Q-LEAP: QUANTUM-INSPIRED LOCALIZED EVOLUTION AND ANCHOR PROJECTION FOR CROSS-SUBJECT SEMG RECOGNITION

Anonymous authors
Paper under double-blind review

ABSTRACT

Cross-subject sEMG recognition is challenging due to substantial physiological variability, electrode shifts, and non-stationary signal drift across individuals. We propose Q-LEAP, a quantum-inspired framework that integrates localized evolution and anchor projection to learn transferable sEMG representations. Q-LEAP maps multichannel signals into normalized complex-valued states in Hilbert space, models cross-channel interactions through learnable localized evolution, and coherently fuses time- and frequency-domain representations via quantum-inspired interference. It further constructs fidelity-based anchor states and projects learned representations onto these anchors through projective measurement, yielding a stable reference space for cross-subject recognition. Extensive experiments on five NinaPro benchmarks demonstrate that Q-LEAP matches or exceeds twelve representative baselines in cross-session recognition, while providing substantial zero-shot and few-shot generalization advantages on unseen subjects and robust cross-database transfer between intact and amputee populations. These findings establish the effectiveness of quantum-inspired localized dynamics and fidelity geometry for robust physiological signal decoding under cross-subject variability.

1 INTRODUCTION

Surface electromyography (sEMG) serves as a non-invasive neurophysiological interface connecting the human neuromuscular system with external intelligent devices, demonstrating significant application value in bionic prosthesis control (Huang et al., 2022), neural rehabilitation (Burrello et al., 2022; Weng et al., 2025), and human-computer interaction (Guo et al., 2024; Kaifosh & Reardon, 2025; Huang et al., 2026).

Although CNNs (Cao et al., 2019; Wan et al., 2026), RNNs (Song et al., 2023; Li et al., 2025), and Transformers (Cui et al., 2026; Duan et al., 2026; Tang et al., 2026) have improved sEMG recognition, they remain poorly matched to the intrinsic structure of multichannel physiological signals. Volume conduction through heterogeneous tissues introduces channel crosstalk, amplitude attenuation, and phase shifts (Yu et al., 2025), whereas conventional real-valued models mainly capture these effects implicitly through convolution or attention. Meanwhile, muscle fatigue induces non-stationary temporal and spectral drift (Billardello et al., 2026), further destabilizing learned representations. Across subjects, physiological variability introduces additional domain shifts. Existing adaptation methods, including GANs (Chen et al., 2026; Wei et al., 2025) and semi-supervised learning (Du et al., 2017; Wang et al., 2023), mainly align absolute feature distributions without explicitly preserving intrinsic amplitude–phase relations or structural geometry. Therefore, robust cross-subject sEMG recognition requires a representation that can jointly encode phase-aware channel interactions, constrain evolution under crosstalk, and establish a subject-stable reference space.

To overcome these limitations, we propose **Quantum-Inspired Localized Evolution and Anchor Projection (Q-LEAP)** for cross-subject sEMG recognition. Q-LEAP represents multi-channel sEMG as complex-valued states in Hilbert space and introduces localized evolution and coherent interference to model structured cross-channel interactions under physiological perturbations. It further establishes a fidelity-based anchor reference space and adopts a decoupled adaptation strat-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

egy, enabling lightweight cross-subject calibration without repeatedly fitting subject-specific feature distributions. The core contributions are threefold.

- **Physics-inspired quantum representation.** We develop a complex-valued representation that jointly encodes amplitude and phase information, and introduce disordered localized evolution and coherent time-frequency interference to constrain cross-channel propagation and attenuate perturbations associated with volume conduction and signal non-stationarity.
- **Fidelity-anchor-based decoupled adaptation.** We construct a geometrically dispersed reference space through quantum-fidelity-based anchor sampling and represent learned states by their anchor-relative overlaps. By freezing the representation backbone and fine-tuning only a lightweight classifier, Q-LEAP enables efficient adaptation while reducing dependence on subject-specific absolute feature distributions.
- **Comprehensive generalization validation.** We evaluate Q-LEAP on five NinaPro benchmarks under cross-session, zero-shot and few-shot cross-subject, and cross-database transfer settings, including intact-to-amputee adaptation. Ablation, robustness, and efficiency analyses further assess the contributions and practical properties of the proposed components.

2 PRELIMINARIES

2.1 QUANTUM STATE

A pure quantum state is represented by a unit vector in a complex Hilbert space and is conventionally written in Dirac notation as a *ket*, $|\Psi\rangle$. Its conjugate transpose is the corresponding *bra*, denoted by $\langle\Psi| = (|\Psi\rangle)^\dagger$. Given an orthonormal basis $\{|i\rangle\}_{i=0}^{n-1}$, a pure state can be expressed as

$$|\Psi\rangle = \sum_{i=0}^{n-1} \alpha_i |i\rangle, \qquad \sum_{i=0}^{n-1} |\alpha_i|^2 = 1, \tag{1}$$

where $\alpha_i \in \mathbb{C}$ denotes the complex amplitude associated with the basis state $|i\rangle$. The normalization condition ensures that $|\Psi\rangle$ has unit norm in Hilbert space.

2.2 QUANTUM MEASUREMENT

For a normalized quantum state $|\Psi\rangle$, a rank-one projective measurement associated with an orthonormal state $|\Phi_k\rangle$ is described by the projector $\Pi_k = |\Phi_k\rangle\langle\Phi_k|$. According to the Born rule (Gleason, 1957), the probability of obtaining the corresponding measurement outcome is

$$P_k = \langle\Psi|\Pi_k|\Psi\rangle = |\langle\Phi_k|\Psi\rangle|^2. \tag{2}$$

For a complete set of orthogonal projectors satisfying $\sum_k \Pi_k = I$, the resulting probabilities satisfy $\sum_k P_k = 1$. This formulation provides a principled mapping from complex-valued quantum states to real-valued measurement probabilities.

2.3 QUANTUM FIDELITY

Quantum fidelity quantifies the similarity between quantum states (Nielsen et al., 2000). For two normalized pure states $|\Psi\rangle$ and $|\Phi\rangle$, it reduces to

$$F(|\Psi\rangle, |\Phi\rangle) = |\langle\Phi|\Psi\rangle|^2, \qquad 0 \le F \le 1. \tag{3}$$

A fidelity of one indicates that the two states are identical up to a global phase, whereas a fidelity of zero indicates orthogonality. Unlike Euclidean distance between complex state vectors, fidelity is invariant to arbitrary global phase transformations, e.g.,

$$F(e^{i\alpha}|\Psi\rangle, |\Phi\rangle) = F(|\Psi\rangle, |\Phi\rangle). \tag{4}$$

This phase invariance makes fidelity a natural geometric similarity measure for normalized complex-valued representations.

3 Related Work

Existing sEMG recognition methods primarily focus on spatial-temporal modeling and cross-domain generalization. Spatial-temporal, Transformer, and graph models (Wei et al., 2019; Rahimian et al., 2020; Fauvel et al., 2021; Karnam et al., 2022; Dai et al., 2023; Ni & Al-Qaness, 2026; Wang et al., 2024; Cha et al., 2026; Lee et al., 2023) effectively capture temporal and inter-channel dependencies, but largely operate in real-valued spaces without explicitly modeling amplitude–phase relations or structured propagation under channel crosstalk. Cross-domain methods (Chen et al., 2025; Lee & Seo, 2026; Tekin, 2026) mainly align subject-dependent feature distributions, while existing quantum-inspired approaches (Chen et al., 2021; Olivia et al., 2026) do not jointly exploit complex representations, localized evolution, and fidelity-based adaptation, motivating Q-LEAP for robust cross-subject recognition.

4 Methodology

4.1 Problem Definition

For a multi-channel sEMG sequence $X \in \mathbb{R}^{T\times C}$ with T time steps and C spatial channels, the quantum-inspired network f_θ learns a mapping $\hat{Y} = f_\theta(X)$ to predict the gesture label Y, where $\hat{Y}$ denotes the output category probability distribution. In cross-subject recognition, the objective is to learn discriminative representations that remain robust to subject-dependent physiological and acquisition variations.

4.2 Motivation

Cross-subject sEMG recognition is challenged by volume-conduction-induced channel crosstalk, electrode shifts, muscle fatigue, and physiological variability, which jointly perturb spatial, temporal, and phase characteristics. Conventional real-valued models mainly capture these dependencies implicitly, making it difficult to preserve stable structural relations across subjects. To address this limitation, we map multi-channel sEMG signals into complex-valued states in Hilbert space, employ localization-inspired unitary evolution and coherent time-frequency interference to model structured channel interactions and suppress perturbations, and use quantum fidelity to construct an anchor-relative representation that reduces dependence on subject-specific absolute feature distributions.

