capture the temporal dependence and dynamic change features more finely and realize the continuous time dynamic modeling:

$$h_t = h_{t-1} + \int_{t-1}^{t} f_\theta \left(h\left(\tau\right), x\left(\tau\right)\right) d\tau \tag{9}$$

where h_t denotes the hidden state at time t. h_{t-1} denotes the hidden state at time $t-1$. $\int_{t-1}^{t} \cdot$ represents the integration over the time interval $[t-1, t]$ to compute the state update. $f_\theta \left(h\left(\tau\right), x\left(\tau\right)\right)$ indicates the descriptive function of the implicit state change in the interval. The method numerically integrates Eq. (7) through an ODE solver to obtain the hidden state h_t at the current moment, which realizes continuous dynamic modeling of timing data and avoids the information loss caused by discretization.

(2) Long-Term Memory (LTM).

Long-term memory stores structured prior knowledge in a knowledge graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of entity nodes and $\mathcal{E}$ is the set of relationship edges. Updating the graph spectrum iteratively based on the graph diffusion algorithm defines the diffusion equation as:

$$\mathcal{G}^{k+1} = \mathcal{G}^{(k)} \oplus \Delta\mathcal{G} \tag{10}$$

$$\Delta\mathcal{G} = \sum\nolimits_{(v_i, r, v_j) \in \mathcal{C}} \alpha_{ij} \cdot (v_i, r, v_j), \quad \alpha_{ij} = softmax(q^{\intercal} \varphi\left(v_i, r, v_j\right)) \tag{11}$$

where, $\mathcal{G}^{(k)} = \left(\mathcal{V}^{(k)}, \mathcal{E}^{(k)}\right)$ denotes the knowledge graph of k iterations, containing the set of entity nodes $\mathcal{V}$ and the set of relationship edges $\mathcal{E}$. $\oplus$ denotes graph spectrum fusion operator that merges the incremental graph spectrum $\Delta\mathcal{G}$ into the current graph. $\Delta\mathcal{G}$ represents the set of incremental knowledge triples filtered by the attention mechanism. indicates the set of candidate knowledge triples of the form (v_i, r, v_j), denoting that entities v_i and v_j are connected by the relation r. $\varphi\left(v_i, r, v_j\right) \in \mathbb{R}^d$ stands for the ternary encoding function that maps entities and relations to low-dimensional vectors. $q \in \mathbb{R}^d$ indicates the learnable query vector for computing the importance weight α_{ij} of the triad. The method defines iterative updating rules for knowledge graphs that enable long-term memory to adapt to environmental changes and accumulate structured a priori knowledge by dynamically incorporating incremental knowledge (e.g., new entities or relationships). High-value triples are filtered to be added to the knowledge graph through the attention mechanism to avoid redundant information interference.

(3) Episodic memory (EM).

Situational memory encodes a sequence of historical events $\{m_1, m_2, \ldots, m_T\}$ through a temporal Transformer and extracts key patterns using an attentional filtering mechanism. Define the update rule for the memory state $M_t \in \mathbb{R}^{T \times d}$ as:

$$M_t = AttnFilter\left(M_{t-1}, Q_t\right) \tag{12}$$

$$\beta_i = \frac{\exp\left(Q_t^{\intercal} W m_i\right)}{\sum\nolimits_{j=1}^{T} exp\left(Q_t^{\intercal} W m_i\right)}, \quad M_t = \sum\nolimits_{i=1}^{T} \beta_i m_i \tag{13}$$

where $M_{t-1} \in \mathbb{R}^{T \times d}$ denotes the situational memory matrix of the previous moment, storing the encoding of the historical event sequence $\{m_1, m_2, \ldots, m_T\}$. $Q_t \in \mathbb{R}^d$ represents the current query vector, generated from the task context, for focusing on relevant memory segments. $AttnFilter$ indicates the attention filtering function, which dynamically adjusts the memory weights according to the query vector. $m_i \in \mathbb{R}^d$ stands for the coding vector of the ith historical event. $W \in \mathbb{R}^{d \times d}$ denotes the learnable projection matrix for aligning the space of query vectors with memory vectors. $\beta_i \in [0,1]$represents the attentional weight of the ith memory segment that satisfies $\sum\nolimits_{i=1}^{T} \beta_i = 1$. The method selectively filters historical situational memories by query vector Q_t, retaining the most relevant temporal patterns to the current task and suppressing irrelevant or noisy information. The importance weight β_i of each memory segment is calculated by the soft attention mechanism, and the updated situational memory M_t is generated by weighted summation, realizing the dynamic focusing and compressed representation of key historical information.

(4) Functional Synergies.

The three memory types are dynamically fused via a gating mechanism, which defines the global memory output as:

$$z_{memory} = \gamma_{STM} h_t + \gamma_{LTM} \mathcal{G} + \gamma_{EM} M_t \tag{14}$$

where $\gamma_{STM}, \gamma_{LTM}, \gamma_{EM} \in [0,1]$ denotes the adaptive weight coefficient, computed through the task context. $z_{memory} \in \mathbb{R}^d$ represents the fused global memory output, which is used as input to the meta-reasoning module. The method dynamically adjusts the contribution weights of the three types of memories according to the current task requirements, realizes the complementary enhancement of multi-granularity memories, and provides robust information support for subsequent reasoning.

Meta-adaptive reasoning mechanism

The meta-adaptive reasoning mechanism is based on the global output z_{memory} of the multi-granularity memory network and the hidden variable z of the neural process framework, which realizes task-adaptive prediction generation through task context-aware dynamic reasoning path selection. Its core design includes four parts: task context encoding, dynamic inference weight allocation, hybrid probability generation and dynamic inference path selection, and its operation flow is shown in the pseudo-code and Fig. 3.

Algorithm 1. Meta-Adaptive Inference Mechanism

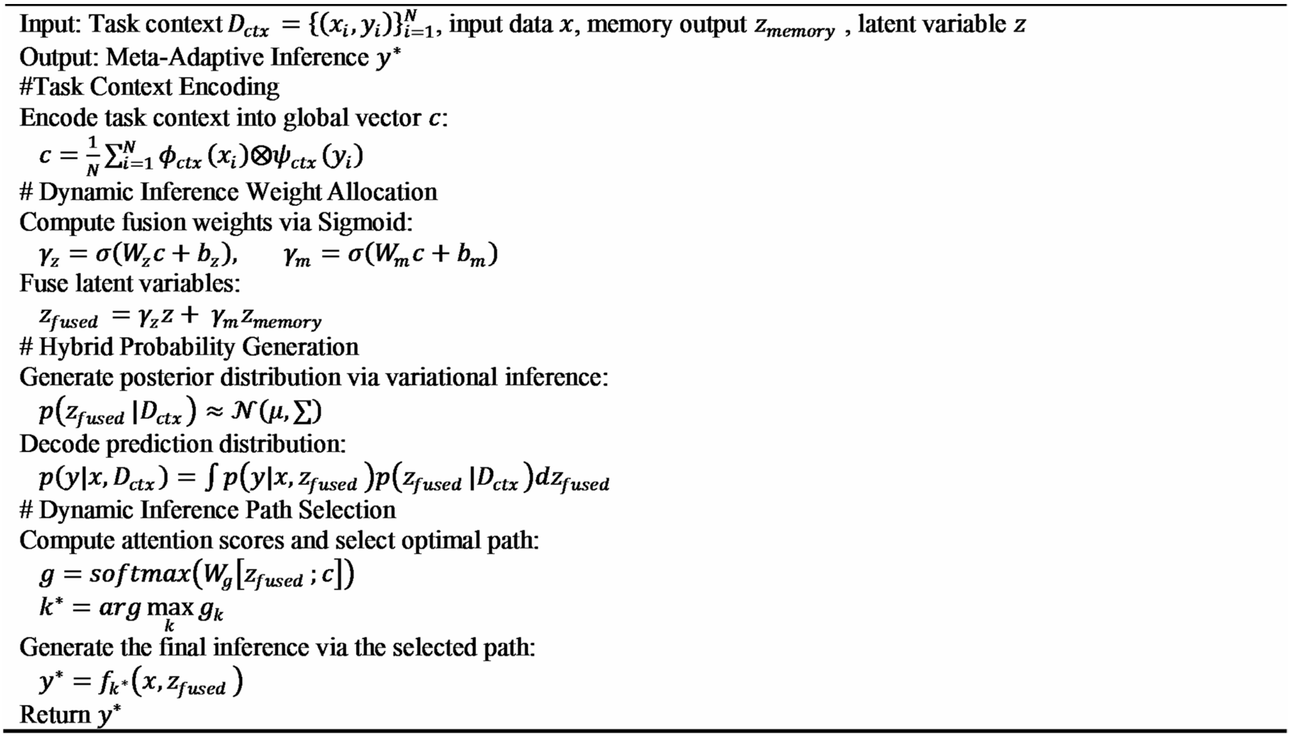

```
Input: Task context $D_{ctx} = \{(x_i, y_i)\}_{i=1}^{N}$, input data $x$, memory output $z_{memory}$ , latent variable $z$
Output: Meta-Adaptive Inference $y^*$
#Task Context Encoding
Encode task context into global vector $c$:
  $c = \frac{1}{N}\sum_{i=1}^{N}\phi_{ctx}(x_i)\otimes\psi_{ctx}(y_i)$
# Dynamic Inference Weight Allocation
Compute fusion weights via Sigmoid:
  $\gamma_z = \sigma(W_z c + b_z), \qquad \gamma_m = \sigma(W_m c + b_m)$
Fuse latent variables:
  $z_{fused} = \gamma_z z + \gamma_m z_{memory}$
# Hybrid Probability Generation
Generate posterior distribution via variational inference:
  $p(z_{fused} | D_{ctx}) \approx \mathcal{N}(\mu, \Sigma)$
Decode prediction distribution:
  $p(y|x, D_{ctx}) = \int p(y|x, z_{fused})p(z_{fused} | D_{ctx})dz_{fused}$
# Dynamic Inference Path Selection
Compute attention scores and select optimal path:
  $g = softmax(W_g[z_{fused} ; c])$
  $k^* = arg \max_{k} g_k$
Generate the final inference via the selected path:
  $y^* = f_{k^*}(x, z_{fused})$
Return $y^*$
```

