

Article

Information Preservation in Communication-Constrained Neural Representation Transfer: Classical and Quantum-Inspired Amplitude Encoding

Naveen Joy ^{1,2} , Aswathi Balakrishnan ¹ , Aparna Rajan ¹ , Lijin Varghese ¹ , Ivin Shajan ¹ ,
Muhammed Shahul ¹ , Levin John ³, Suhail Subair ¹, Neena Xavier ⁴, Sonet D. Thomas ^{1,*} and Rajesh Raju ^{1,*} 

- ¹ Centre for Integrative Omics Data Science (CIODS), Yenepoya (Deemed to be University), Deralakatte, Mangalore 575018, Karnataka, India; naveenjoy.ciods@yenepoya.edu.in (N.J.); aswathib.ciods@yenepoya.edu.in (A.B.); aparnarajan.ciods@yenepoya.edu.in (A.R.); lijinvarghese.ciods@yenepoya.edu.in (L.V.); ivinshajan.ciods@yenepoya.edu.in (I.S.); muhammedshahul.ciods@yenepoya.edu.in (M.S.); suhailsubair.ciods@yenepoya.edu.in (S.S.)
- ² Department of Artificial Intelligence and Machine Learning, Work Integrated Learning Programmes Division (WILP), Birla Institute of Technology and Science (BITS) Pilani, Pilani Campus, Vidya Vihar, Pilani 333031, Rajasthan, India
- ³ Institute for Regeneration and Repair (IRR), The University of Edinburgh, 4-5 Little France Dr., Edinburgh EH16 4UU, UK; levin.john@ed.ac.uk
- ⁴ Department of Instrumentation and Applied Physics, Indian Institute of Science, Bangalore 560012, Karnataka, India; neenaxavier@iisc.ac.in
- * Correspondence: sonetdthomas@yenepoya.edu.in (S.D.T.); rajeshraju@yenepoya.edu.in (R.R.)

Abstract

Neural representations learned by neural networks encode task-relevant information, but their behavior under communication constraints is still poorly understood. We propose an information-theoretic framework for the comparison of classical compression and quantum-inspired amplitude encoding for the transfer of representations under controlled dimensional constraints. We analyze frozen intermediate features extracted from a pretrained ResNet-50 on the ImageNet-1K validation set at different depths in the network. The classical pipeline performs rate-controlled dimensionality reduction and uniform quantization, while the quantum-inspired pipeline does amplitude encoding and computational-basis probability decoding without variational circuits or trainable quantum parameters. Quality of representation was evaluated using differential entropy and a decoder-based lower bound on mutual information. Classical compression achieves higher decoder-based mutual information lower bounds at moderate and high communication budgets. However, the quantum-inspired pipeline retained a larger fraction of decoder-extractable task-relevant information under severe dimensional compression, with resource-normalized information efficiency providing a complementary perspective to the substantially higher absolute decoder-based mutual information lower bounds of the classical pipeline. Robustness analyses with qubit reduction and physically motivated noise models showed predictable degradation. Ablation and SHAP analyses indicate that task-relevant information is structured within representational subspaces. These findings support quantum-inspired representation transfer as a framework for studying communication-efficient neural representation transfer under low-bandwidth constraints.

Keywords: quantum-inspired computing; learned neural representations; representation transfer; information theory; mutual information; amplitude encoding; information preservation



Academic Editors: Muhammad Asad Arshed, Atif Alvi, Christine Dewi and Muhammad Kabir

Received: 3 August 2026

Revised: 28 August 2026

Accepted: 4 September 2026

Published: 22 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Deep neural networks progressively transform raw input signals into structured, high-dimensional representations that capture increasingly abstract and task-relevant information across successive network layers [1,2]. Although these learned representations underpin the success of modern deep learning, they are typically regarded as internal computational objects confined to a single model or device [3–5]. In distributed inference, split computing, edge-cloud architectures, and other bandwidth-constrained environments, however, intermediate representations must often be transmitted between computational components. This raises a fundamental question: how efficiently can learned neural representations be transferred while preserving the semantic information required for downstream inference?

To address this challenge, classical communication pipelines commonly employ dimensionality reduction, quantization, and compression to reduce transmission costs. These techniques perform effectively across moderate and high communication budgets but may progressively discard task-relevant information under severe bandwidth constraints, particularly for deeper and semantically richer representations. Quantum-inspired amplitude encoding offers an alternative nonlinear representation transformation by mapping normalized high-dimensional feature vectors into an amplitude-based state representation that can subsequently be analyzed through its computational-basis probability distribution. Although such transformations have attracted increasing interest within quantum machine learning, their information-preservation properties for learned neural representations remain insufficiently understood [6–8].

In this study, we present an information-theoretic framework for simulating amplitude-encoded representations of learned neural features, enabling a controlled comparison between conventional classical compression and a quantum-inspired representation transformation [8,9]. Frozen intermediate representations extracted from a pretrained ResNet-50 evaluated on the ImageNet-1K validation set are analyzed across multiple network depths. Classical representation transfer is modeled using rate-controlled dimensionality reduction and uniform quantization, whereas the quantum-inspired pipeline applies amplitude encoding followed by computational-basis probability decoding, without variational quantum circuits or trainable quantum parameters. Representation quality is evaluated using differential entropy together with a decoder-based lower bound on the mutual information between transferred representations and class labels [3,10,11].

The present framework evaluates a simulated amplitude-probability representation rather than an operational quantum communication protocol and therefore does not account for state-preparation complexity, finite-shot measurement, or hardware execution costs. Consequently, the objective of this work is not to demonstrate a practical quantum communication advantage, but rather to investigate whether amplitude-probability representations preserve task-relevant information under controlled dimensional constraints and to compare their behavior with established classical compression techniques.

Our framework is intentionally designed to isolate the influence of representation structure from optimization and training dynamics. By using frozen neural representations and excluding fine-tuning, stochastic augmentation, and trainable quantum components, the observed behavior can be attributed primarily to the interaction between learned feature geometry and the applied representation-transfer mechanism [12,13]. Through mutual information analysis, robustness experiments, and feature-level attribution, we evaluate how quantum-inspired amplitude encoding preserves decoder-extractable information under severe dimensional constraints, particularly across deeper and globally pooled representations. Complementary robustness analyses further demonstrate predictable degradation under simulated qubit reduction and physically motivated noise models,

while feature-attribution experiments indicate that semantic information is distributed across structured representational subspaces rather than isolated feature dimensions. Together, these findings establish quantum-inspired representation transfer as a principled computational framework for investigating communication-efficient neural representation transfer under constrained bandwidth conditions and provide a foundation for future studies incorporating finite-shot measurement, explicit state-preparation costs, and quantum hardware implementations [14–16].

2. Materials and Methods

2.1. Data Acquisition

All experiments were conducted using the ImageNet-1K validation set [17]. The dataset comprises 50,000 images spanning 1000 object categories, with 50 validation images per category. The ImageNet training split was not used at any stage of this study. The decision to restrict all experiments to the ImageNet-1K validation set was made for several reasons. First, it eliminates any possibility of training-evaluation leakage because all neural representations are extracted exclusively from a pretrained model operating in inference mode. Second, the validation set provides a balanced and standardized benchmark that facilitates reproducible comparison across representation extraction, compression, and information-theoretic analyses. Finally, limiting the study to the validation split reduces computational complexity while maintaining consistency with established evaluation protocols for pretrained computer vision models [17–19].