4.3 Overall Framework and Training Paradigm

Figure 1 illustrates the proposed quantum-inspired architecture. To simultaneously support standard supervised learning, cross-subject adaptation and zero-shot inference within a unified architecture, the structural modules and the training paradigm are deeply integrated into a decoupled multi-stage optimization process.

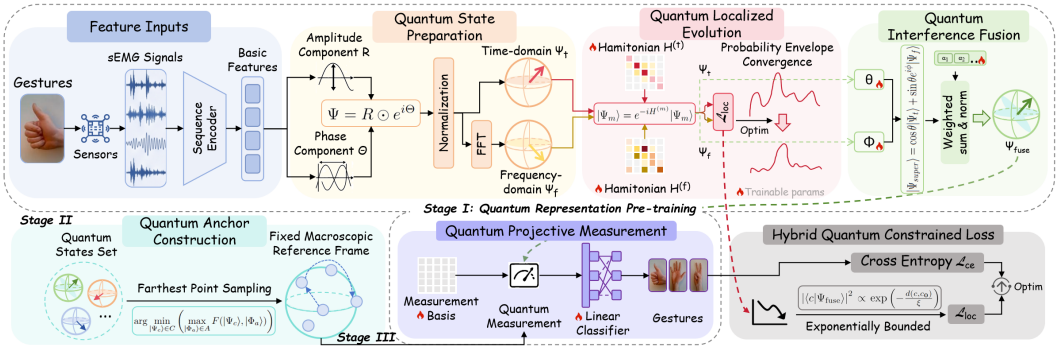


Figure 1: The overall architecture and decoupled multi-stage training paradigm of Q-LEAP.

- **Quantum Representation Pre-training.** This stage drives the quantum state preparation, quantum localized evolution, and quantum interference fusion to achieve convergence.

- **Quantum Anchor Construction.** Upon achieving feature convergence, this stage adopts quantum fidelity as a geometric metric and invokes the farthest point sampling algorithm to sample anchor states with maximized pairwise separation, thereby establishing a fixed macroscopic measurement reference frame.
- **Quantum Representation Fine-tuning.** This stage executes quantum projective measurement, where the high-dimensional complex wave function is collapsed onto the constructed anchors to output a real-valued probability distribution for discrete gesture classification, where the optimizer solely updates the lightweight linear classification layer.

Table 1: Paradigm comparison with Q-LEAP and conventional practice.

Quantum Mechanism	Conventional Practice	Classical Limitation	Quantum Advantage
Localized Evolution	Channel Attention	Lacks spatial physical priors; diffuse crosstalk.	Physically confines spatial crosstalk.
Coherent Interference	Gated Feature Fusion	Magnitude-only scaling; cannot cancel noise.	Exploits relative phase for destructive noise cancellation.
Fidelity Anchor	Euclidean / Cosine Metric	Highly sensitive to electrode phase drift.	Mitigates signal shifts from electrode displacement.
Quantum Measurement	Linear Classifier	Overfits to absolute individual distributions.	Fixed-anchor projection enables rapid few-shot calibration.

4.4 Quantum Representation Pre-training

4.4.1 Quantum State Preparation

The system first encodes multi-channel sEMG signals using a 1D-CNN. Inspired by complex Morlet wavelet analysis (Ward et al., 2021; Stock & Anderson, 2022), two parallel convolutional branches then learn the magnitude $R \geq 0$ and phase Θ, forming the time-domain complex state

$$\Psi_{raw} = R \odot e^{i\Theta}, \tag{5}$$

where $\odot$ denotes the Hadamard product and i is the imaginary unit. The state is normalized as $\Psi_t = \Psi_{raw}/(\|\Psi_{raw}\|_2 + \epsilon)$ to obtain a unit-norm representation in Hilbert space. Finally, an orthonormally normalized FFT transforms Ψ_t into the frequency-domain state Ψ_f, providing complementary time- and frequency-domain representations for subsequent evolution.

4.4.2 Quantum Localized Evolution

Following state initialization, we introduce quantum disordered localization theory (Sanchez-Palencia & Lewenstein, 2010) to construct a cross-channel evolution operator, aiming to model the spatial scattering of wave functions in complex electrophysiological media. This operation is designed to address physical crosstalk and spatial phase shifts among channels caused by volume conduction through heterogeneous biological tissues. We treat variations in skin contact impedance and electrode shifts as internal disordered environments and introduce a learnable Hamiltonian operator, parameterized as a Hermitian matrix $H^{(m)} = W^{(m)} + (W^{(m)})^\dagger$ with complex parameters $W^{(m)}$, to perform spatial scattering evolution on both time and frequency domain, denoted by Eq. 6.

$$|\Psi_m\rangle = e^{-iH^{(m)}}|\Psi_m\rangle = \sum_{j,k}[e^{-iH^{(m)}}]_{j,k}\langle k|\Psi_m\rangle|j\rangle, \tag{6}$$

where $m \in \{t, f\}$ distinguishes the temporal and frequency branches, Hamiltonian satisfies $[e^{-iH^{(m)}}]^\dagger e^{-iH^{(m)}} = I$, while k and j represent the spatial basis indices of the input and output tensors respectively. In addition, to guide this scattering process in accordance with quantum localization principles, optimization of the Hamiltonian H is explicitly constrained by the hybrid quantum constrained loss, which encourages the probability envelope to converge towards the activated core muscle groups, promoting efficient channel denoising, detailed in **Section** 4.6.2.

4.4.3 Quantum Interference Fusion

To dynamically fuse time and frequency representations, the system performs coherent interference fusion according to Eq. 7.

$$|\Psi_{super}\rangle = \cos\theta|\Psi_t\rangle + \sin\theta e^{i\phi}|\Psi_f\rangle, \tag{7}$$

where $|\Psi_{super}\rangle$ represents the superposed intermediate state with learnable parameters θ and ϕ, which is theoretically designed to encourage constructive interference for consistent action features while potentially mitigating random background noise via destructive interference.

The model subsequently aggregates the sequential features to extract a unified global representation for discrete gesture recognition. The model then performs a weighted sum operation on $|\Psi_{super}\rangle$ and normalizes the result according to Eq. 8.

$$|\Psi_{fuse}\rangle = \frac{\sum_\tau \alpha_\tau |\Psi_{super}^{(\tau)}\rangle}{\| \sum_\tau \alpha_\tau |\Psi_{super}^{(\tau)}\rangle \|_2 + \epsilon}, \tag{8}$$

where α_τ represents non-negative learnable weights to evaluate the contribution of each slice τ. The constant ϵ prevents zero division and $|\Psi_{fuse}\rangle$ provides the final discriminative state space for subsequent quantum measurements.

4.5 Quantum Anchor Construction

Before anchor construction, the network employs learnable reference states to classify the fused representations $|\Psi_{fuse}\rangle$, allowing the underlying representation parameters to converge through backpropagation. After pre-training, the network is frozen and used to extract fused states from the entire training set. Quantum anchors are then selected by farthest point sampling (FPS) (Eldar et al., 1997) using fidelity-based dissimilarity. For two normalized states $|\Psi\rangle$ and $|\Phi\rangle$, quantum fidelity $F(|\Psi\rangle, |\Phi\rangle) = |\langle\Phi|\Psi\rangle|^2$ measures their geometric overlap in Hilbert space, and $1 - F$ is used to quantify their separation. Unlike Euclidean distance between complex state vectors, fidelity is invariant to global phase, making it a more suitable similarity measure for the learned complex-valued representations.

Specifically, assuming A is the currently established set of fixed anchors, the algorithm adopts a greedy farthest point sampling strategy to avoid the quadratic computational overhead associated with constructing a massive global correlation matrix. For any unselected state $|\Psi_c\rangle$ in the candidate set C, the system calculates its fidelity with all anchors in A and extracts the maximum value as the shortest physical projection distance from this state to the entire reference frame. The algorithm then selects the quantum state with the greatest shortest distance, meaning the minimum maximum fidelity value, from the candidate set as the new quantum anchor. The mathematical expression is defined as

$$|\Phi_{new}\rangle = \arg \min_{|\Psi_c\rangle \in C} \left(\max_{|\Phi_a\rangle \in A} F(|\Psi_c\rangle, |\Phi_a\rangle) \right) \tag{9}$$

This iterative sampling mechanism reduces the overall time complexity from a quadratic order $\mathcal{O}(N^2)$ to $\mathcal{O}(N \times K)$, where N represents the total number of samples and K denotes the predefined number of anchors. By iteratively minimizing the maximum pairwise fidelity, this mechanism ensures the extraction of well-dispersed anchor states with maximized geometric separation on the Hilbert sphere, establishing a robust macroscopic reference frame for subsequent projection.