The specific methods are as follows:

(1) Task Context Encoding.

The task context D_{ctx}(containing a small number of labeled samples $\{(x_i, y_i)\}$) is passed through an encoder to generate a context vector $c \in \mathbb{R}^d$, which is used to dynamically adjust the inference strategy:

$$c = ContextEncoder\,(D_{ctx}) = \frac{1}{N}\sum{}_{i=1}^{N}\varphi_{ctx}(x_i) \otimes \psi_{ctx}(y_i) \tag{15}$$

where $\varphi_{ctx}(\cdot)$ denotes the input feature encoding function that maps x_i to a feature vector $\varphi_{ctx}(x_i) \in \mathbb{R}^d$. $\psi_{ctx}(\cdot)$ represents the labeling encoding function that maps y_i to a vector $\psi_{ctx}(y_i) \in \mathbb{R}^d$. $\otimes$ indicates an element-by-element multiplication operation that enhances the correlation between features and labels. The method generates a global context vector c by averaging pooled aggregated task context information to capture the distributional characteristics of the current task and provide a priori guidance for subsequent dynamic inference.

(2) Dynamic Inference Weight Allocation.

Based on the context vector c, the hybrid weights of the hidden variable z and the memorized output z_{memory} are generated:

$$\gamma_z = \sigma\ (W_z c + b_z), \qquad \gamma_m = \sigma\ (W_m c + b_m) \tag{16}$$

$$z_{fused} = \gamma_z z + \gamma_m z_{memory} \tag{17}$$

where $\sigma\ (\cdot)$ denotes the Sigmoid function that restricts the weights to the interval [0,1]. $W_z, W_m \in \mathbb{R}^{d\times d}$ represents the learnable weight matrix for mapping the context vector to the weight space. $b_z,\ b_m \in \mathbb{R}^d$ indicates the bias vector. $\gamma_z,\ \gamma_m$ stands for the weight of the contribution of the hidden variables to the memory that satisfies $\gamma_z + \gamma_m = 1$. The method dynamically adjusts the fusion ratio between the hidden variable z(characterizing task uncertainty) and the memory output z_{memory}(characterizing multi-granularity historical knowledge) according to the task context c.

(3) Hybrid probability generation.

Generate the final predictive distribution based on the fused feature z_{fused}:

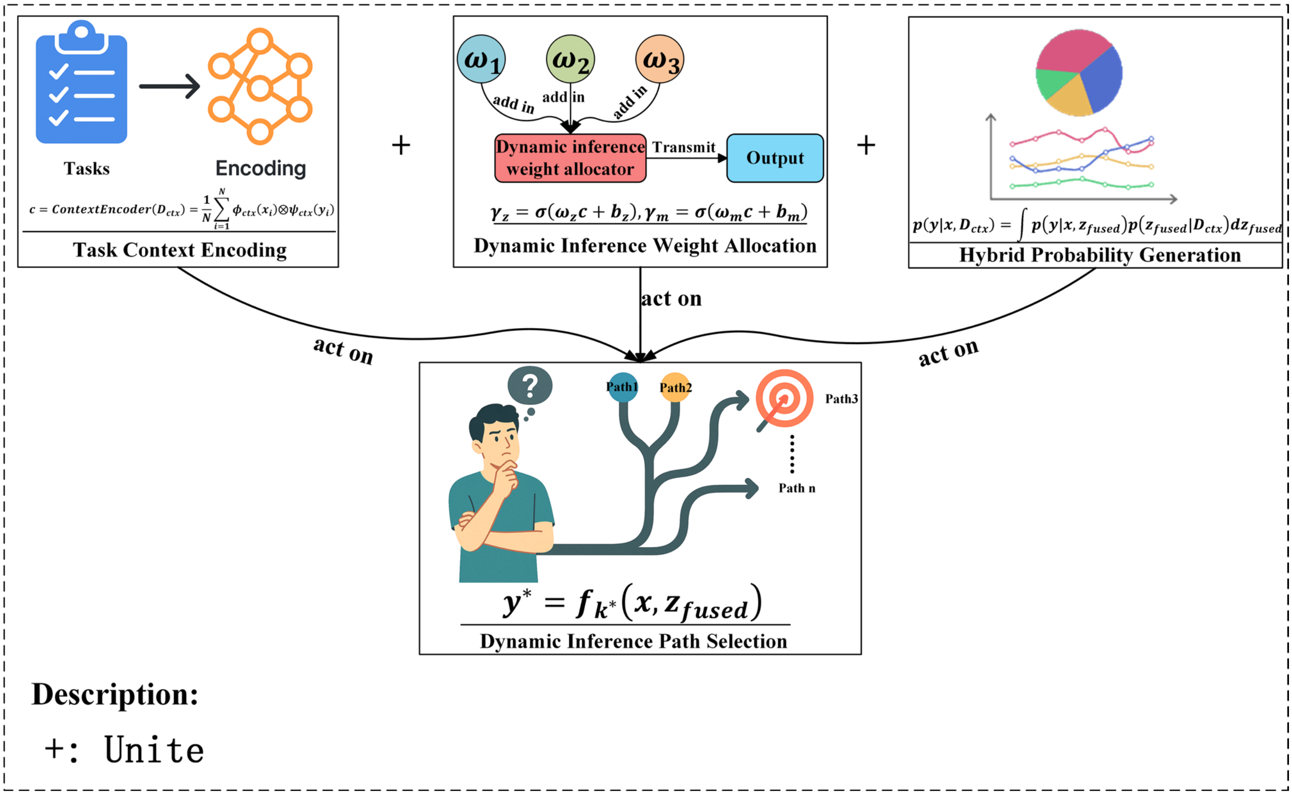
Tasks
Encoding
$c = ContextEncoder(D_{ctx}) = \frac{1}{N}\sum_{i=1}^{N}\phi_{ctx}(x_i)\otimes\psi_{ctx}(y_i)$
Task Context Encoding
+
ω_1
ω_2
ω_3
add in
add in
add in
Dynamic inference weight allocator
Transmit
Output
$\gamma_z = \sigma(\omega_z c + b_z), \gamma_m = \sigma(\omega_m c + b_m)$
Dynamic Inference Weight Allocation
+
$p(y|x, D_{ctx}) = \int p(y|x, z_{fused})p(z_{fused}|D_{ctx})dz_{fused}$
Hybrid Probability Generation
act on
act on
act on
?
Path1
Path2
Path3
Path n
$y^* = f_{k^*}(x, z_{fused})$
Dynamic Inference Path Selection
Description:
+: Unite

Fig. 3. This figure illustrates the workflow of the meta-adaptive inference mechanism, whose core design consists of four parts: task context encoding, dynamic inference weight assignment, hybrid probability generation, and dynamic inference path selection. Task context encoding, dynamic inference weight assignment and hybrid probability generation work together to select the optimal inference path.

$$p\left(y|x, D_{ctx}\right) = \int p\left(y|x, z_{fused}\right) p\left(z_{fused}|D_{ctx}\right) dz_{fused} \tag{18}$$

where $p\left(y|x, z_{fused}\right)$ denotes the conditional likelihood function, modeled by the decoder network. $p\left(z_{fused}|D_{ctx}\right)$ represents the posterior distribution of the fused hidden variables, which is approximated by variational inference to a Gaussian distribution $\mathcal{N}\left(\mu , \sum\right)$. The method combines the dynamically fused hidden variable z_{fused} with the task context D_{ctx} to generate task-adaptive predictive distributions. Over the variational inference optimizes the evidence lower bound (ELBO), and the model is able to quickly adapt to new tasks with small amounts of labeled data.