2.2. Image Preprocessing Pipeline

All input images were preprocessed using the deterministic ImageNet evaluation pipeline implemented with torchvision. Each image was resized such that its shorter side equaled 256 pixels while preserving the original aspect ratio. A centered (224×224) crop was then extracted to obtain the fixed spatial resolution required by ResNet-50. The resulting RGB image was converted to a floating-point tensor with pixel intensities scaled to the interval, $[0, 1]$ yielding

$$X \in \mathbb{R}^{3 \times 224 \times 224}$$

Each RGB channel was subsequently normalized independently according to

$$X'_{c,h,w} = \frac{X_{c,h,w} - \mu_c}{\sigma_c}, \quad (1)$$

where $c \in \{R, G, B\}$ denotes the color channel, $h \in \{1, \dots, 224\}$ and $w \in \{1, \dots, 224\}$ denote the spatial row and column indices, respectively, $X_{c,h,w}$ is the corresponding normalized-input pixel value before channel normalization, and μ_c and σ_c are the ImageNet channel-wise mean and standard deviation. Specifically,

$$\mu = (0.485, 0.456, 0.406), \quad \sigma = (0.229, 0.224, 0.225)$$

No stochastic data augmentation was applied. The preprocessing pipeline was implemented using the torchvision transformation library, and the ResNet-50 backbone was initialized using the ImageNet-1K V1 pretrained weights. This normalization ensures consistency between the network inputs and the preprocessing applied during ResNet-50 pretraining, thereby minimizing distributional mismatch [18,19].

Resizing was used with antialiasing, followed by a centered 224×224 crop; pixel conversion used the standard torchvision tensor conversion without additional rounding. Deterministic preprocessing was adopted to avoid augmentation-induced variability in the extracted feature distributions. Because the objective of this study is to analyze learned

representations rather than improve predictive performance, no stochastic data augmentation was applied. This design reduces preprocessing-related variability in the subsequent entropy, mutual information, compression, and amplitude-encoding analyses.

2.3. Representation Extraction

To investigate the information-preservation characteristics of learned neural representations, feature embeddings were extracted from a pretrained ResNet-50 convolutional neural network trained on the ImageNet-1K dataset [18,19]. The network was used exclusively as a fixed feature extractor, with all model parameters frozen throughout the study. No fine-tuning, gradient computation, or parameter updates were performed.

The pretrained network is represented as

$$f_{\theta} = \mathbb{R}^{3 \times 224 \times 224} \rightarrow \mathbb{R}^{1000}$$

where θ denotes the pretrained model parameters.

For each input image x_i , intermediate feature representations were extracted from selected network layers during the forward pass using forward hooks. The representation at layer l is defined as

$$z_i^{(l)} = f_{\theta}^{(l)}(x_i) \quad (2)$$

where l indexes the selected network layer, $f_{\theta}^{(l)}(x_i)$ represents the mapping from the network input to the output activation of layer l , $z_i^{(l)}$ is the feature representation of the image x_i extracted from layer l .

These intermediate representations constitute the primary data used for all subsequent analyses, including entropy estimation, mutual information computation, representation compression, and quantum communication experiments.

Selected Layers

Representations were extracted from three depths of the pretrained ResNet-50 backbone: layer1, representing an early convolutional stage; layer3, representing a deeper convolutional stage; and avgpool, representing the final global average-pooling layer [20].

For an input image x_i , the corresponding activation tensors are denoted as

$$A_i^{(1)} \in \mathbb{R}^{256 \times 56 \times 5}$$

$$A_i^{(3)} \in \mathbb{R}^{1024 \times 14 \times 14}$$

$$A_i^{(p)} \in \mathbb{R}^{2048 \times 1 \times 1}$$

where the superscripts (1), (3), and (p) refer to layer1, layer3, and avgpool, respectively.

For the convolutional layers, the spatial dimensions were removed using global average pooling. Given an activation tensor

$$A_i^{(l)} \in \mathbb{R}^{C_l \times H_l \times W_l}$$

the pooled representation for channel c is computed as

$$z_{i,c}^{(l)} = \frac{1}{H_l W_l} \sum_{h=1}^{H_l} \sum_{w=1}^{W_l} A_{i,c,h,w}^{(l)} \quad (3)$$

where C_l denotes the number of channels, and H_l and W_l denote the spatial height and width of the activation tensor at layer l . The resulting pooled representation is

$$z_i^{(l)} = [z_{i,1}^{(l)}, z_{i,2}^{(l)}, \dots, z_{i,C_l}^{(l)}]^T \in \mathbb{R}^{C_l}$$

Accordingly, the extracted representation dimensions are $z_i^{(1)} \in \mathbb{R}^{256}$, $z_i^{(3)} \in \mathbb{R}^{1024}$, $z_i^{(p)} \in \mathbb{R}^{2048}$.

For the avgpool layer, the output tensor already has spatial dimensions 1×1 . Therefore, it is directly reshaped into a 2048-dimensional feature vector without additional spatial averaging.

This extraction procedure defines a fixed and deterministic mapping from each input image to its corresponding layer-wise feature representation. All subsequent entropy estimation, mutual information probing, compression, amplitude-probability transformation, robustness analysis, and feature-attribution experiments were performed exclusively on these frozen representations. The extracted vectors were stored in HDF5 format to support efficient processing and reproducibility.

2.4. Entropy Estimation

To quantify the information content of the learned neural representations, we estimate the differential entropy of the layer-wise feature distributions using the non-parametric Kozachenko–Leonenko (KL) k -nearest neighbor estimator [21,22]. For a given network layer l , the extracted representation matrix is denoted by

$$Z^l \in \mathbb{R}^{N \times d_l},$$

where N denotes the number of samples and d_l is the dimensionality of the representation at layer l .

Prior to entropy estimation, each feature dimension is independently standardized using z -score normalization to eliminate scale-dependent effects:

$$\hat{z}_{i,j} = \frac{z_{i,j} - \mu_j}{\sigma_j + \varepsilon} \quad (4)$$

where μ_j and σ_j denote the mean and standard deviation of the feature j , respectively, and ε is a small positive constant introduced for numerical stability.

The differential entropy is then estimated using the Kozachenko estimator, which estimates entropy directly from nearest-neighbor distances without requiring explicit probability density estimation [23]. The Kozachenko–Leonenko estimator was evaluated using $k = 5$ nearest neighbors as the primary setting, with sensitivity assessed at $k \in \{3, 5, 10, 20\}$. Given a set of normalized samples

$$\{x_1, x_2, \dots, x_N\}, \quad x_i \in \mathbb{R}^d,$$

the estimator is defined as

$$\hat{H}(X) = \psi(N) - \psi(k) + \log V_d + \frac{d}{N} \sum_{i=1}^N \log \varepsilon_i \quad (5)$$

where $\psi(\cdot)$ denotes the digamma function, k is the number of nearest neighbors, ε_i is the Euclidean distance between sample x_i and its k^{th} nearest neighbor and V_d denotes the volume of the unit ball in d -dimensional Euclidean space.

Differential entropy provides a descriptive characterization of the statistical structure of the learned representations, whereas the subsequent mutual information analysis quantifies the amount of task-relevant information preserved for downstream prediction. Throughout this work, entropy is used as a comparative descriptor of representation distributions across layers rather than as a direct measure of semantic information, thereby enabling consistent analysis without assuming a parametric probability model [24].

2.5. Decoder-Based Mutual Information Lower Bound

To quantify the task-relevant information extractable from the learned representations, we evaluate a decoder-based lower bound on the mutual information between the representation and the corresponding class labels Y . The mutual information is defined as

$$I(Z; Y) = H(Y) - H(Y|Z), \quad (6)$$

where $H(Y)$ denotes the entropy of the label distribution and $H(Y|Z)$ denotes the conditional entropy of the labels given the learned representations [25].