4.6 Quantum Representation Fine-tuning

4.6.1 Quantum Projective Measurement

To map high-dimensional complex states into real-valued features for gesture classification, we introduce a fidelity-based anchor projection module. Given reference states $\{|\Phi_k\rangle\}_{k=1}^K$, the fused state $|\Psi_{fuse}\rangle$ is represented by

$$P_k = |\langle\Phi_k|\Psi_{fuse}\rangle|^2, \tag{10}$$

where P_k measures the overlap with the k-th reference state. Since the anchors are not required to form an orthonormal basis, $\mathbf{P} = [P_1, \ldots, P_K]^\top$ is treated as an anchor-relative similarity representation rather than a normalized measurement probability distribution.

During pre-training, learnable reference states provide classification supervision. After convergence, they are replaced by fixed fidelity-based anchors. During fine-tuning, the representation backbone and anchors remain frozen, while only the linear classifier is updated, enabling lightweight adaptation while preserving the learned representation geometry.

4.6.2 HYBRID QUANTUM CONSTRAINED LOSS FUNCTION

To prevent unconstrained spatial feature dispersion and encourage energy concentration around dominant muscle activation sites, we formulate a hybrid objective:

$$\mathcal{L}_{total} = \mathcal{L}_{ce} + \lambda \mathcal{L}_{loc}, \tag{11}$$

where $\mathcal{L}_{ce}$ denotes cross-entropy loss and λ balances classification and spatial regularization. Inspired by the exponential spatial decay in disordered localization systems, $\mathcal{L}_{loc}$ regularizes the fused quantum state $|\Psi_{fuse}\rangle$ toward an exponentially bounded spatial envelope over the Hilbert space:

$$\mathcal{L}_{loc} = \sum_{c=1}^{C} \left(|\langle c|\Psi_{fuse}\rangle|^2 - A \exp\left(-\frac{d(c, c_0)}{\xi}\right)\right)^2, \tag{12}$$

where c indexes the spatial basis states of the latent continuous Hilbert manifold and $C =$ Hilbert_dim. The dynamic activation center $c_0 = \sum_{c=1}^{C} c \cdot |\langle c|\Psi_{fuse}\rangle|^2$. The spatial distance is defined as the Euclidean distance $d(c, c_0) = |c - c_0|$. ξ is the localization length controlling the spatial decay rate, and $A = \sum_c |\langle c|\Psi_{fuse}\rangle|^2 / \sum_c \exp(-d/\xi)$ is an energy-conserving scaling coefficient. Under this constraint, $\mathcal{L}_{loc}$ guides feature energy toward core active muscle groups while physically suppressing background noise and channel crosstalk.

Appendix A.7 and A.8 detail the complete computational workflow and tensor dimension tracking.

5 EXPERIMENT

5.1 DATASETS

Table 2 outlines the data acquisition parameters and evaluation protocols for five public sEMG datasets: NinaPro DB1 (Atzori et al., 2012), NinaPro DB2-3 (Atzori et al., 2014), and NinaPro DB4-5 (Pizzolato et al., 2017). All data splittings adhere to the original configurations, ensuring consistency with baselines. Participants were seated in a neutral position and executed specific hand and wrist gestures according to visual prompts, detailed in **Appendix** A.1. To prevent muscle fatigue, each gesture involved a five-second contraction followed by a brief rest interval.

Table 2: Summary of data characteristics, sensor configurations and benchmark evaluation protocols for sEMG databases NinaPro DB1 to NinaPro DB5.

Dataset	Subjects	Sessions	Gestures	Channels	Normality	Sampling Rate	Sensors Type	Train Trials	Test Trials
NinaPro DB1	27	10	53	10	Able-bodied	100 Hz	Otto Bock MyoBock 13E200	1, 3, 4, 6, 8, 9	2, 5, 7, 10
NinaPro DB2	40	10	50	12	Able-bodied	2000Hz	Delsys Trigno Wireless	1, 3, 4, 6	2, 5
NinaPro DB3	11	10	50	12	Amputee	2000Hz	Delsys Trigno Wireless	1, 3, 4, 6	2, 5
NinaPro DB4	10	6	53	12	Able-bodied	2000Hz	Cometa Wave Plus	1, 3, 4, 6	2, 5
NinaPro DB5	10	6	53	16	Able-bodied	200 Hz	Thalmic Labs Myo Armband	1, 3, 4, 6	2, 5

5.2 COMPARISON BASELINES

To evaluate the proposed architecture, we select representative baselines classified into four main paradigms. These encompass spatial-temporal feature mapping methods including MV-CNN (Wei et al., 2019), XceptionTime (Rahimian et al., 2020), XCM (Fauvel et al., 2021), EMGhandNet (Karnam et al., 2022), sEMGXCM (Dai et al., 2023), and ECNN-CS (Ni & Al-Qaness, 2026), global sequence modeling approaches including TDCT (Wang et al., 2024) and AT-Transformer (Cha et al., 2026), cross-domain adaptive networks including I2CNet-DLS (Chen et al., 2025), CLDF (Lee & Seo, 2026), and SMR (Tekin, 2026), along with QED (Olivia et al., 2026) serving as a comparative quantum-inspired approach.

5.3 EXPERIMENT SETTINGS

The experiments were conducted using PyTorch with Python 3.11.4 on an NVIDIA 4090 GPU. The optimal hyper-parameters are determined through grid search on the validation set, where the $Hilbert_dim \in \{32, 64, 128, 256, 512\}$, $\xi \in \{0.5, 1.0, 1.5, 2.0, 3.0\}$, $\lambda \in \{0.01, 0.05, 0.1, 0.2, 0.5\}$. $Quantum_anchors$ is dynamically defined as integer multiples of gesture categories (N), i.e. $\{N, 2N, 3N, 5N, 10N\}$. The final configurations are shown in Table 3.

Table 3: Hyper-parameter settings.

Hyperparameter	Value	Hyperparameter	Value	Hyperparameter	Value
Batch_size	64	Dropout	0.15	Seq_embedder	1D-CNN
Learning_rate	1e-3	Random_seeds	{3407,42,2026}	Quantum_anchors	$3N$
Train_epoch	100	Optimizer	Adam	Hilbert_dim	256
Finetune_epochs	20	Window Size	200ms	Localization_length ξ	1.5
Early_stopping	5	Step Length	10ms	Localization_weight λ	0.2

5.4 Main Results

Table 4 presents the classification accuracies of the proposed model and twelve representative baselines across five public datasets under a rigorous cross-session evaluation protocol, wherein the physical chronological sequence of the training and testing data is strictly isolated. The competitive performance and low standard deviations substantiate the model's superior generalization capability and stability. **Appendix** A.5 illustrates the individual subject accuracy distributions and **Appendix** A.6 shows the statistical significance test.

Table 4: Performance comparison on Ninapro DB1-5 under the cross-session evaluation protocol. The **best** and second-best are highlighted. The results of Q-LEAP are averaged over 3 runs with standard deviation.

Year	2019	2020	2021	2022	2023	2024	2025	2026					Q-LEAP
Methods	MV-CNN	Xception-Time	XCM	EMGhand-Net	sEMGXCM	TDCT	I2CNet-DLS	CLDF	SMR	ECNN-CS	AT-Transformer	QED	AVG±STD
DB1	88.20%	85%	90.50%	91.76%	91.40%	87.39%	94.90%	94.50%	88.30%	-	95.28%	75.45%	**95.69% ± 0.41%**
DB2	83.70%	83.40%	84.80%	92.01%	86.30%	81.32%	-	86.30%	-	-	**93.48%**	72.38%	92.48% ± 2.4%
DB3	64.30%	55%	65%	-	66.50%	-	-	61.50%	-	-	-	48.94%	**66.56% ± 1.12%**
DB4	54.30%	71.70%	78.10%	89.50%	78.70%	-	-	89.40%	87.60%	82.32%	**94.97%**	55.91%	94.75% ± 0.82%
DB5	90%	89%	94%	-	94.20%	72.83%	-	96.30%	-	88.13%	-	79.29%	**96.36% ± 0.41%**

'–' indicates that the respective baseline did not evaluate on that dataset.