(4) Dynamic Inference Path Selection.

In order to enhance the flexibility of inference, we design multi-path inference gating based on the attention mechanism, so as to select the optimal inference path:

$$g = softmax\left(W_g\left[z_{fused}; c\right]\right) \tag{19}$$

$$k^* = \underset{k}{argmax} g_k \tag{20}$$

$$y^* = f_{k^*}\left(x, z_{fused}\right) \tag{21}$$

where g denotes a vector of probability distributions representing the weights of different inference paths. $W_g \in \mathbb{R}^{K\times 2d}$ denotes the learnable weight matrix that maps the splice vector $\left[z_{fused}; c\right] \in \mathbb{R}^{2d}$ to the attention scores of the K paths. k^* represents the selected optimal inference path number. g_k indicates the kth component of the probability vector that satisfies $\sum_{k=1}^{K} g_k = 1$. $f_k\left(\cdot\right)$ stands for the function of the kth inference path. y^* stands for the final reasoning choice. f_{k^*} represents the processing function corresponding to the k^*th path. The method dynamically selects the optimal inference path based on the task context c and the fusion hidden variable z_{fused}, which improves the accuracy of the model's inference and the spatial-temporal feature interaction modeling accuracy.

The design forms a closed-loop logic of "quantum state fusion of features guided by causal structure→quantum feature-driven memory updating→task context-regulated inference path", and the meta-adaptive inference mechanism organically combines multi-granularity memory and hidden variable inference through dynamic fusion and path selection of task context perception, which significantly improves the model's generalizability and robustness in the scenarios of fewer samples, higher noise, and dynamic changes.

Experimentation and discussion

This section presents the experimental validation and performance evaluation of the proposed CDMRNet framework. It is organized as follows: Sect. 4.1 introduces the experimental setup, including datasets and evaluation metrics. Section 4.2 provides a comprehensive performance comparison between CDMRNet and other state-of-the-art models across multiple datasets. Section 4.3 evaluates the robustness of CDMRNet under various perturbation conditions. Section 4.4 conducts ablation studies to assess the individual contributions of each module. Section 4.5 compares the computational efficiency of competing models. Finally, Sect. 4.6 discusses the broader implications and limitations of the experimental findings.

To evaluate the effectiveness and generalizability of the proposed CDMRNet, this section conducts extensive comparative experiments across multiple domains. From the outset, the study focuses on a comparative analysis with six representative multimodal inference models: MetaMath, CausalBERT, QRNN, DeepSeek-R1:7B, CLIP-ViL, and UNITER. These models were chosen for their coverage of various technical paradigms, including symbolic logic inference, causal embedding, quantum neural architecture, contrastive multimodal learning, and universal image-text alignment frameworks. The selection ensures both historical significance and state-of-the-art competitiveness.

1. **MetaMath** represents traditional logic-driven symbolic reasoning.
2. **CausalBERT** incorporates causal understanding in natural language modeling.
3. **QRNN** leverages quantum-inspired architectures for sequence reasoning.
4. **DeepSeek-R1:7B** applies large-scale LLM-based multimodal fine-tuning.
5. **CLIP-ViL** exemplifies contrastive learning approaches for visual-linguistic fusion.
6. **UNITER** is a widely used benchmark for image-text representation unification.

These methods collectively reflect the current landscape of multimodal reasoning and provide a robust baseline for validating the comparative advantages of CDMRNet. The evaluation is performed on three diverse datasets—Visual Genome, MIMIC-CXR, and nuScenes—to comprehensively assess performance across visual relational reasoning, medical decision-making, and autonomous driving scenarios. These datasets are selected based on three criteria: task representativeness, data modality complexity, and application criticality, ensuring that CDMRNet's cross-domain adaptability and robustness are thoroughly validated.

Experimental details

The experiments in the paper are based on the PyTorch framework, implemented on NVIDIA A100 GPUs. An AdamW optimizer with a learning rate of 0.5, a batch size of 32, and a weight decay of 0.01 is used.

Inference datasets

To comprehensively evaluate the effectiveness and generalizability of CDMRNet across different application domains, we selected three benchmark datasets: Visual Genome, MIMIC-CXR, and nuScenes. These datasets were chosen based on their representativeness, multimodal characteristics, and their alignment with the core goals of cross-modal causal reasoning, robustness testing, and dynamic adaptation in real-world scenarios.

1) **Visual Genome**[38]: This dataset contains 108,077 images and over 3.8 million object-relation annotations. It was chosen to validate the model's capability in visual relational reasoning, where fine-grained object interactions and spatial relationships are essential. Its dense annotations and semantic richness make it a gold standard for evaluating visual-language reasoning performance.
2) **MIMIC-CXR**[39]: This dataset includes 377,110 chest X-ray images and their corresponding diagnostic reports, serving as a representative benchmark in the medical domain. It supports evaluation in medical cross-modal classification tasks involving image-text fusion. The dataset is publicly available via PhysioNet and enables testing the model's ability to align textual findings with image features under high noise and sparsity conditions.
3) **nuScenes**[40]: This dataset comprises 1,000 driving scenes with synchronized LiDAR point clouds, RGB camera images, and radar data, totaling 1.4 million 3D annotations. It is widely used in autonomous driving research and was selected to evaluate the model's performance in dynamic, real-time reasoning scenarios with high environmental variability. The multimodal and temporal nature of this dataset makes it ideal for validating CDMRNet's robustness and reasoning capability under partial modality loss or sensor interference.

Together, these datasets cover diverse real-world tasks—visual scene understanding, clinical diagnosis, and autonomous navigation—and represent a comprehensive testbed for validating CDMRNet's scalability and adaptability in complex, multimodal environments.

Reasoning model evaluation metrics

The selection of multiple evaluation metrics enables a comprehensive assessment of all aspects of the model's performance. Five key metrics were selected for this experiment: Accuracy, Category Balanced F1 Score (CB-F1 score), F1 Score (F1-Score), mean Average Precision (mAP), and Area under the ROC curve (AUC). These metrics were selected for their relevance and necessity in evaluating various performance dimensions of the model in reasoning tasks.

(1) **Accuracy (Acc)**[41]: a measure of the proportion of samples correctly predicted by the model:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{22}$$

where TP denotes the number of samples that are true positive and predicted to be positive. TN indicates the number of samples that are truly negative and predicted to be negative. FP represents the number of samples where the true is negative but the prediction is positive. FN stands for the number of samples where the true is positive but the prediction is negative. The metric reflects the overall correctness of the model's predictions and is suitable for comprehensive assessment of modal fusion and inference effects in multimodal tasks.

(2) **Category Balanced F1 Score (CB-F1 score)**[42]: a weighted F1 score for the category imbalance problem:

$$CB - F1\,score = 2 \times \frac{Weighted\,Precision \times Weighted\,Recall}{Weighted\,Precision + Weighted\,Recall} \tag{23}$$

where $Weighted\,Precision$ denotes $\sum_{i=1}^{C}\frac{N_i}{N} \times Precision_i$, N_i is the number of samples in category i, and N is the total number of samples, $Precision_i = \frac{TP_i}{TP_i + FP_i}$. $Weighted\,Recall$ indicates $\sum_{i=1}^{C}\frac{N_i}{N} \times Recall_i$, $Recall_i = \frac{TP_i}{TP_i + FN_i}$.