Direct estimation of $H(Y|Z)$ is computationally challenging in high-dimensional representation spaces. Therefore, we adopt a decoder-based variational lower bound on the mutual information [26]. Specifically, a multinomial logistic regression decoder

$$g_{\phi} : \mathbb{R}^{d_l} \rightarrow \mathbb{R}^C$$

parameterized by ϕ , is trained to predict the class labels from the frozen representations, where d_l denotes the representation dimensionality at layer l and C denotes the number of classes. The parameter set ϕ comprises the learned weights and bias terms of the decoder. The predicted class probabilities are given by

$$\hat{Y} = g_{\phi}(Z)$$

The mutual information is then lower bounded as

$$I(Z; Y) \geq H(Y) - L_{CE}(Y, \hat{Y}) \quad (7)$$

where L_{CE} denotes the cross-entropy loss between the true class labels Y and the predicted class probabilities $\hat{Y}$. The cross-entropy loss provides an empirical upper bound on the conditional entropy $H(Y|Z)$, thereby yielding a tractable lower bound on the mutual information.

The entropy of the ImageNet-1K label distribution is computed as

$$H(Y) = -\sum_y p(y) \log p(y) \quad (8)$$

where $p(y)$ denotes the probability of class y .

Since the ImageNet-1K validation set is balanced across all $C = 1000$ classes,

$$p(y) = \frac{1}{C},$$

which gives

$$H(Y) = \ln(C) \quad (9)$$

$$H(Y) \approx \ln(1000) \approx 6.91$$

During validation, the normalization parameters, PCA transformation, and quantization parameters were fitted exclusively on the training folds and subsequently applied to

the corresponding held-out validation folds to prevent information leakage. The decoder was trained on the resulting training-fold representations, and mutual information lower bounds were evaluated on the held-out validation folds. Consequently, the reported values correspond to decoder-based lower bounds on the mutual information rather than exact mutual information estimates [27].

The decoder was implemented as multinomial logistic regression using the SAGA solver with a regularization strength of $C = 0.1$ and a maximum of 2000 solver iterations. Decoder performance was evaluated using stratified 5-fold cross-validation across five independent random seeds (42, 123, 456, 789, and 2024). The convergence tolerance was set to 0.001. Mutual information lower bounds were calculated as $H(Y) - CE$, where CE denotes the cross-entropy of the decoder predictions, and results are reported as mean $\pm$ standard deviation across the five random seeds.

2.6. Classical Communication Channel

To establish a classical baseline for representation transfer, we model communication as a rate-limited classical pipeline comprising dimensionality reduction, finite-precision quantization, reconstruction, and downstream decoding. Let $Z^{(l)} \in \mathbb{R}^{N \times d_l}$ denote the representation matrix extracted from the network layer l , where N is the number of samples and d_l is the dimensionality of the representation.

To reduce the communication cost while preserving the dominant variance within the representation space, the extracted representations are first compressed using Incremental Principal Component Analysis (Incremental PCA) [28].

For a specified communication budget, the representations are projected onto a lower-dimensional subspace according to

$$Z_{PCA} = ZW \quad (10)$$

where $W \in \mathbb{R}^{d_l \times m}$ contains the first m principal components.

For each communication budget, the training representations are centered using the training-set mean μ and projected onto the first k PCA components:

$$Z = (X - \mu)W_k$$

where W_k contains the first k PCA loading vectors. The PCA parameters and component-wise quantization ranges are fitted using the training partition only. Each PCA component is uniformly quantized to 8 bits using its training-set minimum and maximum, and the same ranges are applied to the held-out validation partition:

$$Q_j = \text{round}\left(255 \frac{Z_j - l_j}{u_j - l_j}\right) \quad (11)$$

After de-quantization, the representation is reconstructed as

$$\hat{Z}_j = l_j + \frac{Q_j}{255}(u_j - l_j) \quad (12)$$

where u_j and l_j are the training-set minimum and maximum values of the j -th PCA component. The de-quantization parameters were estimated exclusively from the training partition and applied unchanged to the corresponding held-out validation partition.

$$\hat{X} = \hat{Z}W_k^T + \mu \quad (13)$$

The decoder uses multinomial logistic regression with the SAGA solver, $C = 0.1$ (inverse regularization strength), tolerance 10^{-3} , and a maximum of 2000 iterations. Stratified five-fold cross-validation is performed using five random seeds (42, 123, 456, 789, and 2024), with the mutual information lower bound evaluated on the held-out validation folds.

The quantization range was likewise estimated from the training portion and applied unchanged to the corresponding held-out validation representation [29,30].

To evaluate the task-relevant information preserved after compression, the reconstructed representations $\hat{z}$ are assessed using the multinomial logistic regression decoder and decoder-based mutual information lower bound defined in Section 2.5. The decoder is configured with a regularization parameter $C = 0.1$ and the SAGA optimization solver [31,32].

This classical pipeline therefore provides a rate-controlled baseline for assessing the information-preservation characteristics of compressed neural representations and for comparison with the proposed quantum-inspired amplitude-encoding pipeline under the same nominal communication-budget settings.

2.7. Quantum-Inspired Amplitude Encoding Pipeline

To investigate quantum-inspired representation transfer, we model a simulated amplitude-encoding pipeline that transforms learned neural representations into simulated amplitude-encoded state representations and evaluates the resulting amplitude-probability representations using the decoder-based information-theoretic framework described in Section 2.5. Rather than implementing an operational quantum communication protocol, the proposed framework simulates the mathematical transformation associated with amplitude encoding, enabling a controlled comparison between classical compression and quantum-inspired representation transformations under the same nominal communication-budget settings.

For each layer-wise representation matrix $Z^{(l)}$ defined in Section 2.6, an individual feature vector is denoted by

$$z_i^{(l)} \in \mathbb{R}^{d_l}$$

where d_l is the dimensionality of the representation at layer l .

Each representation vector is processed independently throughout the quantum-inspired pipeline. The objective of the encoding stage is to transform the real-valued neural representation into an amplitude-based state representation while preserving its relative geometric structure after normalization.

Each feature vector is subsequently mapped to an amplitude-encoded state representation using amplitude encoding [6,8]. Prior to encoding, the vector is normalized using the Euclidean (L_2) norm:

$$\tilde{z}_i^{(l)} = \frac{z_i^{(l)}}{\|z_i^{(l)}\|_2} \quad (14)$$

where $\|\cdot\|_2$ denotes the Euclidean norm.

For a representation containing d features, the corresponding state representation requires

$$n = \lceil \log_2 d \rceil \quad (15)$$

qubits, where n denotes the number of qubits required to represent the normalized feature vector and d is the dimensionality of the transmitted representation [33].

Since amplitude encoding requires a Hilbert-space dimension of 2^n , the normalized feature vector is zero-padded when necessary and embedded into an amplitude-encoded state representation.

$$|\psi\rangle = \sum_{i=1}^{2^n} \tilde{z}_i |i\rangle \quad (16)$$

where $|i\rangle$ denotes the computational basis states and $\tilde{z}_i = 0$ for zero-padded entries [34].

For the simulated amplitude encoding, the number of qubits scales logarithmically with the encoded dimension, while the classical state-vector simulation requires storage proportional to the number of amplitudes.