6 Further Analysis

6.1 Ablation Analysis

Table 5 details ablation and classical substitution experiments to validate the core quantum components. Results indicate that replacing the quantum localized evolution or interference fusion with traditional attention and aggregation strategies incurs performance degradations of 3.10% and 2.57%, respectively. Furthermore, replacing quantum anchors with Euclidean distance or K-Means clustering lowers accuracy by 3.65% and 2.84%. These controlled substitutions confirm that each quantum-inspired formulation provides measurable advantages over its classical counterpart.

6.2 Cross-Subject Generalization

Table 6 evaluates cross-subject generalization using leave-one-out strategy to simulate the clinical scenarios involving unseen target subjects. Specifically, the system samples 1 to 20 training instances for each gesture in the few-shot. Furthermore, an untrained baseline is introduced to demonstrate the performance of random initialization. Results demonstrate that Q-LEAP surpasses comparative baselines in zero-shot (e.g., 26.4% vs. 7.7% on DB1 and 31.6% vs. 14.5% on DB5), indicating that quantum anchors effectively capture a shared feature manifold independent of individual physiological differences. Furthermore, during few-shot adaptation, the model fits the target distribution with minimal fine-tuning overhead to ensure highly stable cross-subject deployment.

6.3 Cross-Database Generalization

To evaluate the domain adaptation capability under extreme physiological heterogeneity, cross-database generalization are conducted on five basic gestures shared between Ninapro DB2 and DB3, which feature identical electrode layouts and recording protocols. Strictly adhering to baselines, the model is pre-trained on data from ten randomly selected source-domain subjects before being evaluated on ten target-domain participants. Results in Table 7 show that transferring from intact

Table 5: Ablation study and classical replacement results. Gray shading indicates the substitution with classical components. **red** and blue highlights denote the most and second most severe average performance degradations (Δ), respectively.

Dataset / Methods	DB1	DB2	DB3	DB4	DB5	Avg	Δ
w/o Phase Θ	93.16	89.82	62.2	92.15	93.84	86.23	-2.94
w/o Localized Evolution	92.49	88.75	61.34	90.9	92.11	85.12	-4.05
w/o Localized Evolution + w/ Channel MLP	92.55	88.93	60.62	90.91	92.05	85.01	-4.16
w/o Localized Evolution + w/ Channel Attention	93.25	89.92	61.55	92.12	93.49	86.07	-3.10
w/o Interference Fusion ($\|\Psi_t\rangle$ only)	93.21	90.15	63.48	92.46	94.27	86.7	-2.47
w/o Interference Fusion ($\|\Psi_f\rangle$ only)	91.83	87.58	60.25	90.14	92.52	84.46	**-4.71**
w/o Interference Fusion + w/ Concatenation	92.63	90.55	62.81	92.37	94.65	86.60	-2.57
w/o Interference Fusion + w/ Gated Fusion	93.95	90.8	64.13	93.05	94.19	87.22	-1.95
w/o Quantum Constraint $\mathcal{L}_{loc}$	93.81	90.65	63.76	92.82	94.3	87.07	-2.10
w/o Quantum Constraint $\mathcal{L}_{loc}$ + w/ L1-norm	93.75	90.66	62.58	92.53	94.33	86.77	-2.40
w/o Quantum Anchor	93.45	90.22	60.84	92.21	93.72	86.09	-3.08
w/o Quantum Anchor + w/ Euclidean Distance	92.31	90.15	59.37	92.2	93.56	85.52	-3.65
w/o Quantum Anchor + w/ K-Means Clustering	94.05	89.74	61.23	92.59	94.05	86.33	-2.84
Q-LEAP (Full)	95.69	92.48	66.56	94.75	96.36	89.17	/

Table 6: Performance comparison of cross-subject generalization under zero-shot and few-shot settings across Ninapro DB1-DB5 datasets.

Strategy	Untrained	Zero-shot		Few-shot									
Dataset/ Methods	Q-LEAP	CLDF	Q-LEAP	1-shot		2-shot		5-shot		10-shot		20-shot	
				CLDF	Q-LEAP	CLDF	Q-LEAP	CLDF	Q-LEAP	CLDF	Q-LEAP	CLDF	Q-LEAP
DB1	6.0%	7.7%	26.4%	37.0%	46.5%	45.5%	53.2%	60.4%	72.8%	72.9%	77.4%	77.8%	87.4%
DB2	3.7%	4.4%	15.8%	28.4%	38.1%	30.4%	45.5%	53.7%	67.1%	63.2%	72.2%	71.0%	83.3%
DB3	2.5%	3.6%	12.3%	24.7%	28.0%	26.5%	32.3%	30.5%	38.6%	38.6%	42.2%	46.5%	50.1%
DB4	9.8%	13.4%	29.5%	33.8%	42.4%	39.5%	54.7%	52.3%	63.2%	62.6%	77.5%	68.3%	79.9%
DB5	12.2%	14.5%	31.6%	34.0%	47.3%	45.3%	56.7%	59.0%	65.0%	69.4%	75.3%	77.5%	81.6%

subjects to amputees (DB2 $\rightarrow$ DB3) yields steady performance gains with increasing target supervision, where Q-LEAP outperforms CLDF across all sample budgets while reducing cross-subject standard deviation to 3.04% (vs. 8.42% for CLDF). Furthermore, despite performance degradation in the reverse transfer (DB3 $\rightarrow$ DB2) caused by altered muscle topologies, the proposed architecture maintains superior robustness, confirming that quantum anchors are reliable to address severe non-stationary domain shifts.

6.4 Robustness Analysis

To evaluate the interference resistance capability of the model under harsh environments, this experiment artificially injects additive white Gaussian noise. Figure 2a demonstrates that under extreme noise interference, the quantum interference mechanism effectively prevents the catastrophic performance collapse observed in classical fusion networks through interference cancellation. **Appendix** A.2.1 shows the effect of sliding windows to verify the robustness against information loss.

Table 7: Cross-database generalization performance in a 5-way classification setting across varying few-shot configurations. DB2 contains intact subjects and DB3 contains amputee subjects.

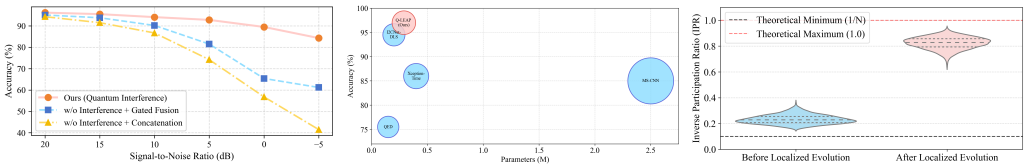
Transfer Direction	Metric	1-shot		2-shot		5-shot		10-shot		20-shot	
		CLDF	Q-LEAP	CLDF	Q-LEAP	CLDF	Q-LEAP	CLDF	Q-LEAP	CLDF	Q-LEAP
DB2 $\rightarrow$ DB3 (Intact-to-Amputee)	Avg	47.84%	54.56%	58.07%	62.15%	64.23%	76.38%	85.60%	88.67%	85.82%	89.34%
	Std	11.73%	9.85%	10.71%	9.28%	13.46%	6.17%	7.18%	4.63%	8.42%	3.04%
DB3 $\rightarrow$ DB2 (Amputee-to-Intact)	Avg	21.10%	25.66%	23.29%	32.53%	24.02%	35.16%	27.36%	41.59%	45.61%	53.25%
	Std	0.94%	0.54%	1.58%	1.38%	5.78%	5.38%	5.20%	5.71%	9.33%	8.01%

6.5 Computational Efficiency Analysis

Figure 2b compares parameter scale against accuracy across baselines. Q-LEAP achieves top-tier accuracy while maintaining a compact footprint of 301.57K parameters and 193.25M FLOPs. **Appendix** A.3 details the multi-hardware and component-wise execution latency.

6.6 Spatial Localization Analysis

To evaluate the spatial localization degree of the quantum state, we calculate the Inverse Participation Ratio (IPR), according to $IPR = \sum_{j=1}^{N} p_j^2 = \sum_{j=1}^{N} |c_j|^4$, where N is the spatial channels, c_j denotes the amplitude of the quantum state at the j-th channel, p_j represents the channel probability. The lower bound $\frac{1}{N}$ corresponds to a diffuse energy state, while the theoretical maximum of 1.0 indicates a highly localized condensation state. The climb of IPR, shown in Figure 2c, confirms that this mechanism can guide the signal to adaptively focus on the core activating muscle groups. **Appendix** A.4 visualizes how this physical evolution affects the feature manifold.