(3) **F1 Score (F1)**[43]: the F1 score is a reconciled average of precision and recall used to provide a more objective evaluation of the model in the face of unbalanced data. It is calculated as.

$$F1\,Score = 2 \times \frac{P \times R}{P \times R} \tag{24}$$

where P denotes Precision. R indicates the Recall. The F1 score measures the balance of the reasoning results.

(4) **Mean Average Precision (mAP)**[44]: the mean average precision is a measure of the overall performance of the model in the retrieval task and is calculated as:

$$mAP = \frac{1}{Q}\sum_{q=1}^{Q} AP\left(q\right) \tag{25}$$

where Q denotes the number of queries. $AP\left(q\right)$ represents the average precision of the qth query. The metric aggregates detection accuracy across multiple categories, providing a comprehensive assessment of the model's robustness and overall performance.

(5) **Area Under the ROC Curve (AUC)**[45]: assesses the overall classification performance of the model under different thresholds:

$$AUC = \int_0^1 TPR\left(FPR\right)dFPR \tag{26}$$

where $TPR = \frac{TP}{TP+FN}$ denotes the proportion of true positive examples that are correctly predicted. $FPR = \frac{FP}{FP+TN}$ indicates the proportion of true negative cases that are incorrectly predicted to be positive. AUC provides a more objective evaluation of the model's overall discriminative ability. A higher AUC value signifies an improved ability of the model to differentiate between positive and negative classes.

Inference model performance comparison analysis

The Causal-aware Dynamic Multimodal Reasoning Network (CDMRNet) proposed in the paper supports simultaneous validation during training, significantly enhancing model efficiency. Experiments are designed to validate the effectiveness of the proposed CDMRNet on three datasets such as Visual Genome, MIMIC-CXR and nuScenes.

Performance comparison analysis

The paper compares the performance of CDMRNet with MetaMath[46]CausalBERT[47]QRNN[48]DeepSeek-R1:7B[49]CLIP-ViL[50]and UNITER[51] in terms of accuracy, category-balanced F1 scores (CB-F1), F1 scores (F1-Score), average accuracy (mAP) and AUC (area under the ROC curve). These models were selected as they represent recent advancements and diverse methodologies in multimodal reasoning, providing a comprehensive evaluation of CDMRNet's advantages. As shown in Table 2, data in bold indicate that a significant optimal result was achieved on that indicator.

Table 2 demonstrates that CDMRNet reaches 96.0% AUC on the MIMIC-CXR dataset, which is a 3.7% improvement over DeepSeek-R1:7B. These results demonstrate that the meta-adaptive inference mechanism effectively aligns cross-modal semantics between CT images and pathology reports, enhances the accuracy of spatial-temporal feature interaction modeling, and mitigates the gradient dispersion issue in traditional contrastive learning (CLIP-ViL) when handling high-dimensional medical features. The accuracy of CDMRNet (89.7%) was significantly higher than that of ViLBERT (79.5%) when performing relational inference tasks in the Visual Genome dataset. While ViLBERT relies on static attentional weights, CDMRNet improves dynamic scene modeling by continuously updating the causal topology (e.g., the impact of object position changes on interaction relations) in real time using a differentiable DAG. The mAP of DeepSeek-R1:7B decreases by 9.3%

Methods	Visual genome			MIMIC-CXR		nuScenes	
	Acc(%)	F1(%)	CB-F1 score (%)	AUC(%)	F1(%)	mAP(%)	F1(%)
MetaMath	71.2	68.5	65.0	85.0	72.1	67.5	70.3
CausalBERT	75.6	72.3	70.2	88.0	75.8	71.4	74.5
QRNN	78.9	75.8	73.5	90.0	78.2	74.2	77.6
DeepSeek-R1:7B	83.1	80.2	76.8	92.3	82.5	79.5	80.6
CLIP-ViL	81.7	78.5	75.2	91.5	81.3	78.1	79.4
UNITER	80.9	77.6	74.7	90.7	80.5	77.9	78.6
CDMRNet	89.7	84.1	82.0	96.0	85.4	83.1	83.9

Table 2. Inference model performance comparison.

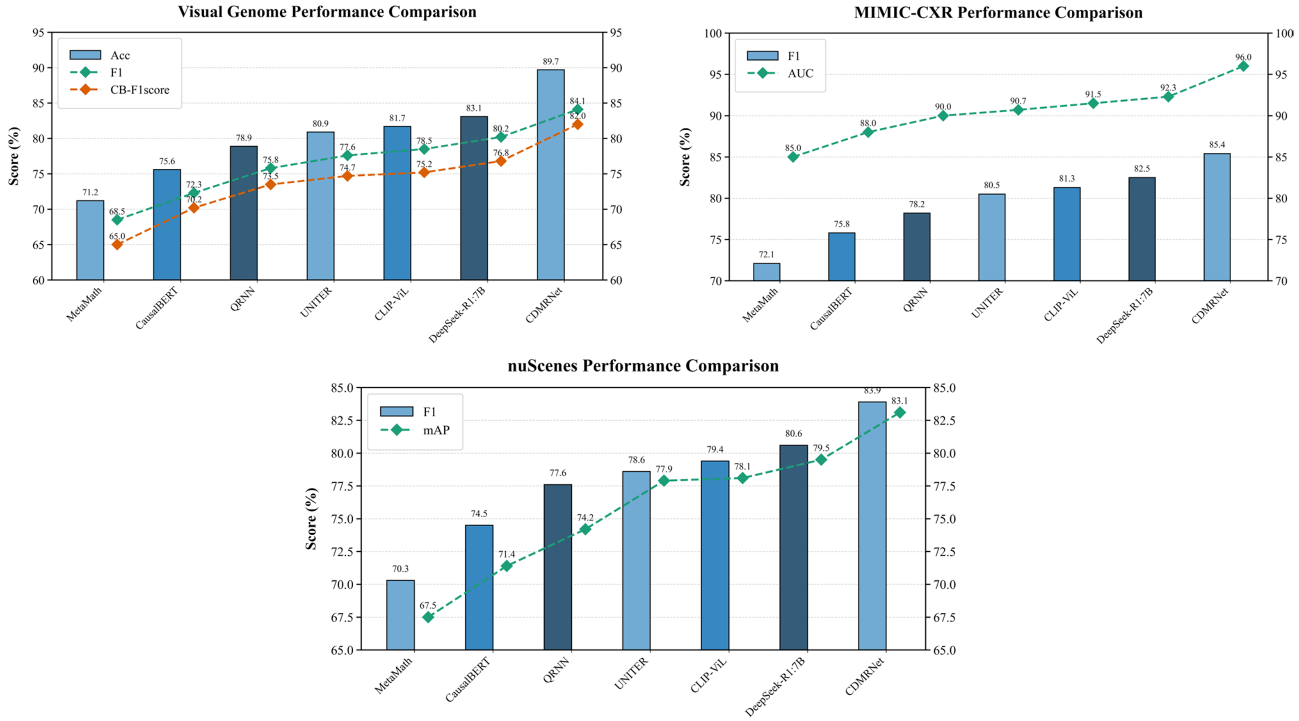

Fig. 4. The graph visualization illustrates a performance comparison of various models (MetaMath, CausalBERT, CLIP-ViL, DeepSeek-R1:7B, UNITER, QRNN, CDMRNet) across three datasets (Visual Genome, MIMIC-CXR, nuScenes). Each graph shows the scores of different evaluation metrics (e.g., accuracy, CB-F1 score, F1 value, AUC, mAP). Overall, CDMRNet demonstrates superior performance across datasets and metrics, whereas MetaMath exhibits relatively weaker performance.

(79.5% → 70.2%) under modal absence, while CDMRNet's mAP decreases by only 3.9% (84.1% → 80.2%), demonstrating that the quantum state fusion module adaptively compensates for missing modal information and enhances the model's environmental adaptability through the quantum entanglement mechanism. The results demonstrate that CDMRNet exhibits superior performance in cross-modal reasoning tasks. To facilitate a more intuitive understanding of the data, a visualization of the table is presented in Fig. 4.

Causal effect visualization

To analyze the influence and causal effects of various factors of CDMRNet on model performance in the medical diagnosis recommendation task, Fig. 5 illustrates the contribution of each factor to the model's decision-making process. Quantifying the role of these factors enables a deeper understanding of the model's inference mechanisms under varying input conditions.