Rather than performing finite-shot quantum measurement, the computational-basis measurement probabilities are computed directly from the simulated state representation [35]:

$$p_i = |\langle i|\psi\rangle|^2 \quad (17)$$

where $\sum_i p_i = 1$. Only the probabilities corresponding to the original (non-padded) feature dimensions are retained, yielding an amplitude-probability representation for subsequent analysis. For each communication budget, the input representation is truncated prior to amplitude encoding according to

$$k = \max(2, \lfloor \beta d_l \rfloor)$$

where k denotes the retained feature dimensionality before encoding. This procedure follows the implementation used in the quantum-inspired pipeline and enables comparison with the classical baseline under the same nominal communication-budget settings.

For fair comparison, the resulting amplitude-probability representations were evaluated using the same 8-bit numerical precision as the classical compression pipeline.

All experiments employ deterministic state-vector simulation without parameterized quantum circuits, variational optimization, or quantum hardware execution [36,37]. Consequently, the framework evaluates the information-preservation characteristics of amplitude-probability representations rather than an operational quantum communication protocol.

The resulting amplitude-probability representations are evaluated using the same multinomial logistic regression decoder and decoder-based mutual information lower bound described in Section 2.5, ensuring methodological consistency across the classical and quantum-inspired pipelines.

For comparison, the retained representation dimension, classical communication cost, and quantum resource requirement are treated as distinct quantities. Because a classical bit and a qubit represent different physical communication resources, their counts are not converted into a common unit in the present analysis. The resulting efficiency values therefore represent scheme-specific resource normalization rather than a physically calibrated bit-to-qubit efficiency ratio. The classical pipeline transmits 8 bits per retained PCA coefficient, whereas the quantum-inspired representation encodes the retained feature vector using $N = \lceil \log_2 d \rceil$ qubits, where d is the retained dimension after truncation. The present simulation accounts for the representation and qubit requirements only; state-preparation and measurement costs are outside the scope of this study. Before amplitude encoding, the first d PCA components are retained, providing a deterministic ordering shared across the compared representations.

2.8. Quantum Robustness and Ablation Analysis

Targeted robustness and ablation analyses are performed to evaluate the stability of the proposed quantum-inspired amplitude-encoding pipeline under resource and noise perturbations. These experiments investigate how the resulting amplitude-probability representa-

tions respond to reductions in representational resources and simulated quantum noise. The analyses are diagnostic in nature and are not used to improve model performance.

To examine the effect of reduced representational resources, the number of qubits available for amplitude encoding is progressively decreased [38]. For a representation of dimensionality d , the corresponding number of qubits is

$$n = \lceil \log_2 d \rceil$$

For lower communication-budget settings, Shannon entropy is used as a descriptive measure of the resulting probability distribution to examine changes in information dispersion following amplitude encoding.

To investigate robustness under simulated hardware imperfections, physically motivated noise models are incorporated into the state-vector simulation, including phase-damping and depolarizing noise channels [39]. Noise is applied uniformly across all qubits using a mixed-state simulator, and the resulting output probability distributions are analyzed as a function of the applied noise probability.

The Shannon entropy of the amplitude-probability representation is computed as

$$H(P) = -\sum_i p_i \log p_i \quad (18)$$

where p_i denotes the computational-basis measurement probability associated with the basis state i .

These experiments are intended to evaluate whether the amplitude-probability representations exhibit smooth and predictable degradation under simulated resource limitations and noise perturbations. The analysis assumes idealized state preparation and probability computation and therefore does not account for practical state-loading complexity, finite-shot measurement, or quantum hardware execution costs [40].

2.9. Feature Attribution via SHAP Analysis

To provide feature-level interpretability of the learned representations, SHAP (SHapley Additive exPlanations) analysis is performed exclusively on the avgpool representations, which capture the highest level of semantic abstraction within the network [41]. The top-100 classes were selected to provide a computationally tractable and controlled setting for feature-attribution analysis while retaining broad category diversity; this restricted class coverage is acknowledged as a limitation of the SHAP analysis.

A linear multinomial logistic regression decoder is trained using a restricted subset comprising the top-100 ImageNet classes, and LinearSHAP is employed to estimate feature attributions [42]. For the linear decoder,

$$f(x) = w^T x + b \quad (19)$$

where w denotes the learned weight vector and b is the bias term.

This analysis serves two purposes. First, it examines whether downstream predictions arise from distributed contributions across multiple features rather than dependence on a small number of isolated dimensions. Second, it provides qualitative confirmation that the observed feature-attribution patterns are consistent with the information-theoretic analyses presented throughout the study. SHAP is used exclusively as an interpretability tool and is not incorporated into any quantitative evaluation, optimization, or comparison.

For the noise robustness analysis, features were stratified into the highest- and lowest-importance 20% according to mean absolute SHAP values, and their responses were

evaluated separately under phase-damping and depolarizing noise at probabilities of 0, 0.05, 0.10, and 0.20.

The robustness, ablation, and feature-attribution analyses are intended solely as complementary validation studies that support the interpretation of the proposed framework. The principal findings of this work are derived from the information-theoretic analyses presented in the Results section.

2.10. Statistical Validation

All decoder-based mutual information results were evaluated across five independent random seeds and are reported as mean $\pm$ standard deviation. Ninety-five percent confidence intervals were calculated for the decoder-based mutual information lower bounds, and paired *t*-tests were used to compare the classical and quantum-inspired pipelines at matched communication budgets.

3. Results and Discussion

3.1. Classical Compression Versus Quantum-Inspired Amplitude-Encoding Performance

We first compare the information-preservation characteristics of the classical compression pipeline and the quantum-inspired amplitude-encoding pipeline under matched communication-budget and precision settings across three representation depths (layer1, layer3, and avgpool). Performance is evaluated using the decoder-based lower bound on mutual information, $I(Z; Y)$. Resource-normalized information efficiency is defined separately for the two representation-transfer schemes as the decoder-based mutual information lower bound divided by the corresponding resource count. For the classical pipeline, this quantity is expressed in nats per transmitted bit, whereas for the quantum-inspired pipeline it is expressed in nats per required qubit. These units are not assumed to be physically equivalent, and no bit-to-qubit conversion factor is imposed because the present study does not model an operational quantum communication channel or establish a physical equivalence between classical bits and qubits. Accordingly, Figure 1 provides a resource-normalized comparison within each communication model rather than claiming a direct physical equivalence between bit- and qubit-based efficiencies. Across all representations and communication budgets, the classical and quantum-inspired pipelines differed significantly in paired comparisons (all $p < 0.001$).

At moderate to high communication budgets, the classical compression pipeline consistently achieves higher decoder-based mutual information lower bounds across all evaluated layers. This indicates that the implemented classical compression pipeline preserves more linearly extractable task-relevant information at moderate and high nominal budgets.

The KL entropy estimates increased systematically with the neighborhood size k , with values of 253.330–287.609 for layer1, 1401.667–1476.186 for layer3, and 2739.267–3061.697 for avgpool across $k \in \{3, 5, 10, 20\}$, indicating that the observed cross-layer entropy ordering is preserved across the tested k values as mentioned in Table 1. For layer1 (256-D), the classical decoder-based mutual information lower bound increased from 0.1759 at 5% budget to 0.4142, 0.7002, and 0.9057 at 10%, 25%, and 50%, respectively. The corresponding quantum-inspired decoder-based mutual information lower bounds were approximately 0.0000, 0.0543, 0.0120, and 0.0089. Thus, the quantum-inspired representation retained approximately 0.0%, 13.1%, 1.7%, and 1.0% of the corresponding classical decoder-based mutual information lower bounds across the four communication budgets.