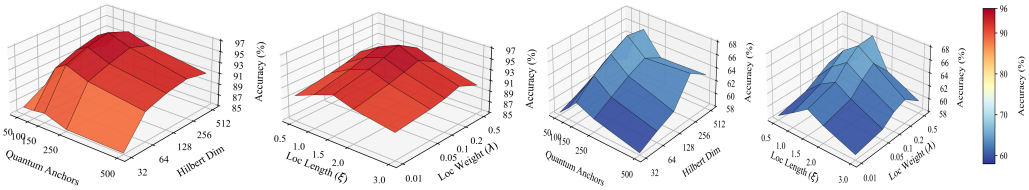


(a) (b) (c)

Figure 2: (a) Robustness analysis under varying signal to noise ratios, (b) the efficiency analysis of parameter scale and accuracy, and (c) the spatial localization degree before and after evolution on DB1.

6.7 Hyper-parameter Sensitivity Analysis

Figure 3 illustrates the sensitivity analysis of core quantum-inspired hyperparameters on the Ninapro DB1 (a, b) and DB3 (c, d) datasets. Experimental results indicate that the network achieves its optimal representation convergence peak at essentially similar hyperparameter values across both intact subject and amputee domains, which validates the generalization capability of the proposed architecture under heterogeneous physiological conditions.



(a) (b) (c) (d)

Figure 3: Joint hyperparameter sensitivity of core quantum-inspired parameters.

7 Conclusion

We proposed Q-LEAP, a quantum-inspired framework for robust sEMG recognition under non-stationarity and cross-subject variability. Q-LEAP represents multi-channel sEMG as complex-valued states in Hilbert space and integrates localization-inspired evolution, coherent time-frequency interference and fidelity-based anchor projection to learn structured and subject-transferable representations. Experiments on five NinaPro benchmarks demonstrate consistent performance across cross-session, zero-shot and few-shot cross-subject, and cross-database transfer settings.

Several limitations remain. Q-LEAP currently focuses on discrete gesture classification and has not been evaluated for continuous kinematic or torque regression, online closed-loop control, or long-term physiological variations. Future work will extend the framework to continuous decoding and assess its robustness in real-time and longitudinal deployment settings.

AI USE STATEMENT

In this work, we used generative AI tools for language polishing and grammatical editing to improve readability. We have not used generative AI tools for any tasks requiring substantive disclosure, including research ideation, mathematical derivations, data collection, or experimental analysis. We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

An anonymized demonstration implementation is provided in the supplementary material. The complete source code will be made publicly available upon acceptance. The datasets used in this study are publicly available. NinaPro DB1-5 are available at `https://ninapro.hevs.ch/`.

REFERENCES

Manfredo Atzori, Arjan Gijsberts, Simone Heynen, Anne-Gabrielle Mittaz Hager, Olivier Deriaz, Patrick Van Der Smagt, Claudio Castellini, Barbara Caputo, and Henning Müller. Building the ninapro database: A resource for the biorobotics community. In *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pp. 1258–1265. IEEE, 2012.

Manfredo Atzori, Arjan Gijsberts, Claudio Castellini, Barbara Caputo, Anne-Gabrielle Mittaz Hager, Simone Elsig, Giorgio Giatsidis, Franco Bassetto, and Henning Müller. Electromyography data for non-invasive naturally-controlled robotic hand prostheses. *Scientific data*, 1(1): 140053, 2014.

Roberto Billardello, Christian Tamantini, Francesca Cordella, Francesco Scotto Di Luzio, Tiwana Varrecchia, Giorgia Chini, Francesco Draicchio, Alberto Ranavolo, and Loredana Zollo. Estimating workers' physical effort during isometric contractions through semg: The role of feature selection. *ACM Transactions on Human-Robot Interaction*, 15(1):1–21, 2026.

Alessio Burrello, Francesco Bianco Morghet, Moritz Scherer, Simone Benatti, Luca Benini, Enrico Macii, Massimo Poncino, and Daniele Jahier Pagliari. Bioformers: Embedding transformers for ultra-low power semg-based gesture recognition. In *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1443–1448. IEEE, 2022.

Xin Cao, Masami Iwase, Jun Inoue, and Eisaku Maeda. Gesture recognition based on convlstm-attention implementation of small data semg signals. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*, pp. 21–24, 2019.

Doheun Cha, Dong-Gyu Lee, and Sangtae Ahn. Adversarially regularized transformer with channel-wise noise for robust hand gesture recognition using surface electromyography. *Biomedical Engineering Letters*, 16(2):463–472, 2026.

Peiji Chen, Wenyang Li, Yifan Tang, Shunta Togo, Hiroshi Yokoi, and Yinlai Jiang. Intra-and inter-channel deep convolutional neural network with dynamic label smoothing for multichannel biosignal analysis. *Neural Networks*, 183:106960, 2025.

Xiangdian Chen, Yan Liu, Heng Jin, Tolulope Tofunmi Oyemakinde, Guanglin Li, and Xiangxin Li. Adawgan: Data augmentation for few-shot hd-semg gesture recognition using single-trial data. *IEEE Journal of Biomedical and Health Informatics*, 2026.

Yiwei Chen, Yu Pan, and Daoyi Dong. Quantum language model with entanglement embedding for question answering. *IEEE Transactions on Cybernetics*, 53(6):3467–3478, 2021.

Wenhui Cui, Christopher M Sandino, Hadi Pouransari, Ran Liu, Juri Minxha, Ellen Zippi, Erdrin Azemi, and Behrooz Mahasseni. Embridge: Enhancing gesture generalization from emg signals through cross-modal representation learning. In *International Conference on Learning Representations*, volume 2026, pp. 149668–149684, 2026.

Qingfeng Dai, Yongkang Wong, Mohan Kankanhali, Xiangdong Li, and Weidong Geng. Improved network and training scheme for cross-trial surface electromyography (semg)-based gesture recognition. *Bioengineering*, 10(9):1101, 2023.

Yu Du, Yongkang Wong, Wenguang Jin, Wentao Wei, Yu Hu, Mohan S Kankanhalli, and Weidong Geng. Semi-supervised learning for surface emg-based gesture recognition. In *IJCAI*, pp. 1624–1630, 2017.

Yahan Duan, Wenzhi Wang, Hao Jia, Wenyu Li, Yankai Yin, and Feng Duan. A hybrid fuzzy temporal transformer method for recognition of finger movements with switching states through semg signals. *IEEE Transactions on Fuzzy Systems*, 2026.

Yuval Eldar, Michael Lindenbaum, Moshe Porat, and Yehoshua Y Zeevi. The farthest point strategy for progressive image sampling. *IEEE transactions on image processing*, 6(9):1305–1315, 1997.

Kevin Fauvel, Tao Lin, Véronique Masson, Élisa Fromont, and Alexandre Termier. Xcm: An explainable convolutional neural network for multivariate time series classification. *Mathematics*, 9 (23):3137, 2021.

Andrew M Gleason. Measures on the closed subspaces of a hilbert space. In *The Logico-Algebraic Approach to Quantum Mechanics: Volume I: Historical Evolution*, pp. 123–133. Springer, 1957.

Weiyu Guo, Ying Sun, Yijie Xu, Ziyue Qiao, Yongkui Yang, and Hui Xiong. Spgesture: Source-free domain-adaptive semg-based gesture recognition with jaccard attentive spiking neural network. *Advances in Neural Information Processing Systems*, 37:36717–36747, 2024.

Jin Huang, Zhijun Li, Haisheng Xia, Guang Chen, and Qingsheng Meng. Cross-modal integration and transfer learning using fuzzy logic techniques for intelligent upper limb prosthesis. *IEEE Transactions on Fuzzy Systems*, 31(4):1267–1280, 2022.

Zhenghao Huang, Huilin Yao, Kaikai Wang, and Lin Shu. Reading your actions: Learning generalizable action representations via pre-training aemg. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20108–20117, 2026.

Patrick Kaifosh and Thomas R Reardon. A generic non-invasive neuromotor interface for human-computer interaction. *Nature*, 645(8081):702–711, 2025.

Naveen Kumar Karnam, Shiv Ram Dubey, Anish Chand Turlapaty, and Balakrishna Gokaraju. Emghandnet: A hybrid cnn and bi-lstm architecture for hand activity classification using surface emg signals. *Biocybernetics and biomedical engineering*, 42(1):325–340, 2022.

Hunmin Lee and Daehee Seo. Upper and lower-limb motor decoding for adaptive and generalized neural rehabilitation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2026.

Hyeyun Lee, Soyoung Lee, Jaeseong Kim, Heesoo Jung, Kyung Jae Yoon, Srinivas Gandla, Hogun Park, and Sunkook Kim. Stretchable array electromyography sensor with graph neural network for static and dynamic gestures recognition system. *npj Flexible Electronics*, 7(1):20, 2023.