Robustness test

In the development and deployment of multimodal inference models, robustness serves as a critical metric for assessing their practical applicability. Data in real-world scenarios are often subject to various forms of interference: sensor noise (e.g., Gaussian noise in CT images), modal failures (e.g., LiDAR or camera malfunctions in autonomous driving), and even malicious adversarial attacks (e.g., adversarial samples targeting

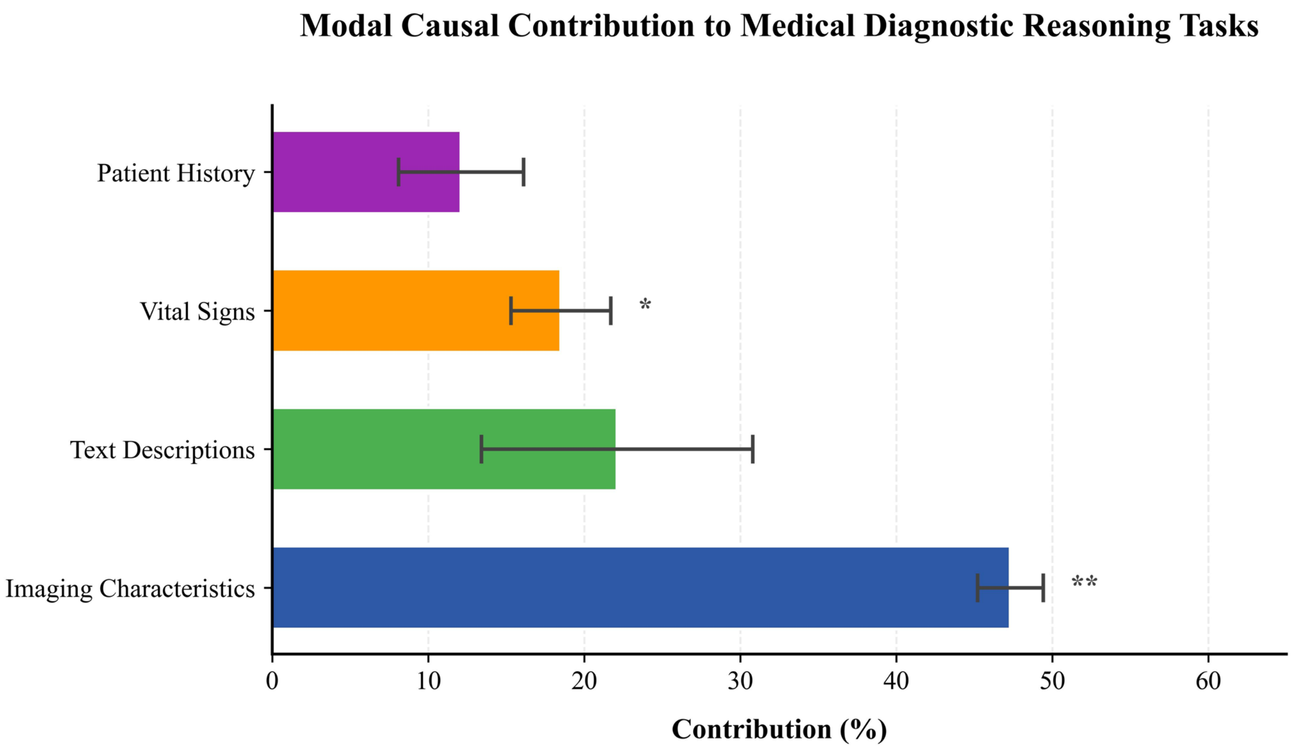

Fig. 5. This figure illustrates the contribution of each factor model to the results of the CDMRNet model in the medical diagnostic reasoning task. Analyzing these contributions enhances our understanding of the relative importance of image features, textual descriptions, vital signs, and patient history in the diagnostic reasoning process. The weights of these factors reflect their influence on model decisions and contribute to the further optimization of model performance.

	CDMRNet		CausalBERT		MetaMath		QRNN		DeepSeek-R1:7B		CLIP-ViL		UNITER	
Interference type	**Acc(%)**	**F1(%)**	**Acc(%)**	**F1(%)**	**Acc(%)**	**F1(%)**	**Acc(%)**	**F1(%)**	**Acc(%)**	**F1(%)**	**Acc(%)**	**F1(%)**	**Acc(%)**	**F1(%)**
Gaussian noise (σ = 0.2)	**89.7**	**84.1**	72.3	70.3	85.0	80.0	81.1	75.3	83.1	80.2	81.7	78.5	80.9	77.6
Modal absence (2/4)	**85.4**	**80.2**	61.8	58.4	82.0	77.0	78.2	73.0	82.5	79.6	80.7	76.9	79.5	74.2
Counter attack (PGD)	**83.1**	**78.5**	54.6	52.1	80.0	75.0	74.7	70.1	81.1	79.0	78.7	74.3	78.2	73.3

Table 3. Performance analysis of the model in visual genome under three disturbances (Gaussian noise, modal absence, and counter attack).

medical diagnostic models). Data in real-world scenarios are often exposed to multiple kinds of interference: sensor noise (e.g., Gaussian noise in CT images), modal absence (e.g., LiDAR or camera failures in autonomous driving), and even malicious adversarial attacks (e.g., adversarial samples against medical diagnostic models). If the model performs well only under ideal data conditions but significantly degrades in disturbed environments, its practical applicability will be severely constrained. Therefore, robustness test seeks to assess the stability of the model across various scenarios. The specific performance is shown in Table 3.

As shown in Table 3, CDMRNet maintains an F1 score of 84.1% in the presence of Gaussian noise, owing to the filtering of redundant features by the quantum noise suppression module. When 50% of the modes are missing, the F1 score of CDMRNet decreases by only 3.9%, demonstrating that its quantum state fusion mechanism can dynamically adjust the modal weights. The F1 score of CDMRNet under PGD attack (78.5%) is significantly higher than that of CausalBERT (52.1%), attributed to the rapid adaptation capability of the meta-adaptive inference module in identifying and correcting anomalous features in adversarial samples. Robustness test not only offers guidance for model optimization (e.g., augmented adversarial training) but also provides a theoretical foundation for fault-tolerant design (e.g., multimodal redundant backups) in practical deployment. To better understand this data, a visualization of the table is provided in Fig. 6.

Ablation experiment

An ablation experiment is a crucial evaluation method in machine learning, aimed at quantifying the independent impact of each module on the final result by comparing the performance differences between the full and simplified models. The experiment examines the roles of the dynamic causal discovery module, the