For layer3 (1024-D), classical MI increased from 1.3681 at 5% budget to 1.8928, 2.5698, and 2.7431 at 10%, 25%, and 50%, respectively. The corresponding quantum-inspired decoder-based mutual information lower bounds were 0.0255, 0.0391, 0.0347, and 0.0240.

These values correspond to approximately 1.9%, 2.1%, 1.3%, and 0.9% of the corresponding classical MI values.




Figure 1. Resource-normalized decoder-based information efficiency of classical and quantum-inspired representations across communication resources. (a) Resource-normalized information efficiency for layer1. (b) Resource-normalized information efficiency for layer3. (c) Resource-normalized information efficiency for avgpool. (d) Cross-layer comparison of the two representation-transfer schemes. Efficiency is calculated as the decoder-based mutual information lower bound per corresponding resource unit: nats per transmitted bit for classical compression and nats per required qubit for quantum-inspired amplitude encoding. Bits and qubits are treated as distinct resource units and are not assumed to be physically equivalent; no bit-to-qubit conversion factor is applied. Therefore, the figure is intended as a resource-normalized comparison within the respective communication models rather than a claim of direct physical equivalence between classical and quantum communication resources. Values are reported in nats per resource unit.

Table 1. Sensitivity analysis of Kozachenko–Leonenko entropy estimates across different nearest-neighbor parameters k for the three neural representation layers.

Representation	$k = 3$	$k = 5$	$k = 10$	$k = 20$
layer1	253.330	262.464	274.844	287.609
layer3	1401.667	1422.088	1449.074	1476.186
avgpool	2739.267	2841.141	2963.054	3061.697

For avgpool, classical MI increased from 3.7119 at 5% budget to 3.9066, 3.9224, and 3.9818 at 10%, 25%, and 50%, respectively. The corresponding quantum-inspired decoder-based MI lower bounds were 0.2665, 0.1841, 0.1114, and 0.0747, corresponding to approximately 7.2%, 4.7%, 2.8%, and 1.9% of the classical values. Thus, avgpool shows the strongest relative retention among the three representations, although the absolute quantum-inspired MI remains substantially below the classical baseline. These ratios represent raw relative information retention and should be distinguished from information efficiency, which normalizes the mutual information lower bound by the actual communication resource. The relative performance gap is calculated as the percentage difference between the classical and quantum-inspired mutual information lower bounds, normalized by the classical value. All mutual information values are reported in nats. Information efficiency and the corresponding raw decoder-based mutual information lower bounds are presented in Figures 1 and 2, and Table 2 respectively.




Figure 2. Decoder-based mutual information lower bound as a function of communication budget.

The decoder-based lower bound on mutual information $I(Z;Y)$ is reported versus normalized budget fraction for classical compression and quantum-inspired amplitude encoding. (a–c) Results for layer1, layer3, and avgpool, respectively. In all cases, the decoder-based lower bound increases with budget for both schemes. The classical compression pipeline yields higher lower-bound values across all nominal budgets under the implemented PCA, quantization, and reconstruction procedure. The quantum-inspired

amplitude-encoding pipeline exhibits slower growth, reflecting the combined effects of feature truncation, normalization, and amplitude-probability transformation.

Table 2. Decoder-based mutual information lower bounds across communication budgets. Values are reported as mean $\pm$ standard deviation across five independent random seeds, with corresponding 95% confidence intervals. p -values are from paired comparisons between the classical and quantum-inspired pipelines at matched communication budgets. Mutual information values are reported in nats.

Representation	Budget	Classical MI Lower Bound (Mean $\pm$ SD)	Classical 95% CI	Quantum-Inspired MI Lower Bound (Mean $\pm$ SD)	Quantum 95% CI	p -Value
layer1	0.05	0.1759 $\pm$ 0.0013	0.1743–0.1775	−0.0000 $\pm$ 0.0000	−0.0000–−0.0000	6.36×10^{-10}
layer1	0.10	0.4142 $\pm$ 0.0024	0.4113–0.4172	0.0543 $\pm$ 0.0008	0.0533–0.0553	2.26×10^{-10}
layer1	0.25	0.7002 $\pm$ 0.0040	0.6952–0.7052	0.0120 $\pm$ 0.0002	0.0118–0.0122	2.69×10^{-10}
layer1	0.50	0.9057 $\pm$ 0.0052	0.8992–0.9122	0.0089 $\pm$ 0.0001	0.0089–0.0090	2.79×10^{-10}
layer3	0.05	1.3681 $\pm$ 0.0042	1.3628–1.3733	0.0255 $\pm$ 0.0001	0.0254–0.0257	2.06×10^{-11}
layer3	0.10	1.8928 $\pm$ 0.0106	1.8796–1.9060	0.0391 $\pm$ 0.0001	0.0390–0.0392	2.62×10^{-10}
layer3	0.25	2.5698 $\pm$ 0.0124	2.5545–2.5852	0.0347 $\pm$ 0.0001	0.0346–0.0349	1.36×10^{-10}
layer3	0.50	2.7431 $\pm$ 0.0146	2.7249–2.7612	0.0240 $\pm$ 0.0001	0.0239–0.0241	1.99×10^{-10}
avgpool	0.05	3.7119 $\pm$ 0.0077	3.7023–3.7216	0.2665 $\pm$ 0.0002	0.2663–0.2668	5.67×10^{-12}
avgpool	0.10	3.9066 $\pm$ 0.0095	3.8948–3.9184	0.1841 $\pm$ 0.0002	0.1839–0.1843	1.02×10^{-11}
avgpool	0.25	3.9224 $\pm$ 0.0084	3.9119–3.9328	0.1114 $\pm$ 0.0001	0.1112–0.1115	6.03×10^{-12}
avgpool	0.50	3.9818 $\pm$ 0.0101	3.9692–3.9943	0.0747 $\pm$ 0.0000	0.0746–0.0747	1.08×10^{-11}

These results show a clear trend that depends on depth: as the number of dimensions in a representation and the level of semantic abstraction rise, the quantum-inspired amplitude-encoding pipeline becomes increasingly competitive, especially in low-bandwidth situations.

3.2. Robustness and Ablation of the Quantum-Inspired Pipeline

To assess the robustness of the simulated quantum-inspired amplitude-encoding pipeline, we perform quantum-specific ablation studies aimed at reducing qubit count and injecting noise, using the Shannon entropy of the simulated computational-basis probability distributions as a descriptive metric. Progressive qubit removal shows smooth, monotonic entropy scaling (from 2.02 at 4 qubits to 4.04 at 7 qubits), indicating a smooth increase in probability-distribution entropy with representational dimensionality. Under phase-damping noise with $p \in [0.0, 0.2]$, the entropy remains unchanged (≈ 4.80 – 4.89). This is expected because phase damping modifies the relative phases of the amplitudes while leaving the computational-basis probabilities $p_i = |a_i|^2$ unchanged. Since entropy in this analysis is calculated from these probabilities, the resulting entropy remains unchanged across the tested noise probabilities. We now explicitly state that this reflects the phase-insensitivity of the measured representation rather than complete robustness of the underlying quantum state. Depolarising noise, on the other hand, causes entropy to rise steadily and clearly (from 4.80 at ($p = 0.0$) to 5.44 at ($p = 0.20$)). The increase in entropy is consistent with progressive randomization of the simulated probability distribution as the depolarizing-noise strength increases. These results, visualized in Figure 3, demonstrate that the simulated amplitude-probability representations exhibit predictable degradation under the adopted noise models.