Meiju Li, Zijun Wei, Zhi-Qiang Zhang, Shuhao Ma, and Sheng Quan Xie. Continuous joint kinematics prediction using gat-lstm framework based on muscle synergy and sparse semg. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 33:1763–1773, 2025.

Sike Ni and Mohammed AA Al-Qaness. Ecnn-cs: Efficient convolutional neural network by channel and spatial fusion for hand gesture recognition using semg sensors. *IEEE Sensors Journal*, 2026.

Michael A Nielsen, Isaac L Chuang, et al. *Quantum computation and quantum information*, volume 1. Cambridge university press Cambridge, 2000.

Giftlin Olivia, Pramod Mathew Jacob, Thejaswini Kishore, and MSP Subathra. Hand gesture recognition using quantum inspired neural networks with semg. *SN Computer Science*, 7(1):71, 2026.

Stefano Pizzolato, Luca Tagliapietra, Matteo Cognolato, Monica Reggiani, Henning Müller, and Manfredo Atzori. Comparison of six electromyography acquisition setups on hand movement classification tasks. *PloS one*, 12(10):e0186132, 2017.

Elahe Rahimian, Soheil Zabihi, Seyed Farokh Atashzar, Amir Asif, and Arash Mohammadi. Xceptiontime: independent time-window xceptiontime architecture for hand gesture classification. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1304–1308. IEEE, 2020.

Laurent Sanchez-Palencia and Maciej Lewenstein. Disordered quantum gases under control. *Nature Physics*, 6(2):87–95, 2010.

Shouliang Song, Anming Dong, Jiguo Yu, Yubing Han, and You Zhou. A multichannel cnn-gru hybrid architecture for semg gesture recognition. In *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 4132–4139. IEEE, 2023.

Jason Stock and Chuck Anderson. Trainable wavelet neural network for non-stationary signals. *arXiv preprint arXiv:2205.03355*, 2022.

Zhichuan Tang, Rui Meng, Huiling Sun, Zhixuan Cui, Long Li, and Xinrui Song. A prosthetic hand system by contralateral-collaborative control based on multi-task learning. *IEEE Journal of Biomedical and Health Informatics*, 2026.

Hazret Tekin. A pioneering channel-pair symmetry descriptor and meta-fusion architecture for robust semg gesture recognition. *Biomedical Signal Processing and Control*, 124:110661, 2026.

Xianglong Wan, Dexin Li, Dandan Fu, and Dong Wen. Dscattenemg: A lightweight semg-based hand gesture recognition model via depthwise separable convolution and multi-head attention. *IEEE Transactions on Mobile Computing*, 2026.

Kang Wang, Yiqiang Chen, Yingwei Zhang, Xiaodong Yang, and Chunyu Hu. Iterative self-training based domain adaptation for cross-user semg gesture recognition. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:2974–2987, 2023.

Zefeng Wang, Junfeng Yao, Meiyan Xu, Min Jiang, and Jinsong Su. Transformer-based network with temporal depthwise convolutions for semg recognition. *Pattern Recognition*, 145:109967, 2024.

Simon J Ward, Rabeb Layouni, Sofia Arshavsky-Graham, Ester Segal, and Sharon M Weiss. Morlet wavelet filtering and phase analysis to reduce the limit of detection for thin film optical biosensors. *ACS sensors*, 6(8):2967–2978, 2021.

Wentao Wei, Qingfeng Dai, Yongkang Wong, Yu Hu, Mohan Kankanhalli, and Weidong Geng. Surface-electromyography-based gesture recognition by multi-view deep learning. *IEEE Transactions on Biomedical Engineering*, 66(10):2964–2973, 2019.

Wentao Wei, Linyan Ren, Ming Zhou, Zhaoyi Ma, Kexin Zhao, and Xun Xu. Improving unimodal semg-based pattern recognition through multimodal generative adversarial learning. *IEEE Transactions on Instrumentation and Measurement*, 2025.

Zihan Weng, Yang Xiao, Peiyang Li, Chanlin Yi, Pouya Bashivan, Hailin Ma, Guang Yao, Yuan Lin, Fali Li, Dezhong Yao, et al. Real-time fine finger motion decoding for transradial amputees with surface electromyography. *Neural Networks*, 190:107605, 2025.

Hang Yu, Yang Gao, Huangliang Wu, and Xiaolin Ning. Synergistic voluntary muscle signal characteristic analysis and comparison of opm-mmg and semg. *IEEE Transactions on Instrumentation and Measurement*, 2025.

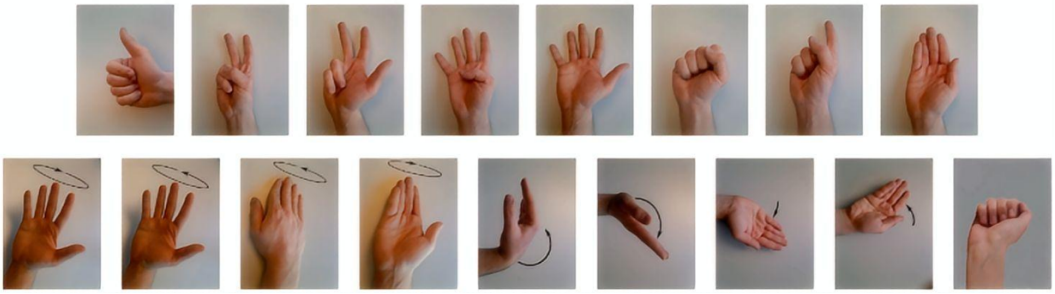
A Appendix

A.1 Dataset Visual Exemplars

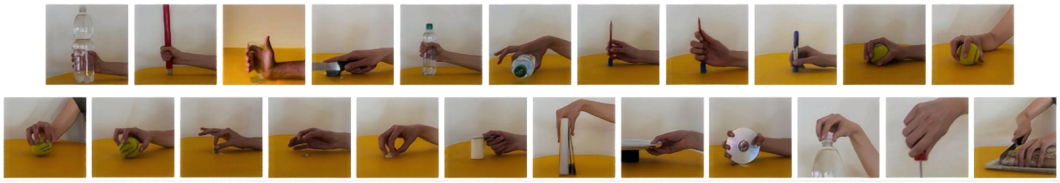
The Ninapro dataset encompasses a comprehensive suite of hand gestures designed to evaluate physiological decoding robustness across diverse clinical and daily living scenarios. The gesture vocabulary is structurally categorized into three distinct subsets to simulate varying degrees of muscular activation complexity. The first subset contains basic finger flexions to test fine-grained single-joint control. The second subset involves isometric hand configurations and wrist kinematics to assess complex multi-joint coordination. The final subset introduces functional grasping tasks interacting with daily objects to simulate real-world bionic control requirements. Details are shown in Figure 4.



(a) Basic finger movements and individual joint flexions.



(b) Isometric and isotonic hand configurations alongside basic wrist kinematics.



(c) Complex functional grasping tasks interacting with everyday objects.

Figure 4: Visual exemplars of the gesture vocabulary from the Ninapro DB5 dataset.

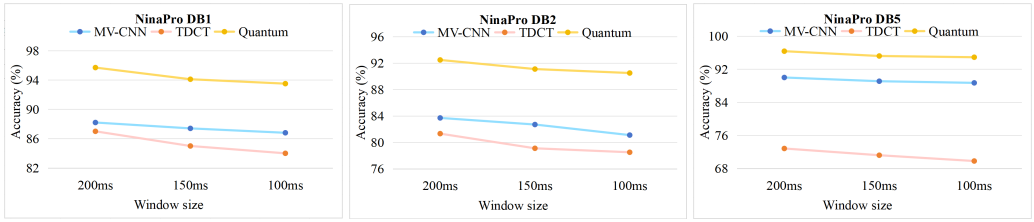
A.2 Hyper-parameter Sensitivity Analysis

A.2.1 Effect of Window Length

Figure 5 evaluates the classification accuracy across varying sliding window sizes to represent distinct hardware acquisition conditions. Experimental results indicate that while shortened windows degrade baseline performance, Q-LEAP resists information loss and maintains significantly higher stability than traditional networks like MV-CNN and TDCT under restricted data constraints.

A.3 Computational Efficiency Analysis

To rigorously evaluate the real-world deployment viability of Q-LEAP, we profile its computational footprint and execution latency under the clinically realistic single-window setting (Batch Size $= 1$, with a 200 ms analysis window and 10 ms step size). Benchmarking was conducted across two heterogeneous hardware environments: an NVIDIA RTX 4090 GPU (CUDA) and an Intel CPU (i9-13900) without hardware acceleration. To eliminate transient cold-start anomalies and dynamic thermal throttling, the model was pre-warmed for 50 iterations, followed by 300 independent inference runs using microsecond-level synchronization on the GPU and high-resolution performance counters on the CPU, reporting the mean $\pm$ std.