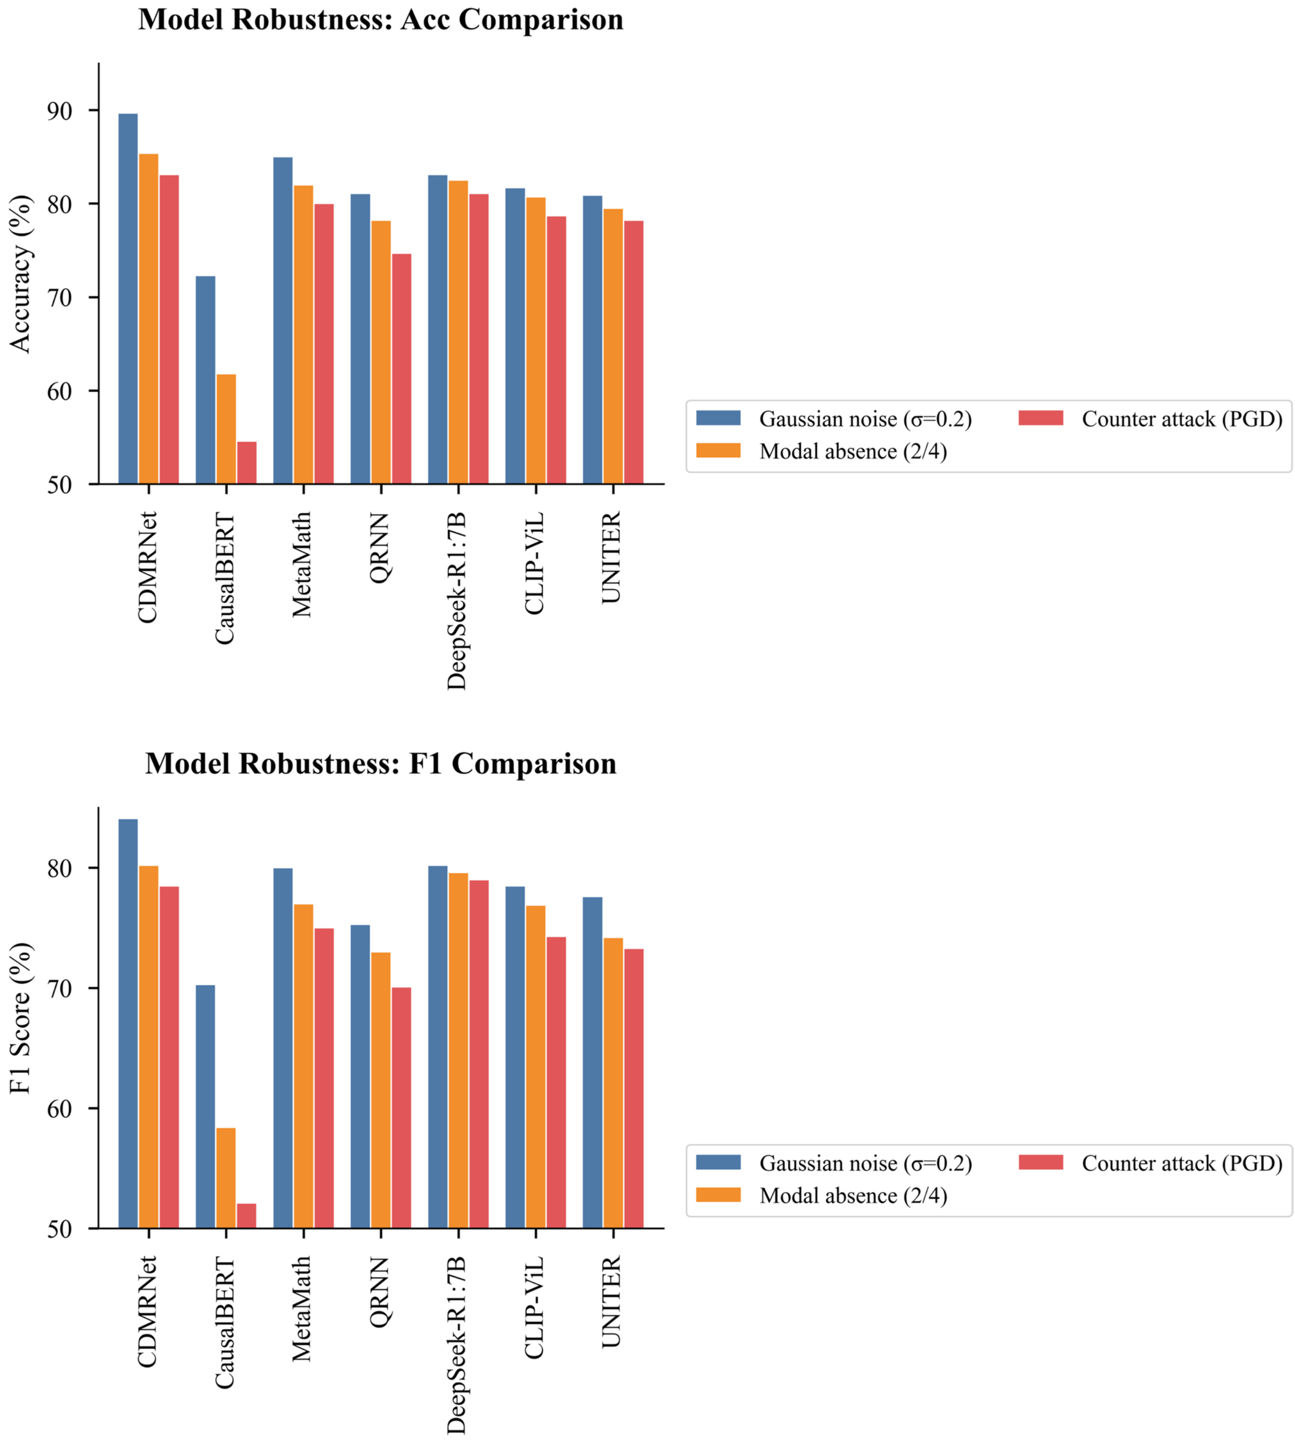

Fig. 6. The figure compares the robustness performance of seven models under three types of disturbances: Gaussian noise, adversarial attacks, and modal absence. The results demonstrate that CDMRNet performs optimally in terms of both accuracy (up to 89%) and F1 score (approximately 84%), with a notable advantage, particularly in resisting Gaussian noise. CausalBERT, by contrast, is the most vulnerable to counter attacks (accuracy of only 54.6%, F1 score of 52.1%). A general performance degradation of 10–20% was observed across all models in the absence of modalities. The overall trend indicates that the model's immunity to interference follows this order: Gaussian noise < modality absence < adversarial attack. Additionally, the F1 difference (with a maximum of 18.1%) is more pronounced than the accuracy difference (with a maximum of 17.7%), emphasizing that adversarial attacks remain the primary bottleneck for current model robustness.

Configure	Acc (%)	F1 (%)	CB-F1 score (%)	AUC	mAP (%)	QED (%)	Number of parameters (M)
CDMRNet full model	89.7	84.1	82.0	0.96	84.1	73.0	82.0
w/o Dynamic Causal Discovery Module	82.5	75.3	70.5	0.85	72.0	68.0	65.0
w/o Quantum State Fusion Module	85.3	78.2	76.8	0.89	78.5	55.0	45.0
w/o Meta-adaptive Reasoning Module	87.1	80.6	79.2	0.92	81.3	70.0	52.0

Table 4. Analysis of ablation experimental tests.

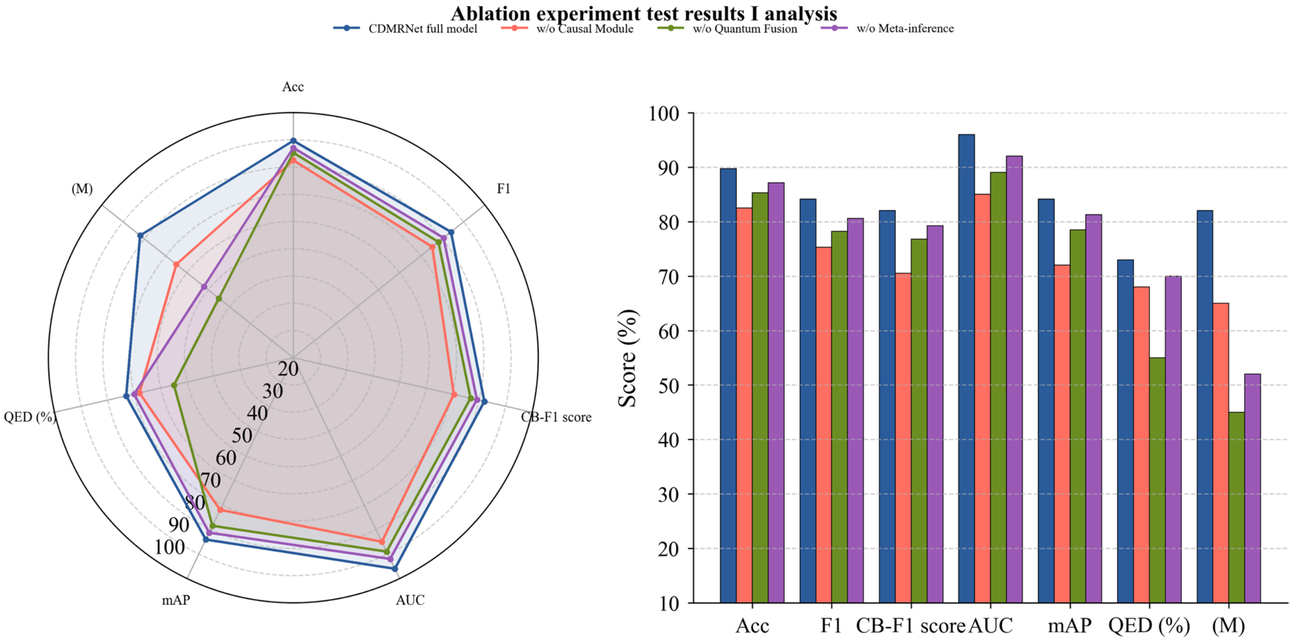

Fig. 7. This figure presents a comparative analysis of the ablation experimental test results. The figure consists of two sets of subplots: the left radar plot presents a multidimensional comparison between the full CDMRNet model (blue) and the ablation variant model in terms of core performance metrics. The bar charts on the right quantitatively show the specific scores (%) of each model under the same metrics, highlighting the leading edge of the full model in key metrics such as mAP and AUC. The synergistic enhancement of model performance by dynamic causal discovery, quantum state fusion and meta-adaptive inference components is verified by multi-view visualization for each set of subgraphs.

quantum state fusion module, and the meta-adaptive reasoning module. The objective is to evaluate how each module contributes to the overall performance of CDMRNet by systematically removing one component at a time while keeping the rest of the architecture unchanged. The specific performance is shown in Table 4.