Figure 3. Robustness of the quantum-inspired amplitude-encoding pipeline under qubit reduction and simulated noise. (a) Qubit-count reduction analysis showing the relationship between the number of transmitted qubits and the corresponding entropy. (b) SHAP-stratified phase-damping noise analysis comparing high- and low-importance feature groups across increasing noise strengths. (c) SHAP-stratified depolarizing-noise analysis under the same noise levels. Error bars represent the standard deviation across repeated runs.

Feature-importance-stratified analysis showed distinct responses between high- and low-SHAP features. Under phase-damping noise, the entropy remained unchanged across the tested noise probabilities, with values of 4.528 ± 0.240 nats for high-SHAP features and 4.257 ± 0.389 nats for low-SHAP features. In contrast, depolarizing noise increased entropy progressively from 4.528 ± 0.240 to 5.358 ± 0.064 nats for high-SHAP features and from 4.257 ± 0.389 to 5.276 ± 0.113 nats for low-SHAP features as the noise probability increased from 0 to 0.20 as shown in Table 3.

Table 3. Feature-importance-stratified robustness of avgpool representations under phase-damping and depolarizing noise. Values represent the mean $\pm$ standard deviation of Shannon entropy across 200 samples for the highest- and lowest-importance 20% of features identified by mean absolute SHAP values.

Feature Group	Noise	0.00	0.05	0.10	0.20
Low-SHAP	Phase damping	4.257 ± 0.389	4.257 ± 0.389	4.257 ± 0.389	4.257 ± 0.389
High-SHAP	Phase damping	4.528 ± 0.240	4.528 ± 0.240	4.528 ± 0.240	4.528 ± 0.240
Low-SHAP	Depolarizing	4.257 ± 0.389	4.711 ± 0.276	4.977 ± 0.205	5.276 ± 0.113
High-SHAP	Depolarizing	4.528 ± 0.240	4.908 ± 0.167	5.125 ± 0.121	5.358 ± 0.064

3.3. Classical Ablation Analysis

We next examine classical feature-ablation experiments to assess how linearly extractable task-relevant information is distributed across representation dimensions. The baseline decoder-based mutual information lower bound in layer1 is 0.393. Keeping only the top-variance features raises the decoder-based lower bound to 0.662 (relative MI = 1.684), while the bottom-variance features retain a substantial decoder-based lower bound of 0.617. Random masking only changes the decoder-based lower bound by a small amount (0.415), but feature shuffling reduces the decoder-based lower bound to -1.737 , indicating that the selected linear decoder no longer provides a nontrivial positive lower bound under this perturbation. The baseline MI for layer3 is 2.557. The top- and bottom-variance features keep about the same amount of information (2.178 and 2.255, respectively). Random masking lowers MI to 1.826, and shuffling yields a negative, and therefore non-informative, decoder-based lower bound of -1.311 . At the avgpool level, the baseline decoder-based lower bound is 4.993. Both the top- and bottom-variance subsets keep most of the information (4.722 and 4.692, respectively), random masking lowers MI to 3.882, and shuffling causes a complete collapse (-1.591). In general, these results show that linearly extractable task-relevant information is distributed non-uniformly and redundantly across the evaluated feature dimensions. Substantial decoder performance is retained after variance-based selection or random masking, whereas feature shuffling substantially disrupts the relationship between feature values and class labels. All reported mutual information values correspond to decoder-based lower bounds and therefore reflect task-relevant information extractable by a linear probe rather than intrinsic representation entropy. This result in Figure 4 shows that there are strong nonlinear dependencies instead of relying on isolated dimensions.




Figure 4. Classical feature ablation and semantic information distribution across representation layers. Decoder-based mutual information lower bounds obtained under different feature-selection and perturbation strategies are shown for layer1, layer3, and avgpool. For each representation, the baseline is compared with top-variance, bottom-variance, random-mask, and shuffled-feature conditions. Positive values indicate retained mutual information, whereas negative values from feature shuffling indicate disruption of the learned semantic structure. Values are reported in nats.

3.4. SHAP Analysis

To complement the information-theoretic analysis with feature-level interpretability, we conduct SHAP (SHapley Additive exPlanations) analysis on the avgpool representations, employing a linear SHAP explainer in conjunction with the multinomial logistic decoder. The SHAP summary plot in Figure 5 indicates a concentrated attribution pat-

tern, with a relatively small subset of dimensions exhibiting consistently large attribution magnitudes. Only a small number of representation dimensions show consistently large SHAP magnitudes across samples. Most remaining dimensions exhibit attribution values concentrated near zero for the evaluated decoder and samples. The most influential dimensions show stable, sign-consistent contributions across samples. This pattern suggests that the selected decoder relies more strongly on a structured subset of representation dimensions than on the full feature space uniformly. This behavior is like the classical ablation result, where high-variance feature subsets retained a comparatively large fraction of the decoder-extractable task-relevant information, and random masking reduced the decoder-based lower bound.




Figure 5. SHAP-based analysis of feature importance and contribution directionality. (a) Global feature importance based on mean absolute SHAP values for the top-ranked representation features. (b) Distribution of SHAP values for the selected features, illustrating the magnitude and spread of feature contributions. (c) Directionality of feature contributions, showing the balance between positive and negative SHAP effects. These analyses characterize the relative contribution of representation dimensions to the decoder output.

Importantly, no single feature completely controls the attribution. Instead, predictive decisions come from small groups of dimensions that have a moderate amount of influence. This shows that there is a balance between sparsity and redundancy. This observation provides no evidence of complete dependence on a single feature within the evaluated decoder. Overall, the SHAP analysis gives qualitative but independent proof that deep neural representations store information in structured, low-dimensional subspaces, which are shown in Figure 5.

Because SHAP analysis was restricted to the avgpool representation and a 100-class subset, the observed feature-attribution patterns should not be assumed to generalize directly to the full 1000-class ImageNet task or to earlier network representations. Earlier layers encode lower-level and more spatially localized features, whereas avgpool provides a higher-level global representation; consequently, their attribution profiles may differ. Similarly, expanding from 100 to 1000 classes could alter the relative importance of repre-

sensation dimensions because of the broader class structure. Therefore, the present SHAP results should be interpreted as an attribution analysis of the selected avgpool representation and 100-class evaluation subset, rather than as a general characterization of feature importance across the entire network or ImageNet label space.

The KL-Entropy details and the communication resource accounting are mentioned in Figures 6 and 7 respectively.




Figure 6. Sensitivity of the Kozachenko–Leonenko entropy estimator to the nearest-neighbor parameter. Entropy estimates are shown for $k \in \{3, 5, 10, 20\}$ across the layer1, layer3, and avgpool representations. The analysis evaluates the stability of the entropy estimates with respect to the choice of nearest-neighbor parameter k , with values reported in nats.




Figure 7. Actual communication resource accounting. (a) Number of retained classical PCA coefficients across the communication-budget fractions for the layer1, layer3, and avgpool representations.

(b) Corresponding number of qubits required for quantum-inspired amplitude encoding, where d is the retained representation dimension. The classical representation retains d PCA coefficients, each quantized to 8 bits, corresponding to an effective classical transmission cost of $8d$ bits. The quantum-inspired resource requirement is expressed in qubits, as determined by the amplitude-encoding dimension.

3.5. Overall Inference

This study recontextualizes learned neural representations as transferable communication objects and evaluates their transmission through a systematic information-theoretic framework. By comparing identical frozen ResNet-50 representations using a classical compression pipeline and a quantum-inspired amplitude-encoding pipeline under controlled communication-budget constraints, we observe that the quantum-inspired approach retains nonzero decoder-extractable task-relevant information under severe dimensional constraints, while the classical pipeline consistently achieves higher absolute decoder-based mutual information lower bounds as the communication budget increases.