(a) (b) (c)

Figure 5: Classification accuracy across varying sliding windows on the Ninapro DB1, DB2, and DB5 datasets.

As detailed in Table 8, Q-LEAP maintains an exceptionally compact architectural scale with merely 301.57 K parameters and 193.25 M FLOPs. In terms of end-to-end execution, on the resource-constrained CPU, Q-LEAP sustains an execution time of 8.576 ± 1.840 ms. A fine-grained inspection of the component-wise breakdown reveals that the feature encoder constitutes the primary computational bottleneck, while quantum-inspired components collectively incur a cumulative latency of merely 0.873 ms (34.84%). This empirical profiling verifies that our quantum-inspired formulations introduce negligible computational overhead.

In the clinical deployment of bionic prostheses and neural human-computer interfaces, the total allowable controller delay between user motor intention and prosthetic actuation must remain below the 100–200 ms imperceptible threshold; exceeding 200 ms introduces severe cognitive fatigue and unnatural control lag. With our single-window inference consuming merely 3.29 ms on GPU and 8.58 ms on CPU, the computational latency of Q-LEAP occupies less than 9% of the stringent 100 ms clinical latency budget, validating the real-time operational feasibility of Q-LEAP on commodity CPU/GPU.

Table 8: Computational footprint, multi-hardware latency, and component-wise execution breakdown of Q-LEAP. Latencies are averaged over 300 runs on GPU and CPU.

Evaluation Metric / Component	Hardware Platform	Latency (ms)	Ratio / Throughput	Params	FLOPs
Component-wise Latency					
Feature Extraction	GPU	1.632 ± 0.117	65.17%	—	—
Quantum State Preparation		0.333 ± 0.042	13.29%	—	—
Quantum Localized Evolution		0.371 ± 0.040	14.80%	—	—
Quantum Interference Fusion		0.169 ± 0.031	6.75%	—	—
End-to-End Latency and Multi-Hardware Comparison					
End-to-End GPU Latency	GPU	3.291 ± 0.214	∼304 FPS	—	—
End-to-End CPU Latency	CPU	8.576 ± 1.840	∼117 FPS	—	—
Model Static Complexity Footprint					
Overall Model Footprint	—	—	—	301.57 K	193.25 M

A.4 T-SNE VISUALIZATION ANALYSIS

To objectively evaluate the feature extraction capability of Q-LEAP in processing highly confusable signals, we select thirteen actions with similar physiological muscle activation patterns from the complete 53 gesture categories for comparative analysis on NinaPro DB1. The classical feature visualization on the left side of Figure 6 indicates that sample points of these similar gestures exhibit a certain degree of overlap in the feature space. Results on the right side of Figure 6 demonstrate the practical effect of the quantum feature evolution mechanism. Following the spatial localization evolution processing, the feature points show a distinct contraction trend towards their respective category centers. This mechanism expands the inter-class distance and reduces the intra-class variance in the feature space to a certain extent, enabling the thirteen test categories to form relatively clear clusters.

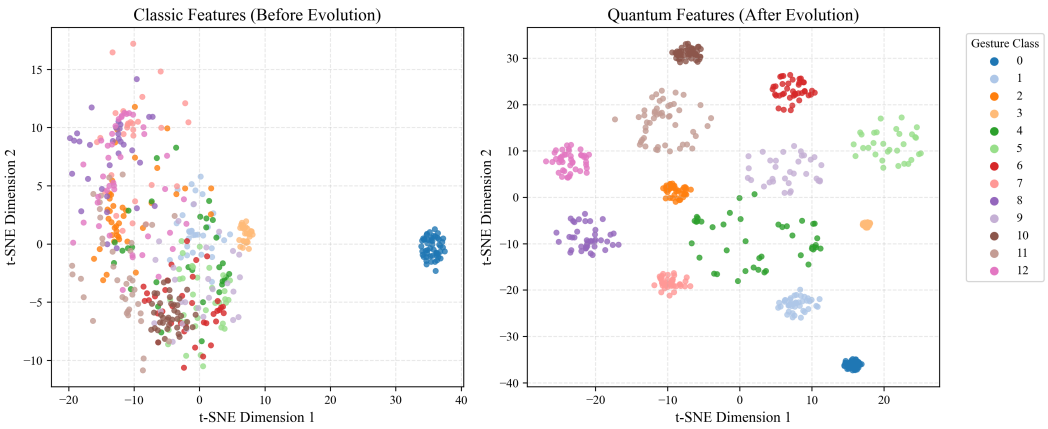


Figure 6: The t-SNE visualization of feature distributions for thirteen similar gesture categories.

A.5 Subject-Wise Performance

To verify that the high recognition rates of Q-LEAP remain robust and equitable across diverse demographic and anatomical profiles, Figure 7 illustrates the distribution of individual subject accuracies across all five public NinaPro benchmarks, with 98 subjects in total.

In intact subjects, Q-LEAP exhibits relatively compact inter-quartile ranges, with close alignment between the median indicators and mean markers suggesting that the performance spread remains comparatively symmetric. Similarly, in amputee subjects, the performance distribution remains relatively cohesive despite the inherent electrophysiological challenges associated with residual muscle variations and altered activation patterns.

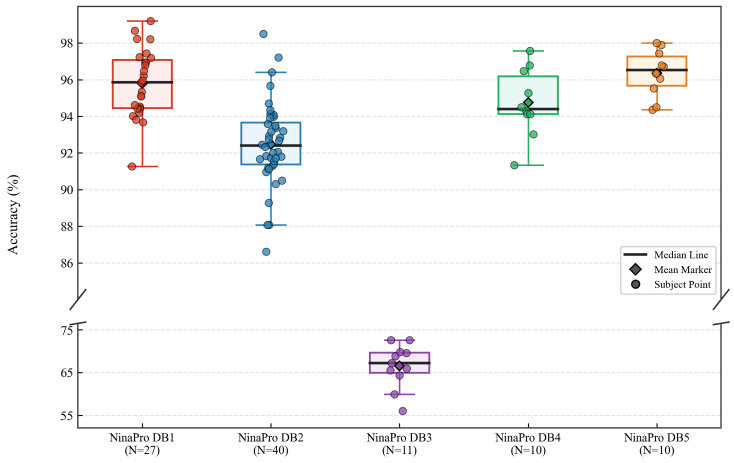


Figure 7: Subject-wise gesture recognition accuracy distributions of Q-LEAP across five benchmarks. Individual subject results are overlaid as jittered scatter points, with mean markers (diamonds) and median lines (solid black).

A.6 Statistical Significance Analysis

To confirm that the empirical improvements of Q-LEAP are statistically robust across the cohort, Table 9 presents formal significance tests on NinaPro DB1 ($N = 27$). For ablation variants, both parametric paired t-tests and non-parametric Wilcoxon signed-rank tests were conducted across all 27 participants. Remarkably, in both evaluations, the Wilcoxon test statistic confirms that Q-LEAP

achieves superior recognition across every single participant in the cohort without a single negative instance.

Furthermore, one-sample t-tests against published baseline averages confirm that Q-LEAP's accuracy significantly outperforms standard CNN models like MV-CNN ($t = 21.56, p < 0.001$) and maintains a statistically significant performance advantage over the competitive domain-adaptive baseline CLDF ($t = 3.66, p < 0.001$).

Table 9: Statistical significance evaluations of Q-LEAP on NinaPro DB1 ($N = 27$).

Evaluation Setting / Target	Comparison Target	t-test Statistic (p-value)	Wilcoxon Test Statistic (p-value)
Ablation Paired Tests	vs. w/o Localized Evolution	$t = 19.15$ ($p < 0.001$)	$W = 378.0$ ($p < 0.001$)
	vs. w/o Quantum Anchor	$t = 13.70$ ($p < 0.001$)	$W = 378.0$ ($p < 0.001$)
Baseline Superiority Tests	vs. MV-CNN	$t = 21.56$ ($p < 0.001$)	-
	vs. CLDF	$t = 3.66$ ($p < 0.001$)	-

A.7 Pseudo-code

To provide a clear and operational overview of the proposed Q-LEAP architecture, the complete algorithmic workflow is summarized in Algorithm 1. Specifically, the procedure systematically unifies three stages: (I) *Quantum Representation Pre-training*, which optimizes the underlying feature extraction, unitary localized scattering, and coherent interference via the hybrid loss; (II) *Quantum Anchor Construction*, which freezes the pre-trained dynamics and applies fidelity-based farthest point sampling on the state manifold to construct an overcomplete, well-dispersed macroscopic reference frame; and (III) *Quantum Representation Fine-tuning*, which projects incoming quantum states onto the fixed anchors via Born-rule measurement, tuning solely the top-level linear classifier with minimal computational overhead for rapid cross-subject and cross-domain adaptation.