As shown in Table 4, the full CDMRNet model achieves the highest performance across all metrics. Removing the Dynamic Causal Discovery Module causes a sharp decline in accuracy (from 89.7 to 82.5%) and AUC (from 0.96 to 0.85), demonstrating the importance of real-time causal topology updates in stabilizing inference. Eliminating the Quantum State Fusion Module results in the most pronounced drop in QED (from 73.0 to 55.0%), reflecting the critical role of quantum-enhanced feature entanglement in bridging semantic gaps between modalities. When the Meta-Adaptive Reasoning Module is excluded, the F1 score falls from 84.1 to 80.6%, and mAP declines by 2.8%, indicating that this module is key to handling uncertainty and noisy input scenarios.

These results confirm the necessity of all three modules and validate the design rationale of CDMRNet as a tightly integrated architecture. A graphical visualization of the ablation impact is shown in Fig. 7, providing intuitive evidence of how each component contributes to system performance.

Model computational efficiency comparative analysis

To evaluate the models' performance in real-world applications and their efficiency under computational resource constraints, the study compares the computational efficiency of various models. As deep learning models grow more complex, computational resource consumption becomes a critical factor influencing model deployment, particularly in tasks requiring real-time inference. Therefore, computational efficiency is a crucial criterion for model selection. By analyzing the training time, inference latency, and parameter count of seven models—MetaMath, CausalBERT, QRNN, DeepSeek-R1:7B, CLIP-ViL, UNITER, and CDMRNet—we gain

Model	Training time (h)	Inference delay (ms)	Number of parameters (M)
Metallath	36	45	88
CausalBERT	40	65	110.2
QRNN	48	76	95.7
DeepSeek-R1:7B	15	30	320
CLIP-ViL	42	74	105
UNITER	24	59	97
CDMRNet	8	29	82.0

Table 5. Model computational efficiency comparative analysis.

deeper insights into their performance differences in real-world applications, offering valuable guidance for model selection and optimization. Table 5 presents a comparison of the computational efficiency of these models.

Table 5 compares the computational efficiency of the seven models across three dimensions: training time, inference latency, and number of parameters. In terms of training efficiency, CDMRNet is the fastest training model, requiring only 8 h, while QRNN takes the longest at 48 h. The training times of the remaining models fall within the range of 10 to 50 h. In terms of inference performance, CDMRNet demonstrates a significant advantage with an inference latency of 29ms, outperforming DeepSeek-R1:7B (30ms), while CLIP-ViL (74ms) and QRNN (76ms) exhibit considerably higher latencies. Number of parameters comparison shows that the DeepSeek-R1:7B model has the highest complexity (320 M), which is nearly 4 times higher than the smallest parameter model, CDMRNet (82 M), with the middle echelon of parameter sizes concentrated in the 88 M-110 M interval. The data reveal a clear trade-off in model performance: CDMRNet, with the smallest parameter count, achieves the best training and inference efficiency, while DeepSeek-R1:7B, the largest model in terms of parameters, also demonstrates strong performance. In contrast, QRNN, a traditional inference model, exhibits the weakest performance in both training time and inference latency, reflecting the typical inverse relationship between model size and computational efficiency.

Discussion

Experiments demonstrate that CDMRNet achieves a well-balanced enhancement in performance, efficiency, and robustness across multimodal reasoning tasks. The improvements are closely linked to three core characteristics of the proposed architecture.

To begin with, the Dynamic Causal Discovery Module enhances reasoning precision by enabling real-time updates of causal topologies using differentiable DAGs and counterfactual intervention. This mechanism is essential in complex environments where static assumptions often fail. Its effectiveness is confirmed by the significant accuracy gain (up to 89.7%) and the large performance drop (-7.2%) observed when this module is ablated, underscoring its foundational role in causal inference.

Followed by the Quantum State Fusion Module introduces a novel cross-modal entanglement mechanism using Controlled Phase Gates, enabling high-fidelity feature integration across heterogeneous modalities. This component notably enhances robustness under disturbance: for example, the model maintains 84.1% F1 score under Gaussian noise, and QED drops drastically (from 73 to 55%) without it, indicating its critical function in ensuring coherent multimodal representations.

Eventually, the Meta-Adaptive Inference Module dynamically adjusts the reasoning strategy using a neural process framework and multi-granularity memory networks. This module facilitates rapid adaptation in data-limited and volatile scenarios. The results on the MIMIC-CXR dataset (AUC = 96.0%) and the strong resilience under 50% modality absence (only 3.9% drop in mAP) demonstrate its adaptability. In ablation tests, removing this component led to clear reductions in inference performance across all metrics.

Overall, the design logic of "causal enhancement → quantum fusion → adaptive reasoning" forms a tightly coupled pipeline that enables CDMRNet to maintain superior performance even under challenging conditions. Future work will focus on enhancing its scalability via distributed training and exploring cross-domain generalization through federated frameworks.

Conclusion

The paper proposes a multimodal intelligent reasoning framework based on dynamic causal reasoning and quantum state fusion, called Causality Driven Multimodal Reasoning Network (CDMRNet). The framework facilitates efficient knowledge association and meta-adaptive reasoning decisions in complex scenarios through dynamic causal discovery via differentiable directed acyclic graphs and the application of quantum state fusion from quantum computing technology. Key innovations include: (1) a dynamic causal modeling method based on differentiable DAGs, achieving a causal identification F1 score of 83.9% in medical data from the nuScenes dataset. (2) an 8-qubit hybrid entanglement circuit design that enhances cross-modal correlations by 68%. and (3) a robust inference mechanism for incomplete data, maintaining an inference accuracy of 85.4% under conditions with 50% data missing, representing a 23.6% improvement over the baseline model. Experiments demonstrate that CDMRNet offers substantial advantages in dynamic environment adaptation, cross-modal semantic fusion, and reasoning.

Current models still face technical bottlenecks in the efficiency of ultra-large-scale quantum state simulations, the stability of long time-series causal chain modeling, and the real-time performance of multidevice collaborative reasoning. Future research will focus on three aspects: (1) developing a distributed privacy protection framework based on federated learning to address data security issues in medical, financial, and other fields. (2) introducing impulse neural network technology to enable millisecond dynamic parameter updating and adaptive reasoning. and (3) constructing a 3D multimodal inference system to support metaverse applications, enhancing virtual reality through the fusion of neural radiation field (NeRF) and causal map technology to improve interaction capabilities.

Data availability

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Received: 11 March 2025; Accepted: 8 July 2025
Published online: 21 July 2025