The observed trends are consistent across multiple analyses. The decoder-based mutual information results demonstrate that the resource-normalized information efficiency becomes comparatively more favorable for deeper representations under low-budget conditions. Classical feature-ablation experiments indicate that task-relevant information is organized within structured representational subspaces rather than isolated feature dimensions, while SHAP analysis provides complementary qualitative evidence that downstream predictions arise from distributed feature contributions. Furthermore, robustness experiments show smooth and predictable changes in amplitude-probability representations under simulated qubit reduction and physically motivated noise models, supporting the stability of the proposed simulation framework.

Taken together, these findings indicate that quantum-inspired amplitude encoding provides a useful computational framework for studying representation transfer under severe bandwidth constraints, but its substantially lower absolute decoder-based mutual information lower bounds indicate that it is not yet competitive with classical compression for practical downstream classification. Although the present study does not implement operational quantum communication or demonstrate quantum advantage on hardware, it establishes a reproducible baseline for evaluating information preservation in amplitude-encoded representations and provides a foundation for future investigations incorporating trainable quantum encoders, finite-shot measurements, and hardware-based quantum communication experiments.

4. Discussion and Conclusions

This paper presents a systematic information-theoretic examination of neural representation transfer using a classical compression pipeline and a simulated quantum-inspired amplitude-encoding pipeline. Neural representation transfer underpins real-world applications such as distributed edge-cloud inference, federated learning, split neural networks, and bandwidth-constrained sensor-to-server communication systems. Using frozen ResNet-50 representations from ImageNet-1K, we show that the simulated amplitude-probability transformation retains a nontrivial amount of decoder-extractable class information under severe nominal dimensional constraints. Nevertheless, the absolute decoder-based lower bounds of the quantum-inspired pipeline remain substantially below those of classical compression across all evaluated budgets. Thus, the present results should not be interpreted as demonstrating practical competitiveness for downstream classification; rather, they establish a computational baseline for assessing information preservation under extreme dimensional constraints. Robustness analyses show systematic changes in the simulated

amplitude-probability distributions under qubit reduction and the adopted noise models. These findings motivate further investigation of hybrid classical-quantum learning systems, although such extensions would require trainable quantum circuits, hardware execution, or other operational quantum-learning components beyond the simulation-based scope of the present study.

The scope of the current study is delineated by several considerations. First, experiments are performed on the ImageNet-1K validation set and pretrained representations, emphasizing representation analysis over end-to-end learning. Second, quantum-inspired experiments use reduced sample and class subsets because state-vector simulation and repeated decoder evaluation impose substantial computational costs. Amplitude encoding provides a compact mathematical state representation; however, alternative encoding strategies may exhibit different preservation and resource trade-offs. Importantly, all reported information values are decoder-based lower bounds on mutual information obtained using a multinomial logistic regression probe. They therefore quantify task-relevant information that is extractable by the specified linear decoder and should not be interpreted as exact mutual information estimates of the underlying representations.

The computational cost of the quantum-inspired simulation also warrants consideration. For an amplitude-encoded representation of dimension dd , the corresponding state-vector simulator uses $n = \lceil \log_2 d \rceil$ qubits and stores 2^n complex amplitudes, giving a state-vector memory requirement of $O(d)$ for power-of-two dimensions. Thus, the present $d = 2048$ representation corresponds to 11 qubits and 2048 complex amplitudes and is computationally tractable in simulation. Nevertheless, state-vector simulation remains memory-intensive as the representation dimension increases, and the present study does not model the computational cost of physical state preparation, transmission, or measurement. These factors should be considered when extending the framework to substantially larger representations or operational quantum systems.

Accordingly, the framework is designed to be model-agnostic; this is a methodological property rather than an experimentally established result in the present study. Validation across alternative architectures, including Vision Transformers, and additional datasets is required to establish the generality of the observed information-preservation and efficiency trends.

Even with these limitations, the fact that the results are consistent across entropy estimation, mutual information scaling, ablation, interpretability, and noise robustness provides convergent support for the observed trends. The framework is designed to be model-agnostic in principle: any pretrained network whose intermediate representations can be extracted can serve as the feature extractor. In the present study, however, the framework is evaluated experimentally only with ResNet-50 on ImageNet-1K. Future work should evaluate the framework across convolutional and transformer-based architectures and diverse datasets to establish the generality of the reported trends.

Author Contributions: Conceptualization, N.J., S.D.T. and R.R.; methodology, N.J., A.B. and S.D.T.; validation, N.J., A.B., A.R., L.V. and S.D.T.; formal analysis, N.J., A.B. and S.D.T.; writing—original draft preparation, N.J. and S.D.T.; writing—review and editing, A.B., A.R., L.V., I.S., M.S., L.J., S.S., N.X., S.D.T. and R.R.; visualization, L.J. and S.S.; supervision, R.R.; project administration, S.D.T. and R.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No new data is generated in the following study. The code used for analysis is publicly available at <https://github.com/naveen-joy-18/Quantum-Enabled-Communication-of-Learned-Neural-Representations.git> (accessed on 3 September 2026).

Acknowledgments: We thank Yenepoya (Deemed to be University), Mangalore, India, for their support in establishing the Centre for Integrative Omics Data Science (CIODS) computational facility.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this manuscript.