A.8 Full-Pipeline Tensor Dimension Specification

Table 10 details the explicit tensor transformations, numerical shapes, and concise mechanisms throughout the three-stage Q-LEAP pipeline on NinaPro DB1 (10 physical electrodes, 100 Hz sampling rate, 200 ms window size yielding 20 temporal points, $C = 256$ Hilbert dimensions, and $K = 159$ anchors).

Table 10: Step-by-step tensor dimension tracking and mechanism specifications for NinaPro DB1.

Pipeline Stage / Module	Input Tensor	Output Tensor	Domain	Physical & Implementation Mechanism
Physical Sampling & Tokenization				
Physical Windowing	Raw sEMG with 10 channels	$(B, 10, 20)$	$\mathbb{R}$	Slices 20 temporal samples (200 ms $\times$ 100 Hz) across 10 physical electrodes.
Transpose	$(B, 10, 20)$	$(B, 20, 10)$	$\mathbb{R}$	Treats 20 temporal samples as channel features, preserving the 10-electrode spatial length.
Sequence Encoder	$(B, 20, 10)$	$(B, 256, 10)$	$\mathbb{R}$	1D-CNN projects 20-dim temporal waveforms into 256-dim continuous Hilbert basis states.
Quantum Representation Pre-training				
Time-domain State	$(B, 256, 10)$	$(B, 256, 10)$	$\mathbb{C}$	Decouples amplitude and phase to construct complex pure states, normalized onto the unit sphere.
Frequency-domain State	$(B, 256, 10)$	$(B, 256, 10)$	$\mathbb{C}$	Unitary FFT maps states to frequency basis.
Localized Evolution	$(B, 256, 10)$	$(B, 256, 10)$	$\mathbb{C}$	Unitary matrix operator simulates volume conduction spatial scattering.
Interference Fusion	$(B, 256, 10)$	$(B, 256, 10)$	$\mathbb{C}$	Relative phase modulation constructively enhances gesture features while canceling random noise.
Weighted Sum	$(B, 256, 10)$	$(B, 256)$	$\mathbb{C}$	Generate a global pure state $\|\Psi_{fuse}\rangle$.
Quantum Measurement	$(B, 256)$	$(B, 53)$	$\mathbb{R}$	Born probabilities are linearly mapped to logits to backpropagate gradients into underlying dynamics.
Hybrid Quantum Constrained Loss	$(B, 256)$	Scalar loc_loss	$\mathbb{R}$	Enforces exponential decay to curb crosstalk.
Quantum Anchor Construction				
Candidate State Extraction	$(N_{total}, 10, 20)$	$(N_{total}, 256)$	$\mathbb{C}$	Extracts pure-state features across the training set using the frozen pre-trained backbone.
Farthest Point Sampling	$(N_{total}, 256)$	$(159, 256)$	$\mathbb{C}$	Iteratively samples states maximizing quantum infidelity to construct a dispersed reference frame.
Quantum Representation Fine-tuning				
Target Subject Forwarding	$(B, 10, 20)$	$(B, 256)$	$\mathbb{C}$	Forwards target samples through the frozen backbone to extract high-dimensional pure states.
Frozen Quantum Measurement	$(B, 256)$ & $(159, 256)$	$(B, 159)$	$\mathbb{R}$	Collapses states onto anchors via Born rule.
Linear Classifier	$(B, 159)$	$(B, 53)$	$\mathbb{R}$	Fine-tunes the lightweight linear head.

Algorithm 1 Pseudo-code for Q-LEAP.

Require: Training dataset $\mathcal{D}_{train} = \{(X_n, Y_n)\}_{n=1}^{N}$ with $X \in \mathbb{R}^{T \times C}$; Target test sample X_{test}; Anchor count K; Loss hyperparameter λ; Localization length ξ.

Ensure: Predicted gesture probability distribution $\hat{Y}$; Fine-tuned linear classifier.

// Stage I: Quantum Representation Pre-training

Initialize sequence encoder, Hamiltonians $H^{(t)}, H^{(f)}$, superposition parameters θ, ϕ, temporal weights α_τ, dynamic measurement basis, and linear classifier.

for each pre-training epoch and mini-batch $(X, Y) \subset \mathcal{D}_{train}$ **do**

Extract temporal representation via sequence encoder: $H_{seq} \leftarrow \text{SequenceEncoder}(X)$

Project H_{seq} through parallel 1D-CNNs into amplitude $R \geq 0$ and phase Θ

$\Psi_{raw} \leftarrow R \odot e^{i\Theta}, \quad \Psi_t \leftarrow \frac{\Psi_{raw}}{\|\Psi_{raw}\|_2 + \epsilon}, \quad \Psi_f \leftarrow \text{FFT}(\Psi_t)$ ▷ Sec 4.4.1

$|\Psi_t\rangle \leftarrow e^{-iH^{(t)}}|\Psi_t\rangle, \quad |\Psi_f\rangle \leftarrow e^{-iH^{(f)}}|\Psi_f\rangle$ ▷ Sec 4.4.2

$|\Psi_{super}\rangle \leftarrow \cos\theta|\Psi_t\rangle + \sin\theta e^{i\phi}|\Psi_f\rangle$ ▷ Sec 4.4.3

$|\Psi_{fuse}\rangle \leftarrow \frac{\sum_\tau \alpha_\tau |\Psi_{super}^{(\tau)}\rangle}{\|\sum_\tau \alpha_\tau |\Psi_{super}^{(\tau)}\rangle\|_2 + \epsilon}$ ▷ Sec 4.4.3

Compute spatial activation center: $c_0 \leftarrow \sum_{c=1}^{C} c \cdot |\langle c|\Psi_{fuse}\rangle|^2$

$\mathcal{L}_{loc} \leftarrow \sum_c \left(|\langle c|\Psi_{fuse}\rangle|^2 - A \exp\left(-\frac{d(c,c_0)}{\xi}\right)\right)^2$ ▷ Sec 4.6.2

Predict distribution $\hat{Y}$ via dynamic measurement basis; Compute $\mathcal{L}_{total} \leftarrow \mathcal{L}_{ce}(\hat{Y}, Y) + \lambda \mathcal{L}_{loc}$

Update all Stage I underlying parameters via backpropagation on $\mathcal{L}_{total}$.

end for

// Stage II: Quantum Anchor Construction

Freeze all Stage I parameters; Extract fused states for all training samples to form candidate set $C \leftarrow \{|\Psi_{fuse}^{(n)}\rangle\}_{n=1}^{N}$.

Initialize fixed anchor set $A \leftarrow \{|\Phi_1\rangle\}$ with an arbitrary candidate state from C.

while $|A| < K$ **do**

$|\Phi_{new}\rangle \leftarrow \arg\min_{|\Psi_c\rangle \in C} \left(\max_{|\Phi_a\rangle \in A} F(|\Psi_c\rangle, |\Phi_a\rangle)\right)$ ▷ Sec 4.5, where $F(\Psi, \Phi) = |\langle\Phi|\Psi\rangle|^2$

$A \leftarrow A \cup \{|\Phi_{new}\rangle\}$ ▷ Establish fixed macroscopic reference frame $A = \{|\Phi_k\rangle\}_{k=1}^{K}$

end while

// Stage III: Quantum Representation Fine-tuning & Inference

Set A as the fixed measurement basis; Keep all underlying dynamics parameters strictly frozen.

for each calibration / fine-tuning sample (X, Y) **do**

Forward propagate X to obtain $|\Psi_{fuse}\rangle$ using frozen backbone

Compute Born-rule projection vector $P \leftarrow [P_1, P_2, \ldots, P_K]^\top$, where $P_k = |\langle\Phi_k|\Psi_{fuse}\rangle|^2$ ▷ Sec 4.6.1

Predict $\hat{Y} \leftarrow \text{LinearClassifier}(P)$; Update **only** the top-level linear classifier via $\mathcal{L}_{ce}(\hat{Y}, Y)$.

end for

Inference on unseen sample X_{test}:

Extract $|\Psi_{fuse}\rangle$, compute projection vector P_{test}, output predicted label $\hat{Y}_{test} \leftarrow \arg\max(\text{LinearClassifier}(P_{test}))$.

return $\hat{Y}_{test}$