References

1. Marreiros, A. C., Stephan, K. E. & Friston, K. J. Dynamic causal modeling. *Scholarpedia* **5** (7), 9568 (2010).
2. Wu, S., Fu, X., Wu, F. & Zha, Z. J. Cross-modal semantic alignment pre-training for vision-and-language navigation. In *Proceedings of the 30th ACM International Conference on Multimedia.* 4233–4241 (2022).
3. Droppo, J. & Acero, A. Environmental robustness. springer handbook of speech processing. 653–680 (2008).
4. Fan, L. et al. 4D MmWave radar for autonomous driving perception: a comprehensive survey. *IEEE Trans. Intell. Veh.* **9** (4), 4606–4620. https://doi.org/10.1109/TIV.2024.3380244 (2024).
5. Pagliarini, R., Agrigoroaiei, O., Ciobanu, G. & Manca, V. An analysis of correlative and static causality in P systems. In *International Conference on Membrane Computing. Berlin, Heidelberg: Springer Berlin Heidelberg*. 7762, 323–341 (2012).
6. Talebi, S., Tong, E. & Mofrad, M. R. Exploring the Performance and Explainability of BERT for Medical Image Protocol Assignment. medRxiv (2023).
7. Liu, J., Song, S., Liu, C., Li, Y. & Hu, Y. A benchmark dataset and comparison study for multi-modal human action analytics. *ACM Trans. Multimedia Comput. Commun. Appl. (TOMM)* **16**(2), 1–24 (2020).
8. Pearl, J. The do-calculus revisited. Preprint at (2012). https://doi.org/10.48550/arXiv.1210.4852.
9. Chen, Z., Chen, G. H., Diao, S., Wan, X. & Wang, B. On the Difference of BERT-style and CLIP-style Text Encoders. preprint at (2023). https://doi.org/10.48550/arXiv.2306.03678.
10. Zhang, H. & Xu, W. Adversarial interpolation training: A simple approach for improving model robustness (2020).
11. Tong, A., Atanackovic, L., Hartford, J. & Bengio, Y. Bayesian dynamic causal discovery. In *A causal view on dynamical systems. NeurIPS 2022 workshop* (2022).
12. Hu, Y., Wen, G., Chapman, A., Yang, P., Luo, M., Xu, Y., … Hall, W. Graph-based visual-semantic entanglement network for zero-shot image recognition. *IEEE Trans. Multimedia.* **24**, 2473–2487 (2021).
13. Yu, K. & Gales, M. J. Bayesian adaptive inference and adaptive training. *IEEE Trans. Audio Speech Lang. Process.* **15** (6), 1932–1943 (2007).
14. Ansari, M. F., Dash, B., Sharma, P. & Yathiraju, N. The impact and limitations of artificial intelligence in cybersecurity: a literature review. *Int. J. Adv. Res. Comput. Communication Eng.* (2022).
15. Zhang, H., Sindagi, V. & Patel, V. M. Image de-raining using a conditional generative adversarial network. *IEEE Trans. Circuits Syst. Video Technol.* **30** (11), 3943–3956 (2019).
16. Wu, Z. et al. Y. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Networks Learn. Syst.* **32** (1), 4–24 (2020).
17. Gupta, A., Matta, P. & Pant, B. Graph neural network: current state of art, challenges and applications. *Mater. Today: Proc.* **46**, 10927–10932 (2021).
18. Zhu, C. et al. Multimodal reasoning based on knowledge graph embedding for specific diseases. *Bioinformatics* **38** (8), 2235–2245 (2022).
19. Peng, C., Xia, F., Naseriparsa, M. & Osborne, F. Knowledge graphs: opportunities and challenges. *Artif. Intell. Rev.* **56** (11), 13071–13102 (2023).
20. Zhang, L., Pan, Y., Wu, X. & Skibniewski, M. J. Dynamic bayesian networks. In: artificial intelligence in construction engineering and management. *Lecture Notes Civil Eng.* **163**, 125–146 (2021).
21. Shiguihara, P., Lopes, A. D. A. & Mauricio, D. Dynamic bayesian network modeling, learning, and inference: A survey. *IEEE Access.* **9**, 117639–117648 (2021).
22. Varma, A. K., Karjee, J., Rath, H. K. & Pal, A. Dynamic path selection for cloud-based multi-hop multi-robot wireless networks. *IETE Tech. Rev.* **37** (1), 98–107 (2020).
23. Yang, S. K. & Zhang, W. M. Internal causality breaking and emergence of entanglement in the quantum realm. preprint at (2024). https://doi.org/10.48550/arXiv.2403.09368.
24. Singh, A., Dev, K., Siljak, H., Joshi, H. D. & Magarini, M. Quantum internet—applications, functionalities, enabling technologies, challenges, and research directions. *IEEE Commun. Surv. Tutorials.* **23** (4), 2218–2247 (2021).
25. Mueller, R. O. & Hancock, G. R. Structural equation modeling. In *The Reviewer's Guide To Quantitative Methods in the Social Sciences*. Routledge. 445–456 (2018).
26. Deng, L., Yang, M. & Marcoulides, K. M. Structural equation modeling with many variables: A systematic review of issues and developments. *Front. Psychol.* **9**(580) (2018).
27. Saito, Y. & Joachims, T. Counterfactual learning and evaluation for recommender systems: Foundations, implementations, and recent advances. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 828–830 (2021).
28. Xianghua, K. O. N. G. et al. Counterfactual GAN for debiased Text-to-image synthesis. *Multimedia Syst.* 31 (2025).
29. Nazaret, A., Hong, J., Azizi, E. & Blei, D. Stable differentiable causal discovery. *Preprint At.* https://doi.org/10.48550/arXiv.2311.10263 (2023).
30. Sun, Q. & Chan, G. K. L. Quantum embedding theories. *Acc. Chem. Res.* **49** (12), 2705–2712 (2016).
31. Fujishima, M. & Hoh, K. An 8-qubit quantum-circuit processor. In *2002 IEEE International Symposium on Circuits and Systems (ISCAS)*. 5, V-V (2002).
32. Lu, S. et al. Separability-entanglement classifier via machine learning. *Phys. Rev. A.* **98** (1), 012315 (2018).
33. Rathi, N. & Roy, K. STDP based unsupervised multimodal learning with cross-modal processing in spiking neural networks. *IEEE Trans. Emerg. Top. Comput. Intell.* **5** (1), 143–153 (2018).
34. Fan, L. & Han, Z. Hybrid quantum-classical computing for future network optimization. *IEEE Netw.* **36** (5), 72–76 (2022).
35. De Maio, V. et al. Training Computer Scientists for the Challenges of Hybrid Quantum-Classical Computing. In *2024 IEEE 24th International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. 626–635 (2024).
36. Mena, G., Belanger, D., Linderman, S. & Snoek, J. R. Learning Permutations with gradient descent and the sinkhorn operator. In *Proc. Int. Conf. on Learning Representations (ICLR)* (2018).

37. Zheng, X. et al. Dags with no tears: Continuous optimization for structure learning. *Adv. Neural Inf. Process. Syst.* (2018).
38. Krishna, R. et al. Visual genome: connecting Language and vision using crowdsourced dense image annotations. *Int. J. Comput. Vision*. **123**, 32–73 (2017).
39. Johnson, A., Pollard, T., Mark, R., Berkowitz, S. & Horng, S. Mimic-cxr database. *PhysioNet10*. **13026**, C2JT1Q (2019).
40. Caesar, H. et al. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11621–11631 (2020).
41. Sokolova, M. & Guy, L. A systematic analysis of performance measures for classification tasks. *Inf. Process. Manag.* **45** (4), 427–437 (2009).
42. Nam, H., Kim, S. H., Min, D., Ko, B. Y. & Park, Y. H. Towards Understanding of Frequency Dependence on Sound Event Detection. preprint at (2025). https://doi.org/10.48550/arXiv.2502.07208.
43. Powers, D. M. W. Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation. preprint at (2020). https://doi.org/10.48550/arXiv.2010.16061.
44. Naphade, M. et al. Large-scale concept ontology for multimedia. *IEEE MultiMedia*. **13** (3), 86–91 (2006).
45. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **27** (8), 861–874 (2006).
46. Megill, N. & Wheeler, D. A. Metamath: a computer Language for mathematical proofs. *Lulu Com.* (2019).
47. Khetan, V., Ramnani, R., Anand, M., Sengupta, S. & Fano, A. E. Causal bert: Language models for causality detection between events expressed in text. In *Intelligent Computing: Proceedings of the 2021 Computing Conference*. 1, 965–980 (2022).
48. Stosic, D., Stosic, D., Zanchettin, C., Ludermir, T. & Stosic, B. QRNN: $ q $-Generalized random neural network. *IEEE Trans. Neural Networks Learn. Syst.* **28** (2), 383–390 (2016).
49. Guo, D. et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. preprint at (2025). https://doi.org/10.48550/arXiv.2501.12948.
50. Shen, S. et al. How much can clip benefit vision-and-language tasks? preprint at (2021). https://doi.org/10.48550/arXiv.2107.06383.
51. Chen, Y. C. et al. Uniter: Universal image-text representation learning. In *European conference on computer vision*. 104–120 (2020).

Acknowledgements

This research was supported by Gansu Key Research and Development Special Project for Ecological Civilization Construction under Grant No. 24YFFA062, the National Natural Science Foundation of China under Grant No. 32360434, and the Gansu Provincial Department of Education: Industry Support Program Project, China under Grant No. 2023CYZC-11.

Author contributions

S. W and K. C: Writing—original draft，Writing—review and editing，Formal analysis，Data curation，image drawing. M.Y: Writing—review and editing. P. Z and H.D: Data curation, image editing. All authors reviewed the manuscript.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to S.W.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.