References

1. Tetkova, L.; Brusch, T.; Dorszewski, T.; Mager, F.M.; Aagaard, R.O.; Foldager, J.; Alstrom, T.S.; Hansen, L.K. On convex decision regions in deep network representations. *Nat. Commun.* **2025**, *16*, 5419. [\[CrossRef\]](#) [\[PubMed\]](#)
2. Johnston, W.J.; Fusi, S. Abstract representations emerge naturally in neural networks trained to perform multiple tasks. *Nat. Commun.* **2023**, *14*, 1040. [\[CrossRef\]](#) [\[PubMed\]](#)
3. Zhang, J.; Zhang, Y.; Ji, B.; Chen, A.; Liu, A.; Xu, H. Feature-Driven Semantic Communication for Efficient Image Transmission. *Entropy* **2025**, *27*, 369. [\[CrossRef\]](#) [\[PubMed\]](#)
4. Kutay, E.; Yener, A. Harnessing the Power of Pre-Trained Models for Efficient Semantic Communication of Text and Images. *Entropy* **2025**, *27*, 813. [\[CrossRef\]](#) [\[PubMed\]](#)
5. Yang, H.; Wang, L.; Wang, P.; Li, J.; Qing, L.; He, X. Semantic Adaptive Communication Based on Double-Attention Phase and Compress Estimator for Wireless Image Transmission. *Sensors* **2025**, *25*, 7201. [\[CrossRef\]](#) [\[PubMed\]](#)
6. Ranga, D.; Rana, A.; Prajapat, S.; Kumar, P.; Kumar, K.; Vasilakos, A.V. Quantum Machine Learning: Exploring the Role of Data Encoding Techniques, Challenges, and Future Directions. *Mathematics* **2024**, *12*, 3318. [\[CrossRef\]](#)
7. Long, C.; Huang, M.; Ye, X.; Futamura, Y.; Sakurai, T. Hybrid quantum-classical-quantum convolutional neural networks. *Sci. Rep.* **2025**, *15*, 31780. [\[CrossRef\]](#) [\[PubMed\]](#)
8. Eisinger, J.; Gauderis, W.; Huybrecht, L.; Wiggins, G.A. Classical Data in Quantum Machine Learning Algorithms: Amplitude Encoding and the Relation Between Entropy and Linguistic Ambiguity. *Entropy* **2025**, *27*, 433. [\[CrossRef\]](#) [\[PubMed\]](#)
9. Buhrmester, V.; Münch, D.; Arens, M. Analysis of Explainers of Black Box Deep Neural Networks for Computer Vision: A Survey. *Mach. Learn. Knowl. Extr.* **2021**, *3*, 966–989. [\[CrossRef\]](#)
10. Hanson, S.J.; Negishi, M. On the emergence of rules in neural networks. *Neural Comput.* **2002**, *14*, 2245–2268. [\[CrossRef\]](#) [\[PubMed\]](#)
11. Devendra, R.; John, T.C.; Kappil, S. What is a Gaussian channel, and when is it physically implementable using a multiport interferometer? *Philos. Trans. A Math. Phys. Eng. Sci.* **2025**, *383*, 20240422. [\[CrossRef\]](#) [\[PubMed\]](#)
12. Nguyen-Van-Yen, B.; Del Moral, P.; Cazelles, B. Stochastic Epidemic Models inference and diagnosis with Poisson Random Measure Data Augmentation. *Math. Biosci.* **2021**, *335*, 108583. [\[CrossRef\]](#) [\[PubMed\]](#)
13. Ikusan, A.; Dai, R. Deep Feature Compression with Rate-Distortion Optimization for Networked Camera Systems. In *Proceedings of the 14th ACM Multimedia Systems Conference, Vancouver, BC, Canada, 2023, 7–10 June 2023*; ACM: New York, NY, USA, 2023; pp. 86–96.
14. Arunachalam, S.; de Wolf, R. Guest Column: A Survey of Quantum Learning Theory. *SIGACT News* **2017**, *48*, 41–67. [\[CrossRef\]](#)
15. Dahan, N.A.; Jaha, E.S. IRIS-QResNet: A Quantum-Inspired Deep Model for Efficient Iris Biometric Identification and Authentication. *Sensors* **2025**, *26*, 121. [\[CrossRef\]](#) [\[PubMed\]](#)
16. Shirmovsky, S.E. Modeling of the quantum entangled state transfer protocol in the cell microtubules. *Biosystems* **2024**, *235*, 105100. [\[CrossRef\]](#) [\[PubMed\]](#)
17. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Kai, L.; Li, F.-F. ImageNet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, USA, 20–25 June 2009*; IEEE: New York, NY, USA, 2009; pp. 248–255.
18. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016*; IEEE: New York, NY, USA, 2016; pp. 770–778.
19. Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.* **2015**, *115*, 211–252. [\[CrossRef\]](#)
20. Lin, M.; Chen, Q.; Yan, S. Network in Network. *arXiv* **2013**, arXiv:1312.4400.
21. Cover, T.M.; Tomas, J.A. Information Theory and Statistics. In *Elements of Information Theory*; John Wiley & Sons, Inc.: Hoboken, NJ, USA, 2005; pp. 347–408.
22. Cover, T.M.; Tomas, J.A. Information Theory and Portfolio Theory. In *Elements of Information Theory*; John Wiley & Sons, Inc.: Hoboken, NJ, USA, 2005; pp. 613–656.
23. Wang, L.; Zhang, Y.; Feng, J. On the Euclidean distance of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **2005**, *27*, 1334–1339. [\[CrossRef\]](#) [\[PubMed\]](#)
24. Gao, W.; Oh, S.; Viswanath, P. Demystifying fixed k-nearest neighbor information estimators. *IEEE Int. Symp. Inf. Theory* **2018**, *64*, 5629–5661. [\[CrossRef\]](#)

25. Cover, T.M.; Tomas, J.A. Entropy, Relative Entropy, and Mutual Information. In *Elements of Information Theory*; John Wiley & Sons, Inc.: Hoboken, NJ, USA, 2005; pp. 13–55.
26. Xiang, G.; Dian, S.; Du, S.; Lv, Z. Variational Information Bottleneck Regularized Deep Reinforcement Learning for Efficient Robotic Skill Adaptation. *Sensors* **2023**, *23*, 762. [[CrossRef](#)] [[PubMed](#)]
27. Wieczorek, A.; Roth, V. On the Difference between the Information Bottleneck and the Deep Information Bottleneck. *Entropy* **2020**, *22*, 131. [[CrossRef](#)] [[PubMed](#)]
28. Balzano, L.; Chi, Y.; Lu, Y.M. Streaming PCA and Subspace Tracking: The Missing Data Case. *Proc. IEEE Inst. Electr. Electron. Eng.* **2018**, *106*, 1293–1310. [[CrossRef](#)] [[PubMed](#)]
29. Yoo, J.; Ban, G. Efficient Deep Learning Model Compression for Sensor-Based Vision Systems via Outlier-Aware Quantization. *Sensors* **2025**, *25*, 2918. [[CrossRef](#)] [[PubMed](#)]
30. Cheng, H.; Zheng, J.; Wu, H.; Qi, K.; He, L. A communication-efficient distributed deep learning remote sensing image change detection framework. *Int. J. Appl. Earth Obs. Geoinf.* **2024**, *129*, 103840. [[CrossRef](#)]
31. Qin, T.; Li, Z.; Zhao, J.; Yan, Y.; Du, Y. Mixed precision quantization based on information entropy. *Sci. Rep.* **2025**, *15*, 12974. [[CrossRef](#)] [[PubMed](#)]
32. Tezuka, T.; Namekawa, S. Information Bottleneck Analysis by a Conditional Mutual Information Bound. *Entropy* **2021**, *23*, 974. [[CrossRef](#)] [[PubMed](#)]
33. Schuld, M.; Killoran, N. Is Quantum Advantage the Right Goal for Quantum Machine Learning? *PRX Quantum* **2022**, *3*, 030101. [[CrossRef](#)]
34. Nielsen, J.A.H.; Kicinski, M.J.; Arge, T.N.; Vijayadharan, K.; Foldager, J.; Borregaard, J.; Meyer, J.J.; Neergaard-Nielsen, J.S.; Gehring, T.; Andersen, U.L. Variational quantum algorithm for enhanced continuous variable optical phase sensing. *npj Quantum Inf.* **2025**, *11*, 70. [[CrossRef](#)]
35. Albantakis, L.; Prentner, R.; Durham, I. Computing the Integrated Information of a Quantum Mechanism. *Entropy* **2023**, *25*, 449. [[CrossRef](#)] [[PubMed](#)]
36. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2018**, *2*, 79. [[CrossRef](#)]
37. Wang, S.; Fontana, E.; Cerezo, M.; Sharma, K.; Sone, A.; Cincio, L.; Coles, P.J. Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.* **2021**, *12*, 6961. [[CrossRef](#)] [[PubMed](#)]
38. Rath, M.; Date, H. Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy. *EPJ Quantum Technol.* **2024**, *11*, 72. [[CrossRef](#)]
39. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information*; Cambridge University Press: Cambridge, UK, 2010.
40. Lidar, D.A.; Brun, T.A. *Quantum Error Correction*; Cambridge University Press: Cambridge, UK, 2013.
41. Lundberg, S.M.; Lee, S.-I. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 4768–4777.
42. Lundberg, S.M.; Erion, G.; Chen, H.; DeGrave, A.; Prutkin, J.M.; Nair, B.; Katz, R.; Himmelfarb, J.; Bansal, N.; Lee, S.-I. From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* **2020**, *2*, 56–67. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.